

Author: Bennett Cyphers

A publication of the Electronic Frontier Foundation, 2019.

“The Case for Fiber to the Home, Today: Why Fiber is a Superior Medium for 21st Century Broadband” is released under a Creative Commons Attribution 4.0 International License (CC BY 4.0).

View this report online:

<https://www.eff.org/wp/case-fiber-home-today-why-fiber-superior-medium-21st-century-broadband>



The Case for Fiber to the Home, Today

**WHY FIBER IS A SUPERIOR MEDIUM FOR 21ST
CENTURY BROADBAND**

BENNETT CYPHERS
Staff Technologist

OCTOBER 11, 2019

Table of Contents

Executive Summary	5
Glossary	6
Technical Background	10
DOCSIS 3.1 and the Future of Coax	15
5G and the Future of Wireless	18
Fiber Today and in the 21st Century	21

Executive Summary

The debate over the best infrastructure to deliver fixed last-mile broadband service in the 21st century is settled, and fiber is the undisputed winner. Fiber-to-the-home deployments are a better option for consumers today, and they are the only option that will allow expansive, efficient upgrades to America's networks for a generation.

This is not to say that no broadband technology will ever surpass fiber-optics, but we know the limitations of existing technologies in use today. Currently, the alternatives to fiber face headwinds that fiber does not, including limited bandwidth, attenuation, noise, upstream/downstream asymmetry, and latency. While other means of delivering high-speed broadband are not too far behind fiber right now, the properties of each technology will allow fiber deployments to scale up quickly and easily while copper and wireless broadband networks will struggle to keep up. If we install fiber-to-the-home connections today, we'll be able to upgrade the transmitters at each end without touching the underlying cables, yielding massive performance increases at low cost for decades to come. Fiber will enable the next generation of applications that depend on high-throughput, low-latency, high-reliability connections. There is an identifiable "speed chasm" between fiber and everything else that is only going to grow more pronounced in time.

This whitepaper gives a brief technical background and explains key concepts for understanding internet services such as bandwidth, latency, channel capacity, and noise. Understanding these concepts is essential in order to assess and compare broadband networks. This whitepaper then evaluates three different classes of last-mile broadband connections—coaxial cable, wireless, and fiber—from a technical perspective. It argues that through this lens, fiber is indisputably the best option for consumers today. New wireless technologies, like mmWave 5G, will supplement rather than compete with fiber-to-the-home technology. And aging wireline technologies like DOCSIS are already being incrementally replaced by fiber.

This paper focuses on the "last mile" of broadband connections because a vast majority of the internet infrastructure before the last mile has already transitioned to fiber. Lawmakers and regulators in positions of determining infrastructure policy must understand the realities of networking technologies in order to properly assess the capability of networks to absorb greater user demand.

This paper does not explore policy mechanisms to address the fiber deficit currently facing the United States market. EFF intends to publish such material at a future date. The purpose of this paper is to educate policymakers as to the technological differences between different broadband networks and as to the future proof nature of fiber networks. With the advent of cloud computing, virtual reality, gaming, telehealth, remote services, and high capacity services we have not yet imagined yet, policymakers must grapple with updating the Internet's infrastructure for the 21st century so that the American people are not left behind.

Glossary

4G - The fourth generation of cellular network technology. 4G was standardized in 2008 by the International Telecommunication Union (ITU) as “IMT-Advanced,” and is specified to support speeds up to 1 Gb/s down. However, real-world systems have more commonly achieved a maximum of a few hundred megabits, and an average of a few tens of megabits.

5G - The fifth generation of cellular network technology. 5G is still in the process of being standardized by the ITU as “IMT-2020.” 5G will use low- and mid-band frequencies below 6GHz for mid- and long-distance communication, as well as millimeter wave frequencies above 24GHz for short-range, high bandwidth communication. 5G promises to support high-throughput communication up to 10Gb/s as well as “last-hop” latencies as low as 1-4ms.

Absorption - A type of attenuation that occurs when signal-carrying photons are absorbed by other matter. Wireless signals may be absorbed by walls, foliage, or the air; beams of light in a fiber-optic cable are absorbed by tiny imperfections in the fiber’s glass core. When signal-carrying photons are absorbed by environmental obstructions, the signal becomes weaker.

