FTTH council | fiber to the home

DOCSIS 3.0 and FTTH

For More Information Contact:
David St. John, FTTH Council
315.849.3800 / media@ftthcouncil.com

DOCSIS 3.0 and FTTH: The Essential Differences

Prepared by the Fiber-to-the-Home Council North America
June 2008

Table of Contents

Executive Summary

Amid aggressive competition between cable television companies and traditional telephone companies for voice, data and video subscribers, early evidence shows that the quality and capability of the access network greatly influences subscriber decisions on who will provide their triple-play services.

The country's major cable television providers offer their services over Hybrid Fiber-Coax (HFC) access networks, while many Telcos are deploying fiber to the home (FTTH) networks. FTTH uses all-optical transport – capable of boosting data speeds well beyond those provided by today's typical cable modem – and thereby is seeing market success competing against cable HFC systems.

In response to this, the cable companies have unveiled a new technology enhancement to their HFC systems, *DOCSIS 3.0*.

By deploying DOCSIS 3.0, cable operators can divert existing video transport capability in their systems toward increasing downstream data rates to their cable modem subscribers. In other words, DOCSIS 3.0 allows a service provider to redistribute existing capacity from video delivery to high speed data downstream. This is enabling cable operators to claim that in the future they will be delivering Internet transmission speeds of 50 Mb/s and above – which seem to be competitive with high-end data rates offered by many FTTH providers.

It is important, however, to note that there are major differences between the architectures employed by DOCSIS 3.0 and FTTH networks. And, as a result, there are key differences in the capabilities of the two types of systems and how they provide service to customers. As will be discussed at length in this paper, those differences are in brief…

- While DOCSIS 3.0 does give cable operators the ability to increase data transmission speeds over current HFC technology, the most optimistic expectations are that it can boost downstream speeds to only 15% of what FTTH can provide with current technology, and upstream speeds to only 12% of FTTH. In other words, even at its highest capability, DOCSIS 3.0 does not have the potential to deliver downstream data rates that are even close to what can be provided by FTTH systems.

- With DOCSIS 3.0, increasing data speeds comes at a price to the cable operator, who must sacrifice a significant part of video capacity – 10 standard definition video programs for every 38 Mb/s of downstream speed. FTTH operators have the capability to provide more video programs while boosting data speeds well beyond the capacity of cable TV systems, and therefore do not have to make such choices.

- With DOCSIS 3.0, upstream rates are likely to remain similar to those provided over current HFC technology, given the difficulties and trade-offs that would be required to appropriate more upstream bandwidth to data. FTTH networks are typically capable of providing symmetrical service – that is, equally high data rates on both the downstream and the upstream – without major trade-offs. This is a significant difference, given the ongoing evolution toward increased user-generated content (particularly work-at-home, peer-to-peer and online gaming traffic) among subscribers of home Internet services.

What is DOCSIS?

The cable television industry led the world in the rollout of high speed data (HSD) to consumers in the late 1990s with the introduction of standardized cable modems. Cable Television Laboratories (CableLabs), a joint venture of most cable operators, sponsors the development of cable modem standards (among others), relying on guidance from member cable TV operating companies and the expertise of equipment suppliers. Once they are developed, these standards are turned over to the Society of Cable Telecommunications Engineers (SCTE) to complete the standardization process. SCTE is a membership organization of technologists practicing in the cable TV area, and is recognized as a standards development body by the American National Standards Institute (ANSI). After the standards are created by SCTE, they are referred to ANSI for adoption as American National Standards.

The cable modem series of standards are referred to under the umbrella name *Data Over Cable Service Interface Specification* (DOCSIS). The first standard, DOCSIS 1.0, was adopted in the late 1990s, though the final and current version of the formal specification is dated 2002. It provided for basic downstream and upstream data transmission. The DOCSIS 1.1 standard was adopted a while later to add certain technical parameters that improved performance and security, and to require certain operating modes that were optional in DOCSIS 1.0. DOCSIS 2.0 was the next standard to be adopted, and added additional modulation modes intended to improve the upstream data speed in the presence of known impairments, as well as to add certain other features. DOCSIS 2.0 is considered the workhorse standard today, though some features of it have not been turned on by most cable TV operators.

The DOCSIS 3.0 specification was being worked on at a rather leisurely pace until the middle of 2006, when the cable TV industry recognized the threat posed by fiber to the home systems, with their vastly superior data capabilities as well as other greatly improved features. Realizing the urgency of the situation for cable operators, the specification was quickly completed and published. The current ANSI/SCTE-approved documents are dated 2007, with subsequent documents released by CableLabs in 2008, but not approved by SCTE and ANSI as of this writing. Current documents released to the public by CableLabs are available at www.cablemodem.com. ANSI/SCTE documents are available at www.scte.org.

DOCSIS 3.0 Enhances Cable Modem Download Speeds

This paper compares the triple-play delivery capability of cable TV systems equipped with DOCSIS 3.0 cable modems with the triple-play delivery capabilities of fiber to the home (FTTH) networks. DOCSIS systems operating over all-fiber networks – known as *Radio Frequency over Glass* (RFoG) – are included in the discussion. We show that, while DOCSIS 3.0 modems are the first cable modems capable of delivering 100 Mb/s downstream to subscribers, DOCSIS 3.0 provides less access bandwidth than FTTH systems can deliver.

DOCSIS 3.0 and other cable TV services are delivered using a technology called *hybrid fiber-coax* (HFC). With this technology, radio frequency (RF) signals from the *headend* (functionally equivalent to a central office in telephone terminology) travel most of the way to a neighborhood on fiber optic cable. Then at a *node*, they are converted to electrical RF signals for the remaining distance. All signals are modulated onto RF carriers at the headend, and remain in RF form until they reach the subscriber's set top box. Upstream signals are modulated onto RF carrier at the home, and stay in RF form all the way to the headend. Thus, the HFC node is an optical receiver in the downstream direction, and an optical transmitter in the upstream direction. It converts between RF signal on coax and RF signals modulated onto optical carriers. Different frequency bands are used in the two directions, 5-42 MHz upstream from subscribers and 54-750 MHz (sometimes higher) downstream.

In contrast, FTTH delivers the entire broadcast video spectrum over optical fiber all the way to the customer on one wavelength of light – 1550 nanometers (nm). Rather than share that bandwidth with data, FTTH provides much more upstream and downstream bandwidth for high speed data on different wavelengths ("colors") – 1310 nm upstream and 1490 nm downstream.

The capabilities of DOCSIS 3.0 come at a potentially high price for the cable system. With use of a technique called *channel bonding* (which is explained later in this document), for every 38 Mb/s of downstream speed, a cable operator loses the ability to program 10 or more video streams. The cable operator is forced to make a decision – is it better to add more downstream bandwidth, or more video? Or, put another way, how many video streams shall I sacrifice in order to increase my downstream data rate to subscribers?

Because FTTH systems have additional delivery capacity for high speed data access bandwidth, more services and higher speeds are possible. Therefore, the FTTH operator is not forced into making these types of trade-offs – rather, he can supply as much access bandwidth as needed without removing any video channels.

HFC and DOCSIS Architecture are Maturing Together

When the DOCSIS standard was developed, it was determined that approximately 43 Mb/s (mega, or million, bits per second) was the maximum amount of data that could be fit into one 6 MHz channel, as a practical matter. (The 6 MHz channel width derives form the width of a TV channel in North America, going back to the first successful commercial TV systems. See Appendix A for background on RF modulation and channels.)

After removing the overhead needed to support the DOCSIS channel itself, roughly 38 Mb/s of downstream bandwidth is left for users' data (the *payload*). The number of subscribers served by this 38 Mb/s bandwidth is normally determined by the number of subscribers served by one physical subset of the HFC network, also called a *node*. The data signals in the downstream DOCSIS channel flow to all subscribers in the node, so those subscribers share the data bandwidth.[1] A common size of a modern node is 500 homes passed. Assuming that one-third (167) of those homes take data service, the average data bandwidth per subscriber is given by simple division:

$$\frac{38\ Mb/s}{167\text{ homes}} = 0.23\, Mb/s\text{ per subscriber.}$$

This number is used below for average downstream bandwidth per subscriber. The actual number of subscribers served could be more, or it could be less. If the operator sells more cable modem subscriptions on that node, then the average bandwidth per subscriber will drop. It is possible to put more than one DOCSIS channel on the node and serve fewer subscribers with the channel. It is also possible to combine several nodes and serve all subscribers on those nodes with one DOCSIS channel.