Amplifier - In coaxial cable deployments, a device which amplifies information-carrying signals. Amplifiers are installed along coaxial cable running between a headend and customer terminals in order to boost signal power. Amplifiers can add noise to the system as well and decrease the signal-to-noise ratio.

Amplitude - A measure of the power of an electromagnetic wave. Waveforms generated with more power will have greater amplitudes; this makes signal-bearing waves easier to detect relative to background noise.

Attenuation - Loss of signal power over distance. Attenuation is a factor in all methods of signal propagation. Wireless signals attenuate according to the inverse-square law in free space. Electrical signals in coaxial cables attenuate primarily due to electric impedance. Guided beams of light in fiber-optic cables attenuate primarily due to absorption.

Bandwidth - The range of frequencies available in a given channel. Bandwidth is defined as the difference between the maximum frequency available in a channel and the minimum frequency available.

Base station - In cellular networking, an installation that generates and receives wireless signals in order to provide wireless service to cellular phones and other mobile devices. Also known as a “cell site” or “cell tower.”

Bits per second (B/s) - Measure of information throughput. A bit is a single value, 1 or 0. Modern broadband channels can transmit many millions (megabits or Mb) or billions (gigabits or Gb) of bits per second.

Cable headend - A facility for processing television and internet signals from a service provider's regional network and transmitting them over the "last mile" from the larger network to customers' buildings.

Cellular network - A network in which the last-mile link is wireless. Cellular networks use cell sites, or base stations, to broadcast wireless signals "over the air" and provide internet service over a wide area. Cell sites usually communicate with cellular consumer devices, like phones, using radio-frequency signals. Cellular network standards are generally referred to by "generation;" the newest generation to be implemented is the fifth, or 5G.

Channel - A logical connection over which information-carrying signals may be transmitted. A channel comprises a transmitter, a receiver, and the medium over which signal travels between those two. Examples of channels are the connection between a WiFi transmitter and a laptop computer, as well as the connection between a cable headend and a customer's modem.

Coaxial cable - A copper cable consisting of a central conducting wire and an outer conducting tube separated by an insulating sheath. The central wire carries current in one direction and the conducting sheath carries it in the other direction. Most coaxial cables can carry radio frequencies up to around 3 GHz over relatively long distances, and are designed to minimize electrical interference.

Crosstalk interference - Interference that occurs when an electrical signal in a medium interacts with another signal from outside the medium. For example, in unshielded twisted pair wiring, signal from one pair can interact with signal in a nearby wire, adding noise and diminishing the channel's information capacity.

DOCSIS - Short for Data Over Cable Service Interface Specification. DOCSIS is the international standard for carrying internet signals in last-mile networks over coaxial cable. The most recent version of DOCSIS is 3.1.

Electromagnetic spectrum - Commonly referred to as spectrum, it refers to the full range of frequencies that can characterize electromagnetic waves. Portions of spectrum are often referred to as "bands" and described by their middle frequency; for example, the "5 GHz band" might refer to the section of spectrum between 4.95 and 5.05 GHz. Different bands are used for transmitting different kinds of signals, both in guided media (like cables) and "over the air" as unguided waves.

Electromagnetic wave - The oscillation of an electromagnetic field. Electromagnetic "waves" are the representation of electromagnetic radiation in classical theory. Electromagnetic waves always propagate at the speed of light. Waves are measured by their amplitude (power) and frequency (speed of oscillation).

Fiber-optic cable - A transparent thread made of high quality glass utilized for fiber optic communications. Fiber-optic cables operate as waveguides for beams of light. A beam of light shined down one end of a fiber-optic cable will reflect off the insides of the cable and be completely contained within the glass core. "Single-mode" fibers are used for links longer than a few meters. These cables are extremely thin, around 9 micrometers, and only allow light to travel in one path or "mode" through the fiber in order to minimize noise.

Forward error correction - Method for encoding information in a signal with some redundancy so that the signal is robust to noise. Forward error correction uses error-correcting encoding to send information such that, if small portions of the signal are transmitted incorrectly, the receiving end of the channel can recognize and correct the errors.

Frequency - A measure of the speed of oscillation of an electromagnetic waveform. Frequency is usually measured in oscillations per second, or Hertz. For example, a radio station operating at 88.9 Megahertz (MHz) has electrons oscillating on its antennae at 88,900,000 cycles per second. Frequency is inversely proportional to wavelength, meaning the higher the frequency, the shorter the wavelength—and vice versa.