Oversubscribing DOCSIS Relies on Statistical Subscriber Behavior Over Shared Bandwidth

In the example above, each user has on the average 0.23 Mb/s of bandwidth. So how can cable operators offer packages of up to about 10 Mb/s to subscribers? *Statistics* is the answer. Years of experience and analyzing data flow on their networks have taught cable operators a lot about the statistics of data usage. If you are surfing the web, for example, you might use high speed to download a page, but then you will spend time reading that page. While you are reading the page, you are not using network bandwidth at all. So your neighbor can use the bandwidth. You might need your speed to download your email, but after it is downloaded, you don't use any bandwidth as you read it. This property of sharing the bandwidth statistically can be quantified using the terms *oversubscription*, or *overbooking*.[2] The cable industry has found that if it oversubscribes the bandwidth by 100:1, few if any subscribers will notice that they are not getting the full bandwidth they are paying for. In the example, if all 167 data subscribers get 10 Mb/s contracted data rate (a very high number today), then the total data rate that would be needed is 167 times 10 Mb/s, or 1670 Mb/s. But at 100:1 oversubscription, only 16.7 Mb/s, or 44% of one

[1] The DOCSIS standards include provisions to protect one subscriber's data from the prying eyes of another subscriber.

[2] Mathematically, oversubscription is equal to the total contracted bandwidth for all subscribers on the node, divided by the peak bandwidth.

DOCSIS channel is needed. Two nodes could easily be combined, with subscribers all getting what they pay for. Of course, the same assumptions on oversubscription apply to FTTH, with its much faster data rates.

Thus, we have three concepts here. The *peak bandwidth* is the amount of bandwidth that can be supplied on the network. In the (pre – 3.0) DOCSIS example above, the peak bandwidth is 38 Mb/s, set by the speed of the DOCSIS channel. The *average bandwidth per user* is the peak bandwidth divided by the number of subscribers sharing that bandwidth. It is a figure of merit only, and has limited significance in the real world. The *contracted bandwidth* is the bandwidth the operator sells to each user, and the total can and will exceed the peak bandwidth. When you total the bandwidth per user for all users, it will far exceed the peak bandwidth, but this works because of statistics.

This concept of using statistics to allow overbooking is by no means limited to data applications. To *not* rely on oversubscription would make many things uneconomical. Airlines routinely overbook flights, relying on studies that show a certain percentage of fliers don't show up. The telephone system cannot handle all phones being in use at the same time, but the phone companies know from many years of experience how much they can oversubscribe their facilities. The roadways are not sized to accommodate all cars at the same time (to do so would probably take more than all available land), but they don't need to, because all cars are not on the road at the same time.

Statistics allow FTTH systems to overbook bandwidth in the same way, but due to the wider bandwidth, the numbers go up by nearly a couple of orders of magnitude.

Comparison of the Technologies

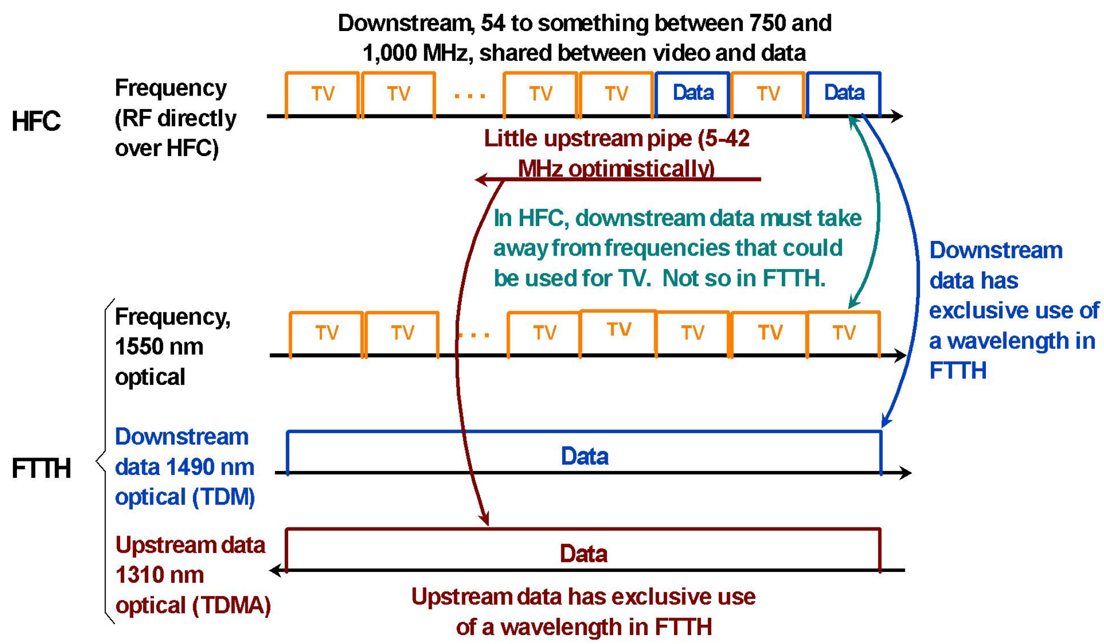

Figure 1. Graphical representation of HFC vs. FTTH Transmission

Figure 1 illustrates the differences in capacity of HFC and FTTH networks. HFC, at the top, has only a single data path in each direction: 54-750 (or more, or less) MHz downstream, and 5-42 MHz upstream (a "little pipe"). The downstream bandwidth must be shared between TV programs and data, as indicated. (Voice is included in the data bandwidth, since in nearly all cable voice systems today DOCSIS is used for voice transport.) The maximum frequency of most cable systems is 750 MHz, because that was state-of-the-art when most systems were rebuilt during the late 1990s and early 2000s. Some more recently rebuilt systems extend the maximum frequency to 860 MHz, and today 1,000 MHz is being put in when needed. HFC systems require significant modifications to operate at 1,000 MHz, while FTTH can carry that spectrum over the all-optical 1550 nm wavelength easily.

HFC Upstream Bandwidth is limited

To handle upstream bandwidth needs, the spectrum from 5 to 42 MHz is used. Commonly today there are two or three services that share this band. DOCSIS upstream channels occupy most of the spectrum. Return signals from set top boxes occupy some of the bandwidth, and status monitoring of nodes, power supplies and sometimes RF amplifiers take more of the available bandwidth. Since the band is so narrow, less data can be put in it than can be accommodated in the downstream direction.

Furthermore, the upstream spectrum has some serious limitations in usage. Because of the tree-and-branch nature of HFC systems, there are typically several network branches feeding back to a node in the upstream direction. Each branch carries noise generated in the network and noise ingressing into the network. The summation of noise from all branches is called *noise funneling*. The bandwidth from 5 to about 15 or 20 MHz is of limited usefulness due to the amount of noise that exists, mostly from man-made sources. These include high powered signals from amateur radio and shortwave broadcast stations (plus still some military communications), electrical noise from motor commutation, power switching, electrical dimmers, florescent lighting, and numerous other sources. Because of this noise, the lower frequencies are limited to low data rates and robust forms of modulation. These limitations mean that the signal can get through in the presence of a lot of noise, but it also means that the data rates are low, so not that much data can get through. From about 20 – 40 MHz the situation regarding noise improves, so more data can get through. (Above about 40 MHz other limitations apply.)

Even in the choice band of about 20 – 40 MHz, there are still some noise funneling and other interference problems, so the modulation formats used for upstream DOCSIS were limited to relatively robust formats, which imply that the amount of data that can be transmitted per unit of bandwidth is also lower. However, as the state-of-the-art progressed and the pressure for more upstream bandwidth mounted, less robust modulation formats were considered, written into the DOCSIS 2.0 specification and carried over into the 3.0 specification. These higher density but less robust modulation formats theoretically can deliver up to about 30 Mb/s of *wire rate* data in a 6.4 MHz upstream channel. (Note the upstream channel is wider than is the standard 6 MHz downstream channel.) *Wire rate* refers to the actual data rate on the system. The payload data rate is lower, how much lower depending on the amount of error correction used, as well as the overhead. A maximum rate of 30 Mb/s is often quoted. Most cable operators appear reluctant to use this higher-density modulation because it is more susceptible to the perils of the upstream path.

FTTH uses additional wavelengths of light for additional capacity

FTTH systems provide an RF downstream path on 1550 nm as shown, nearly identical to that provided by HFC. FTTH being newer, almost all systems operate to 860 MHz, and some are specified to 1,000 MHz. There is no need to remove any bandwidth for data because data is carried on different wavelengths on the fiber – it is like having separate cables for data, but rather than having to use different cables, FTTH simply uses different wavelengths of light. This would be analogous to using, say, red and blue lights to signal from ship-to-ship in the old Morse code light-signaling days. If you needed higher signaling speeds than one signalman could handle in each direction, you could use two signal lights, one red and one green. At the receiving ship, you could assign two signalmen, one to copy only the red signal and the other to copy only the green signal. Thus, you use the same medium (in the case of FTTH a single fiber optic strand) to send two signals, one an RF broadcast signal and the other a digital (on/off) data signal on the same fiber strand at the same time, but using a different wavelength ("color") of light. You use a third wavelength to send data back upstream.