Hertz (Hz) - A unit for measuring frequency, equal to one oscillation per second.

Interleaving - Transmission technique which makes forward error correction more effective. Errors in DOCSIS systems tend to occur in bursts. Forward error correction is better at dealing with errors that are spread out over time, so operators can "interleave," or mix up, symbols before they are sent. This increases the effectiveness of error correction at the expense of more latency.

Internet backbone - High-capacity portion of the internet where large amounts of data are exchanged between different regional networks and different internet service providers. Links in the internet backbone are typically extremely low latency and high throughput, and may span oceans or continents.

Inverse-square law - Physical law governing the rate at which wireless signal power attenuates in a vacuum. For every doubling in distance from a signal's source, the power of the signal is reduced by a factor of 4 (75%).

Jitter - Deviation from expected timing in a series of packets. Jitter is caused by sudden, random spikes in latency. In broadband systems, it may be caused by dropped packets, sudden delays due to congestion on shared networks, or delays in upstream traffic due to bandwidth allocation. Jitter can negatively impact time-sensitive applications like video chat or online gaming.

Last mile - The portion of the internet which connects service providers' shared infrastructure to end users, such as homes or businesses. In a DOCSIS cable network, the last mile is the connection between the cable headend and the customer's building. In a cellular wireless network, the last mile is the wireless connection between a base station and a mobile device. Sometimes also called the "first mile."

Latency – The time it takes for a signal to be transmitted over a channel. This includes encoding time, travel time, and decoding time.

Millimeter wave – Refers to signals between 30GHz and 300GHz, designated by the ITU as “Extremely High Frequency (EHF)” signals. 30GHz waves have a wavelength of approximately 1 millimeter.

Modem – Short for “modulator-demodulator.” A consumer device for receiving and transmitting internet signals over a last-mile wireline connection. Modems are usually sold by internet service providers and used to connect customers to the wider network.

Noise – Any unwanted or unintended modifications to a signal that occur during transmission. Noise can come from a variety of factors, including crosstalk (interference with other signals), ambient radiation, and errors in transmitters or receivers.

Optical line terminal (OLT) – The headend of a fiber-optic Passive Optical Network. A single OLT may serve internet to several dozen optical network terminals (ONTs). Signals from the OLT are directed to individual ONTs by passive optical splitters (lenses) that duplicate and redirect optical signals.

Optical network terminal (ONT) – The consumer end of the last mile connection in a Passive Optical Network. An ONT receives downstream signals generated by its OLT, interpret the packets meant for it, and responds with its own upstream signals on the shared fiber optic cable.

Passive optical network (PON) – Network architecture for last-mile internet over fiber optic cable. A single optical line terminal (OLT) drives signals to several optical network terminals (ONTs). The OLT sends a single stream of downstream traffic that is seen by all ONTs. Each ONT reads the content of only those packets that are addressed to it; packets can be encrypted to prevent eavesdropping. ONTs respond to the OLT by taking turns, known as “time-division multiplexing.”

Scattering – Related to absorption; scattering occurs when photons are reflected, or absorbed and re-emitted, by matter. Scattering is one of the chief causes of attenuation and noise in wireless signals, especially when they pass through obstructions like buildings and foliage.

Shannon limit – The absolute upper bound on the amount of information a channel can carry, in bits per second. The Shannon limit is a function of the bandwidth of a channel and its signal-to-noise ratio.

Signal – Any time-varying wave or function that carries information. Electromagnetic waves can transmit signals using amplitude modulation (AM), frequency modulation (FM), binary pulsing, or other means.

Signal-to-noise ratio (SNR) – The ratio between the power of information-carrying signal and the average power of the noise in a channel. Along with bandwidth, the SNR determines the maximum theoretical information capacity of a channel.

Symbol – The smallest coherent unit of a signal. Every signal can be thought of as a sequence of symbols. Each symbol takes on one out of a possible set of values. In the simplest case, a symbol is a bit: either a 1 or a 0. Symbols may be represented as high or low voltage values, as pulses of light, or as different shapes of electromagnetic waveform.

Throughput – The rate of information that a channel can carry, usually measured in bits per second.

Wavelength – A measure of the distance between peaks in an electromagnetic wave. Inversely proportional to frequency. Higher-frequency waveforms have shorter wavelengths.

Technical Background

All information is transmitted via *signals*. Telegraphs, radio, land-line telephones, the spacecraft Voyager 1, and 5G-enabled phones all rely on signals transmitted by some kind of *electromagnetic wave*. Signals can either be analog, as with AM radio or traditional phone service, or digital, like the signals used to carry data over the internet. A digital signal is a sequence of information-carrying symbols, like letters in a string of text.

Signals are carried over *channels*. A channel is a connection that can carry a signal from one place to another. Different channels are useful for different purposes, and there are tradeoffs involved with choosing to use one kind of channel over another. Land-line telephone signals are transmitted by electricity over copper wires, which are cheap and reliable. Analog radio is transmitted “through the air” by radio waves, which can carry simple signals in all directions over long distances. And the backbone of the internet uses guided light waves in fiber-optic cables to transmit huge amounts of information for hundreds of miles, but building and installing these cables can be expensive. In all of the aforementioned channels, specially-formed electromagnetic waves are used to carry the signal.

Bandwidth and noise

Electromagnetic waves are described by their *amplitude* (power) and *frequency*. The frequency of a waveform is measured in Hertz (Hz), or oscillations per second. Different channels can carry different frequencies of EM waves. For example, old-school analog phone lines were designed to carry frequencies from 300 to 4,000 Hz, approximately the range audible to the human ear. The range of frequencies a channel can carry is called its

bandwidth. The bandwidth is calculated simply by subtracting the minimum frequency a channel can carry from the maximum. A channel spanning 0 to 1,000 Hz has 1,000 Hz of available bandwidth, and a channel spanning 100,000 to 101,000 Hz has the same. Bandwidth helps determine how much information a channel can transmit: more bandwidth means more information capacity.

Noise is a general term for all the random, chaotic, and meaningless disruptions that information-carrying signals in a channel might suffer. Electromagnetic noise is everywhere; radio waves are constantly being pumped into the air by cell towers, police radios, power lines, and the sun. These sources of radiation can interfere with individual signals traveling from one device to another through the air, and are part of the reason wireless signals can't travel over infinite distances. In shielded media like coaxial cables and fiber optics, imperfections in shielding or connections can allow noise to "leak" in; signal transmitters and receivers can also add noise by themselves. The *signal-to-noise ratio* (SNR) in a channel is the ratio of the power of the signal to the power of the noise.

All signals degrade over distance; this is referred to as *attenuation*. Wireless signals, like radio waves, lose power according to the inverse-square law: that is, if you travel twice as far away from the source, the signal will be at least four times as weak. Wireless signals also attenuate due to interactions with the environment, including *absorption* and *scattering*. Just as a beam of light can be blocked by a wall in its way, wireless signals can be disrupted by buildings, trees, and people. Wireless signals at higher frequencies degrade much more quickly than lower-frequency signals. As soon as a wireless signal's power falls below that of the average background radiation, it becomes impossible to decipher, so high-bandwidth wireless signals generally can't travel very far.

In wires, signals aren't subject to the inverse-square law, so signal power attenuates more gradually. However, signals in traditional twisted-pair copper wires become noisy over distance due to *crosstalk interference* and other factors. Coaxial cables suffer from less noise, but still aren't perfect. Modern fiber optic cables are even better, and have exceptionally low noise. In fiber optic communication systems, most noise comes from imperfections in transmitters and receivers.¹ Still, light beams in fibers attenuate over distance due to interactions with small imperfections in the glass. Signals can travel much further in some channels than in others, but the SNR always increases with distance.

Channel capacity and the Shannon limit

Given a fixed amount of bandwidth and a constant signal-to-noise ratio, there is a theoretical limit to the amount of information throughput a channel can carry. This limit is captured by the Shannon-Hartley theorem, often referred to as the *Shannon limit*. The Shannon limit expresses the maximum information *capacity*, C , of a channel in bits per second. C is a function of B , the bandwidth, S , the power of the signal, and N , the average

¹ A. Demir, "Noise Analysis for Optical Fiber Communication Systems," Institute of Electrical and Electronics Engineers, available at <https://ieeexplore.ieee.org/abstract/document/1257814>.

power of the noise. The relationship S/N is often referred to as the *signal-to-noise ratio*, or SNR. The exact equation is shown below.

$$C = B \log_2 \left(1 + \frac{S}{N} \right)$$

The Shannon–Hartley theorem, describing the theoretical limit to the information capacity of a channel as a function of bandwidth (B), signal power (S) and average noise power (N).