Thus, FTTH offers the best of both worlds: the RF bandwidth of the widest HFC systems, *plus* lots of bandwidth for data – up to 2.4 Gb/s downstream and 1.2 Gb/s upstream with today's GPON systems. Use of the RF wavelength is optional – you can build FTTH systems without using broadcast if you wish, by placing video on IPTV using the data wavelengths at 1490nm and 1310 nm.

Downstream Bandwidth is still highest for FTTH

Several standards are used for FTTH, with the most common today being EPON and GPON (BPON is still being used, but is being replaced with GPON, and is not covered here). EPON specifies a bandwidth of 1 Gb/s (1,000 Mb/s) both downstream and upstream. GPON in it's most common form provides a *wire rate*[3] of 2.488 Gb/s downstream and 1.244 Gb/s upstream. Most commonly today, this bandwidth is shared among 32 or fewer subscribers. Knowing these numbers, we can prepare a table comparing what FTTH can do vs. what cable TV can do today.

Figure 2 (below) shows the high speed downstream bandwidth that can be offered by FTTH and by cable TV. Cable TV lacks very badly compared with FTTH, even when we use the lower EPON bandwidth. Each operator may provide different contracted bandwidths, but the rates available with FTTH are vastly superior to those that can be offered by cable TV using DOCSIS. The four-channel bonding line applies to DOCSIS 3.0 and is described further below. It is better, but still not really competitive.

	Peak Rate (Mb/s)	Average Rate per subscriber (Mb/s)	Contracted Bandwidth	Programs Lost
FTTH	1,000+	31+	By policy	0
Cable TV, 1 ch.	38	0.23	By policy	10
Cable TV, 4 ch. bonding (see below)	150	0.92	By policy	40

(Note: data rates in Mb/s, see text for assumptions)

Figure 2. Downstream bandwidth comparison

Upstream Bandwidth Limitations with DOCSIS 3.0

In terms of upstream bandwidth, DOCSIS falls even further behind FTTH. We showed above that the upstream bandwidth of a DOCSIS signal is limited by problems with the return path in an HFC system. The return path is limited in bandwidth, parts of that bandwidth are very unreliable due to noise, and even the reliable parts are not reliable enough to allow the high levels of modulation provided in the downstream direction.

The most common data rate for a single DOCSIS upstream channel today is about 10 Mb/s (16-QAM modulation), though we showed above that data rates up to about 30 Mb/s are included in the specifications. We can construct a table similar to Figure 2 but for the upstream direction, making the same set of assumptions.

[3] Wire rate is the actual data rate on the medium, in this case fiber. The payload, or usable, data rate is somewhat less, but is still quite high. The wire rate of EPON is actually 1.25 Gb/s, but it is usually expressed as 1 Gb/s, the difference being due to the encoding overhead used.

	Peak Rate (Mb/s)	Average Rate per Subscriber (Mb/s)	Contracted bandwidth
FTTH	1,000 +	31+	By policy
Cable TV, 1 ch	10	0.06	By policy
Cable TV, 4 ch bonding	120 optimistically	0.72	By policy

Figure 3. Upstream bandwidth comparison

Figure 3 (above) illustrates the upstream bandwidth of FTTH and two scenarios for cable TV upstream bandwidth. The first scenario, with a single upstream channel, assumes common parameters in use today: a single channel of 16-QAM modulation operating at the highest specified baud rate. This yields a single upstream channel of 10 Mb/s shared among the assumed 167 subscribers on the node. This shows that the average upstream bandwidth per subscriber is only about 0.06 Mb/s, compared with over 31 Mb/s for FTTH. If we assume DOCSIS 3.0 bonding of four channels (not available commercially as of this writing), and if we assume the most optimistic modulation (not likely to be practical in most systems), we still get less than 1 Mb/s for cable, vs. 31 Mb/s for FTTH.

Channel Bonding is the technique used to shift capacity from video delivery to downstream data delivery

Compared with earlier versions of DOCSIS, the 3.0 version contains a number of features intended to lower costs and improve efficiency in the operation of HFC networks. Some of the improvements relate to the physical configuration of the headend portion of the DOCSIS network, and some permit marginal improvement in the data rate. But the feature that creates most of the talk and most of the bandwidth improvement is called *channel bonding*. It is analogous to pair bonding in telephone's twisted pair cable world.

In conventional telephone practice, a pair of wires is used to carry signals to and from telephone sets. The two wires are twisted together for technical reasons, so the wire pair has come to be known as a *twisted pair*, or sometimes just a *pair*. It is common to install cables consisting of many twisted pairs, some of which are used and some of which are held in reserve. Techniques have been developed to use the twisted pair for data as well as for voice, with the best-known techniques being codified into a set of standards collectively known as DSL, for *digital subscriber line*. Physical limitations restrict the maximum data rate that can be carried on a single twisted pair. In order to overcome this limit, one technique that has been developed is called *pair bonding*, in which extra sets of twisted pairs are used to carry part of the data. At the transmitting point, a single data stream is taken apart. For example, half the bits may go to twisted pair #1 and half to twisted pair #2. At the receiving end, the two halves of the data are reassembled into one data stream at twice the rate that could be carried over a single pair. More than two pairs can be bonded.

In HFC access systems used by the cable TV industry, there are no twisted pairs. Rather, signals are transmitted on radio frequency carriers (RF) on fiber optic cable most of the way to the subscriber, then transferred to coaxial cable (still modulated on RF carriers) for the remainder of the journey. Only a single physical path is used to transmit signals to a group of subscribers, as opposed to use of multiple twisted pairs of wires. Different signals are transmitted on different frequencies, called *channels*, after the common TV usage. The concept of channels dates back at least to the first television transmissions. In North America, channel 2 consists of the 6 MHz frequency band from 54 to 60 MHz. Channel 3 is 60-66 MHz and so on. (Not all channels are contiguous in frequency, but such details are unimportant in the present context.) When you tune a TV from channel 2 to channel 7, for example, you are changing the frequency the TV is receiving, from 54-60 MHz to 174-180 MHz.

RFoG is emerging as a replacement for the coaxial cable portion of the HFC network

The cable industry is developing a new standard for fiber to the home that is unique to cable TV. As of early 2008 work had started on the standard, but was a long way from completion. Its name may change, but as of this writing it is called RFoG, for *radio frequency over glass*. It is really standard cable TV hybrid fiber-coax (HFC) plant, with a node size of one. That is, RFoG takes the fiber to the side of the home. But unlike true FTTH technologies like GPON or EPON (or BPON), RFoG doesn't have the additional wavelengths for upstream and downstream data transmission. RFoG relies on all RF-modulated carriers in both directions. There are some serious implications in the upstream direction related to the turn-on and turn-off of individual optical transmitters, and there are design issues related to use of RF-modulated optics that don't exist in true PONs.

Compared with normal HFC, RFoG does offer some advantages. In some concepts, the plant is all-passive, which does improve reliability, and the lack of amplifiers can also lower operational expense. The quality of the upstream signals is likely to be higher, meaning that the higher bandwidth modes of DOCSIS are more likely to work than they are in normal HFC plant. But the data rates will not be higher than they are with any other DOCSIS-based data transmission, and can never approach what true FTTH offers.

Cable's Limited Answer to FTTH Includes DOCSIS 3.0 and RFoG

Figures 2 and 3 showed dramatically why cable needed to have a better answer than DOCSIS 2.0. This need was one basis for creation of the DOCSIS 3.0 specification. DOCSIS 3.0 uses *channel bonding*, introduced above, to partially answer FTTH data rates. In round numbers, you can multiply the peak and average data rates available through cable modems, by the number of channels bonded together. For example, as shown in the last row of the table, if you bond four channels, you get a peak rate of 150 Mb/s, though you only have an average rate (for what that means) of just under 1 Mb/s.

The last column of the table tells a strong story: each 38 Mb/s of data bandwidth in the downstream direction requires one RF channel as shown above. One RF channel today will typically be enough to transmit about 10 digitally compressed video programs. So the cable operator must sacrifice the ability to transmit 10 video programs for every 38 Mb/s of data he wants to offer. Only a certain total amount of spectrum (number of RF channels) is available due to limitations in the HFC system. The operator must figure out how to optimize use of his limited channel resources.

Note that for FTTH, there are *no* programs lost to data. This is because, for RF broadcast, the video programs travel on a separate wavelength from the data. There is no relationship between them. Of course, it is also possible to transmit video on the 1490 nm data path, using a technology that is in the news a lot right now, IPTV. In this case, one simply uses some of the considerable bandwidth of the 1310/1490 FTTH system for additional video capability.

True Next-Generation Broadband is Achievable via FTTH, not via HFC.

In 2007, citing the country's need to secure its economic future in the Information Age, Senator John D. Rockefeller IV (D-WV) introduced a Senate resolution calling for universally available 100 Mb/s Internet connections by 2015. The FTTH Council wholeheartedly endorses this goal, and strongly believes that reaching it requires a rapid shift to end-to-end fiber networks.