You don't need to understand the math behind the theorem to get the basics: more bandwidth means more capacity, as does a better signal-to-noise ratio. If bandwidth and signal power of a channel are fixed, more noise means less capacity. The Shannon limit is important to understand because it means we can take the physical properties of a medium, like copper wire or fiber optics, and figure out how much capacity we might someday squeeze out of it—even if we can't do it yet.

Generally, the longer the distance a signal has to travel, the weaker the signal power becomes due to attenuation. This reduces the SNR and, according to Shannon's theorem, the total information the signal can carry. Therefore, it's not possible to talk about the capacity of a channel without knowing how far a signal has to go. The channel capacity of 10 yards of cable might be 10 Gb/s, but the capacity of 10 miles of the same cable might only be 5 Mb/s.

To recap: the *bandwidth* and *signal-to-noise ratio* (SNR) of a channel determine the maximum rate of data it can carry. The longer a link needs to be, the worse the channel's SNR will become. Most channels can carry high-capacity signals for short distances, but few can support the same capacity over many miles.

Latency and jitter

Channel capacity is only half the story. The Shannon limit describes how many bits per second a channel can carry, but it says nothing about how fast a bit actually gets from point A to point B. *Latency* is the time it takes for a message to make the trip from one end of a channel to the other. *Jitter* describes variations in latency; it occurs when portions of a signal arrive out of sync from their expected schedule. Think of a video call over the internet. Latency is responsible for the constant small delay between you speaking and the other person registering your voice, while jitter is responsible for glitches, freezes, and other distortions in the stream.

The ultimate lower bound on latency is determined by the speed of light: no signal can travel faster than light in a vacuum. The speed of light limits how fast signals can be transmitted across oceans and continents, but in last-mile connections (the subject of this whitepaper), latency is almost always dominated by the time it takes to process a signal at each end of a channel. For example, the latency between a phone and a 4G LTE

tower a mile away is approximately 9 milliseconds;² however, the radio waves that carry the signal can travel that distance in around 5 *microseconds* (0.005 ms). That means over 99.9% of the latency is incurred by the transmitting and receiving devices.

In low-bandwidth and error-prone channels, messages need to be encoded with layers of error-correcting codes, and signal encoding/decoding can take some time. On the other hand, channels with lots of bandwidth and low error rates can be generated and processed with little latency. Error rates in modern fiber-optic channels are typically very low, and signals can be transmitted and received with minimal delays for processing and error correction.

Jitter occurs when packets sent over a channel are delayed or dropped. Jitter is experienced as spikes in latency: instead of *all* packets being delayed by a fixed amount, *some* packets are delayed, while others arrive on time. For example, even with error correction, some parts of a signal may be dropped entirely, which can cause higher-level protocols like TCP/IP to pause and retransmit old packets. This results in an uneven or “choppy” connection. Latency can be constant and predictable, but jitter is always random. Channels that are subject to jitter may be fine for tasks like downloading large files or streaming video, but will cause noticeable issues with applications like video chat or online gaming.

From channels to networks

Modern, high-speed network links comprise many different parts, with different technologies used to transmit data for different stages of the journey. Data networks, like the internet, use a hierarchical “tree” structure: high-capacity links at the “trunk” carry data from many people across long distances, while lower-capacity links in the “branches” carry a few connections to smaller regions. Eventually, the branching links are subdivided into “leaves” that each link to a single network participant, like a computer or mobile phone.

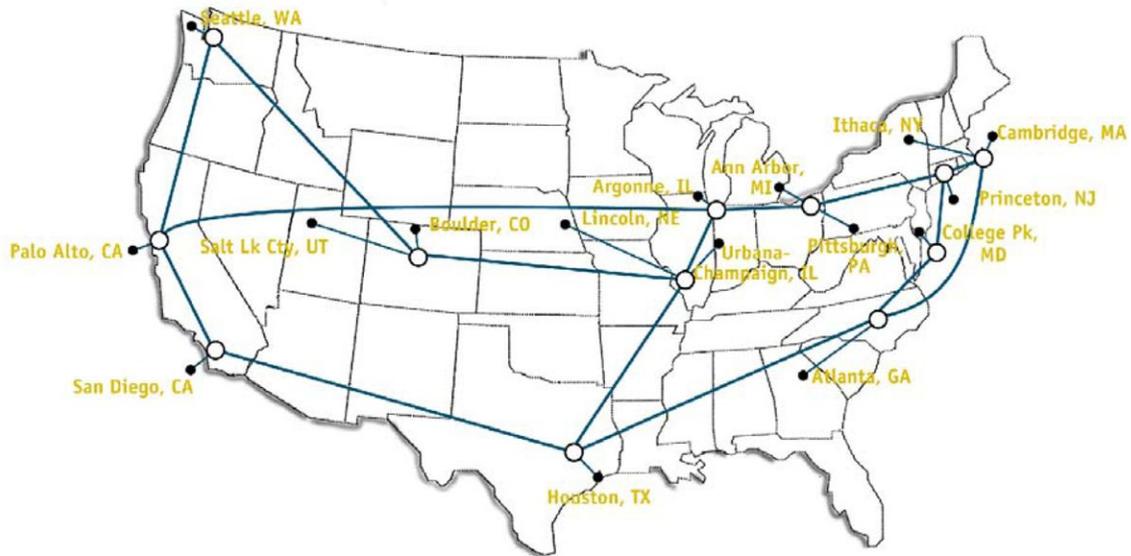
Let’s consider an example. When you connect to Google using a laptop on your home’s WiFi network, the data first travels from your computer to your WiFi router via radio waves in the 2.4 GHz or 5 GHz bands. Next, it travels over ethernet, which probably uses short (<100m) copper wires to carry the data from your router to your modem. If you have cable internet, the signal then travels over a coaxial cable from your house to a small “cabinet” or “node,” a box on the curb that serves a few dozen or few hundred people in your neighborhood. From there, it travels along with your neighbors’ traffic through a fiber to a “cable headend,” the local service center where your cable company operates. The connection from your home to your local cable headend is known as the “last mile” connection.

From there, data from you and all the other customers in your neighborhood travels along one or more higher-capacity connections, using fiber-optic cables, until it reaches

² Wireless One, *LTE Latency Today 9 ms. Down to 2 ms ~2019*, March 17, 2018, available at <http://wirelessone.news/10-r/1007-lte-latency-today-9-ms-down-to-2-ms-2019>

the “backbone” network connection for your region. The backbone carries thousands of connections from one regional subnetwork to another, which could be across the country or across the world. Backbone networks nearly always use high-capacity fiber optic cables, which are, by far, the most effective way to carry high-bandwidth signals over long distances. The backbone connection will carry your data (along with data from thousands of others) to the regional subnetwork where the nearest Google server is located, where it will be routed back down through the “branches” of that network to the “leaf,” a server in a datacenter that will process and respond to your request.

NSFNET T3 Network 1992



A diagram showing the “backbone” of the early Internet. Today, the backbone has many more connections.

With the technical background explained, this whitepaper will now turn to the “last mile” connections that link local subnetworks to individual internet subscribers. While “middle mile” and backbone connections have been systematically converted to fiber-optic cable over the past three decades, last mile connections still use a diverse set of technologies: DSL, DOCSIS, 4G (and soon, 5G) wireless, and fiber-to-the-home. This paper gives a brief overview of the dominant last-mile technologies in use today. It argues that while there are advances to be made in DOCSIS and wireless internet technology, they are not in a position to surpass fiber. In fact, future advancements in other technologies will rely on fiber. Fiber-to-the-home is the best option for reliable, high-throughput, and future-proof last mile connections today.

DOCSIS 3.1 and the Future of Coax

Coaxial cable, or “coax” (pronounced co-*axe*), is the standard conduit for cable TV. It is made up of a core copper wire and an outer copper tube separated by an insulating sheath. The design of coaxial cable makes it much more resistant to “crosstalk” and other noisy interference than traditional twisted-pair copper wiring. Coax can carry much higher-bandwidth signals with less interference than other copper cables, which is why it is preferred to twisted-pair cables for broadband internet.

Although coax has much better resistance to noise than copper alternatives, some noise is still present due to reflections and radio-frequency interference.³ In addition, each coaxial cable has a “cutoff frequency” above which signals become muddled and hard to recover.⁴ Most commercial cables are rated to carry up to a few GHz of bandwidth.⁵ High-powered signals cause more noise, and cables are usually rated for a maximum signal power. Coax also experiences signal attenuation (weakening over distance) due to electrical impedance, and higher-frequency signals suffer from more attenuation.