Even using the enhancements provided by DOCSIS 3.0, HFC systems cannot deliver on this goal. If every subscriber in our earlier example was supplied with 100 Mb/s, you can see that DOCSIS *must* do channel bonding. Without it, there is simply no way to attain this goal. On the other hand, FTTH easily fits into this model. With 32 subscribers each having 100 Mb/s, the oversubscription ratio is only 3.2:1, far from the 100:1 that cable TV uses today. On the other hand, with DOCSIS 3.0 and the example above (167 data subscribers), you could achieve a peak data rate of 100 Mb/s, but if everyone was taking that data rate, your oversubscription ratio would be 111:1, worse than the 100:1 ratio used earlier, and much, much worse than FTTH's oversubscription ratio.

Of course, you can bond more channels in DOCSIS 3.0 (if and when more channel bonding is commercially available) and improve the oversubscription ratio, but you lose even more video programs when you do. Given the important role video service offerings play in consumer take-up of triple-play services, it is difficult to imagine this sacrifice being made on a regular basis. In addition, even with channel bonding, DOCSIS 3.0 is nowhere near FTTH in terms of what kind of data it can really serve. Even when we transmit IPTV on FTTH, we will have much more capability than will the traditional HFC cable operator.

DOCSIS 3.0 does not result in more overall cable TV system capability; rather it allows cable operators to choose to use existing capacity for either the delivery of more high speed data bandwidth, or the delivery of more video channels, but not both. Of course, you can do the same thing without channel bonding if you don't need to deliver more than 38 Mb/s to any one subscriber. FTTH delivers the entire video spectrum over optical fiber all the way to the customer on one wavelength of light – 1550 nanometers. Rather than share that capability with data, FTTH provides much more upstream and downstream bandwidth for high speed data access on different colors of light – 1310 nm upstream and 1490 nm downstream.

Conclusion – FTTH provides superior capability

DOCSIS 3.0 was developed in part to answer the speed of FTTH. While it does increase downstream speed to a single user, it does so at a cost in terms of the number of video programs that can be transmitted. DOCSIS 3.0 can allow the industry to say it takes us toward the "100 megabit nation," but its capability is not by any stretch of the imagination the same as that of FTTH. DOCSIS upstream is in even more trouble, and upstream is becoming a valuable feature with work-at-home, peer-to-peer and gaming traffic. DOCSIS 3.0 is a valiant effort to eke more performance out of existing cable TV HFC plant, but if falls well short of being as good as FTTH.

Appendix A: Explanation of RF Modulation and Video Channels

This section contains in introduction to some comments referenced above, which may be less familiar to readers not from a technical background. With the exception of IPTV, all broadcast radio and TV signals are *modulated* onto *RF carriers* of differing *frequencies*. For TV, the carriers are organized by a *channel* number, which is simply an easier way of talking about frequency. What's an RF carrier? Well, let's start with the ways we can send electricity, and then we'll extend that to signals, which will get us to RF carriers, and in turn that will get us to channels. Quite a bit for a few paragraphs!

The simplest form of electricity transmission is called *direct current*, or dc. This is what you get from a battery. It simply means that a steady voltage is applied from the battery to a load, for example a light in a flashlight. That voltage causes a current to flow, and the current heats the light filament and makes light. Let's use a water pipe example. Suppose we have a water reservoir and a pump connected to it, such that the pump tries to pump water out of the reservoir, through some sort of apparatus, and back to the reservoir (we don't want to waste water). The pump puts the water under pressure, and this pressure is analogous to the voltage produced by the battery. Then when we pump water through some device, say a water wheel, the volume of water being pumped is analogous to the current in the electrical circuit. The pump produces a steady water pressure ("voltage"), which results in a steady flow of water ("current"). We would refer to this as direct current, or dc. We can refer to the "positive pressure" voltage terminal as + (plus or positive, analogous to pushing out water) and the negative voltage terminal as – (minus or negative, analogous to pulling the water back into a reservoir). (If you know a bit about electron flow, you may argue on which way we define the flow, but it doesn't matter here.)

The analogy is pretty straightforward so far, but it starts breaking down a bit at the next point. It turns out that with electricity, maybe unlike with water, there are advantages to NOT operating the electricity system at a steady voltage (and hence current). If we alternate the direction of the water pressure, this is analogous to *alternating* the voltage of the electrical circuit. This results in current flowing first in one direction then in the other. We call this *alternating current*, or ac. You can do a lot of interesting things with ac that you cannot do with dc. Our house current is ac, though in the early days of electrification there was a debate over whether ac or dc should be used. In North America the voltage goes through 60 sets of + and – each second, so we say the frequency is 60 Hertz (Hz). Hertz is the unit of frequency used, named after Heinrich Hertz, a pioneer in developing radio technology.