All of that means trying to send a high-frequency signal over a long distance is a tough proposition. The signal power drops off drastically over distance, but the power at the transmitter can only be raised to a certain point before it starts adding too much noise. As a result, high-throughput signals can only be carried over shorter cables or using *amplifiers* installed along the cable.

The standard used by cable companies to deliver internet service over coax is called DOCSIS (Data Over Cable System Interface Specification). DOCSIS signals are served from a “cable headend,” a station that generates signals and transmits them along cables to subscriber homes. On the other side, modulator-demodulators, or “modems,” allow cable customers to interpret the signals produced by the headend and generate their own digital signals in return. A single cable headend can serve customers up to a few miles away. Older DOCSIS setups sent signals strictly over coax, but modern headends usually drive signals down fiber optic lines to smaller “nodes,” each of which uses coax to serve just a few subscribers. In each node, the signal from the fiber is “split” and sent down coax for the final few meters to subscriber homes. These kinds of deployments are known as “hybrid fiber-cable” (HFC) networks.

³ Radio-frequency interference is usually the greatest source of noise in coaxial cables. Though the design of coax cancels out most noise, electrical resistance in the outer shield can induce noise and holes in the shield allow high-frequency signals to “leak” through. See Howard Johnson & Martin Graham, *High-Speed Signal Propagation: Advanced Black Magic* (Prentice Hall, 2003).

⁴ Above a cable’s cutoff frequency, waves begin to propagate in different “modes” and at different speeds, causing interference and making it much harder to recover a useful signal. Cables with smaller diameters have higher cutoff frequencies, but also have much worse power handling capabilities. See Peter McNeil, *How High is a Coaxial Cables Max Frequency?* See Peter McNeil, *How High is a Coaxial Cable Max Frequency?*, Pasternack Blog (Oct. 11, 2018) available at <https://blog.pasternack.com/coaxial-cable/how-high-is-a-coaxial-cables-max-frequency>.

⁵ Helukabel, *Cable specifications overview*, available at http://biakom.com/pdf/RG-coaxial_cables_Helukabel.pdf.

The latest version of the standard is DOCSIS 3.1.⁶ DOCSIS 3.1 was first deployed in early 2016. By 2019, much of the U.S.'s cable infrastructure had been upgraded from DOCSIS 3.0.⁷ DOCSIS 3.1 uses 1.2 GHz of bandwidth and, in theory, it can support 10 Gb/s download speeds and 1Gb/s upload speeds over a single cable. While these numbers represent the *theoretical* throughputs available to individual subscribers, they do not reflect the reality of DOCSIS performance on the ground. The 10Gb/s maximum is the amount of data that can be sent down a single cable; most deployments use one cable to reach multiple houses, so the total capacity is shared between dozens or hundreds of customers. Furthermore, the maximum speeds can only be reached with “deep fiber” HFC setups, where most of the last mile is fiber and a relatively short length of high-quality coax connects the node to subscribers. Although Comcast finished deploying DOCSIS 3.1 in October 2018,⁸ independent tests from around that time show that it offered average real-world speeds around 100Mb/s down and 15Mb/s up.⁹

The first major drawback of DOCSIS 3.1 is the tremendous discrepancy between upload and download speeds. In the recent past, internet users have demanded much more data capacity for downloads than they have for uploads. Activities like browsing the web and watching videos pull lots of data down from servers without sending much back, so DOCSIS has evolved to prioritize downstream throughput. Most DOCSIS deployments allocate less than 85 MHz of the 1.2 GHz of available bandwidth for upstream service. The 3.1 standard only supports using up to 200 MHz of bandwidth, about 1/6 of the total, for upstream traffic.¹⁰ But usage patterns are changing, and operators expect to see major growth in demand for upstream throughput over the next few years.¹¹ Cable operators will have to upgrade their systems sooner rather than later if they want to keep up with the requirements of modern applications and demand driven by fiber-to-the-home competitors. And the upgrades will involve laying lots of new fiber.¹²

DOCSIS 3.1 deployments also suffer from issues related to latency and jitter. There is a good deal of variation in the quality and conditions of cable networks. Older cable may have higher noise rates or significant attenuation, especially when carrying high-frequency signals that it was not originally intended to handle. To deliver consistent throughputs in the face of these discrepancies, DOCSIS employs sophisticated

⁶ CableLabs, *DOCSIS 3.1 Technology*, available at <https://www.cablelabs.com/technologies/docsis-3-1>.