Now it turns out that speech is represented as a complex, ever-changing series of ac signals, that happen to occupy the frequency range from about 300 Hz to 3500 Hz or so. Music can occupy the frequency range of about 20 to 15,000 Hz, which we write 15 kHz, for kilo (thousand) Hertz. This is roughly the range of human hearing, though some folks will claim hearing up to about 20 kHz. Video is a much more complicated thing than is sound, but it also occupies a certain range of frequencies – traditional analog video as used in North America occupies from maybe 10 Hz to 4 MHz (mega, or million, Hertz). One thing about the frequencies occupied by both audio and video, is that for practical reasons, you cannot transmit them directly over the air or over cable TV or satellite or the broadcast portion of FTTH. There are technical reasons, and there is the practical issue that you cannot transmit multiple sets of speech or video on that same set of frequencies ("frequency band"). What you can do, though, is to *modulate*, or somehow change, a higher frequency with the audio or video, and you can transmit that higher frequency with the modulation on it.

For example, let's start with AM (amplitude modulated) radio, the oldest form of radio communications that most of us are familiar with. Each radio station is assigned a frequency on which to send out signals. For example, WSB radio in Atlanta is assigned to send out signals on 750 kHz, so they send speech and music to their transmitter, which changes the amplitude, or strength, of it's *carrier* at 750 kHz, according to the material being sent. We call this process *modulation*, which simply means we are changing something about the signal, in this case, its amplitude. The frequency of 750 kHz in this case is called the *carrier*, or *carrier frequency*, because it carries the modulation, or intelligence you want to hear. It corresponds to a real signal generated by the transmitter, and sent to an antenna, which radiates the signal to radios. You will often see a radio station identified by its frequency ("AM 750"), as this tells you how to tune to the station on your radio (you tune to 750). Similarly, FM radio station WABE is assigned to send its signal on 90.1 MHz. Being FM, it changes, slightly, the frequency of its carrier with the audio supplied from the studio. They often identify themselves as 90.1, because, again, you tune your radio to read 90.1 if you want to listen to them.

Now we could carry the same concept of frequency over to TV signals if we wanted to, but it gets more complicated. The reason is that, due to the nature of a TV signal and the frequencies involved, we need a rather wide *set* of frequencies in which to transmit a TV signal. It worked out reasonably well to define frequency bands that are 6 MHz wide for use in North America (Europeans use wider bands – the reasons are way too much information for now.) The lowest TV frequency set used in North America is 54-60 MHz. The next set is 60-66 MHz and so on. But unlike the AM and FM radio bands, not all TV frequencies are contiguous – that is, they don't lie next to each other. That fact and the complexity of describing the set of frequencies used, drove people to develop a different paradigm for describing the frequency on which a TV station transmits. We called the frequency set 54-66 MHz channel 2, the 60-66 MHz set channel 3 and so on. Channel 4 is 66-72 MHz, but then we get into the first non-contiguous channel: you'd expect channel 5 to be 72-78 MHz, but you'd be wrong! It is defined to be 76-82 MHz. What's the deal with the 4 MHz gap? It is assigned to certain non-TV services for technical reasons dating back to the beginning of TV broadcasting. There are yet other gaps in channel assignments, and the channel vs. frequency table for cable TV is only partially the same as for off-air. Satellite channel assignments have no relation in frequency to anything else.

All that complexity is hidden from you, the user. You merely have to remember the channel you want, and today with electronic program guides, even the channel number is losing its significance. What about digital TV? When the digital system was being developed in the early 90s, an early decision was to keep the same channel vs. frequency definitions, and to tell the engineers to figure out the most efficient way to fit the digital video into those channels, which they did.

So that's how we get to the channels we referred to above: they simply are a convenient way of keeping track of a set of frequencies used for some purpose, be it to send analog video, digital video, or a bunch of data in a DOCSIS channel. Note that not all channels are 6 MHz wide. In the downstream direction, they pretty much are all 6 MHz wide in order to allow different services to intermingle efficiently. In the upstream direction, things are not so standardized. Some lower-bandwidth upstream DOCSIS channels are 3.2 MHz wide, and some higher-bandwidth upstream DOCSIS channels are 6.4 MHz wide. But the concept is always the same: a channel is a set of frequencies occupied by a signal carrying some sort of intelligence. (At least we hope there is intelligence in the signal, though sometimes we are disappointed. But we maintain that failure to carry intelligence is *not* the fault of the engineers.)