⁷ Press Release, Comcast, Comcast to Introduce World's First DOCSIS 3.1-Powered Gigabit Internet Service in Atlanta, Chicago, Detroit, Miami, and Nashville (Feb. 2, 2016); See also Tech News Today, *Cable Companies Can Save Money Now That DOCSIS 3.1 Upgrade is Mostly Done* (Jun. 15, 2019), available at <https://latesttechnewsblog.com/2019/06/15/cable-companies-can-save-money-now-that-docsis-3-1-upgrade-is-mostly-done>.

⁸ Daniel Frankel, *Comcast Reaches the Finish Line on DOCSIS 3.1 Deployment*, Multichannel News (Oct. 18, 2018), available at <https://www.multichannel.com/news/comcast-reaches-the-finish-line-on-docsis-3-1-deployment>.

⁹ Speedtest, *United States Fixed Broadband Speedtest Data Q2-Q3 2018*, available at <https://www.speedtest.net/reports/united-states/2018/#fixed>.

¹⁰ John Ulm, *Making Room for D3.1 & FDX*, in SCTE & ISBE Journal of Network Operations, 4, 1, 2018, available at <https://www.scte.org/SCTEDocs/Journals/SCTE-ISBE%20Network%20Operations%20Journal%20N4V1.pdf>.

¹¹ Ayham Al-Banna, Tom Cloonan, and Jeff Howe, *Network Migration Strategies for the Era of DAA, DOCSIS 3.1, and New Kid on the Block... Full Duplex DOCSIS!*, SCTE-ISBE and NCTA, 2017. Available at <https://www.nctatechnicalpapers.com/Paper/2017/2017-network-migration-strategies/download>.

¹² See *supra* 10, table 2 on page 19. (The only viable options for significantly improving upstream capacity involve going “fiber deep” or transitioning to fiber-to-the-home entirely.)

encoding schemes¹³ which offer better robustness at the expense of up to 3.5ms of extra latency.¹⁴ For example, “interleaving” involves scrambling portions of a signal before sending it over the wire, allowing forward error correction to more effectively deal with bursts of noise. This scrambling and unscrambling means that symbols cannot be processed in real time, and interleaving can add milliseconds of latency to the system.¹⁵ Headend operators can choose how to configure their networks: simpler encoding schemes add less guaranteed latency but are worse at correcting for noise, which leads to more dropped packets and jitter. More complex encoding schemes add milliseconds of latency, but deliver more consistent throughput. Furthermore, “media acquisition” protocols in DOCSIS 3.1—which are used to grant individual modems access to upstream traffic on shared cables—add an additional 2–8 ms of latency to the system.¹⁶

The next generation of DOCSIS technologies includes a proposal for “Low Latency DOCSIS” (LLD).¹⁷ LLD would primarily improve latency for certain applications, like video chat or online games, by prioritizing some types of traffic over others at the modem level. While this doesn’t improve average latency, it does offload latency to applications (like downloads or streaming video) where it doesn’t matter as much. LLD will also improve on the media acquisition protocols currently used in DOCSIS 3.1. This change will improve average latency, but it won’t address the delays caused by encoding and decoding traffic. As DOCSIS advances and transmission technologies improve, they will remain subject to tradeoffs: better throughput will only be possible with more complex encoding schemes and over shorter coax cables.

Planned future versions of DOCSIS will support “full duplex” speeds of 10Gb/s for both uploads and downloads, and may use up to 3 GHz of spectrum down the road.¹⁸ The next version of DOCSIS, known as 4.0, is still in early stages of development and will not be standardized until the mid or late 2020’s. In the long term, coax may be able to deliver speeds up to 25 or even 50 Gb/s, but the technology will run up against the Shannon limit sooner rather than later.

¹³ John Downey, *Understanding DOCSIS Data Throughput and How to Increase it*, available at http://piedmontscte.org/resources/DOCSIS_Throughput.doc

¹⁴ In DOCSIS 3.1, the simple Reed–Solomon error correction encoding used for versions 1.0 to 3.0 was replaced with a concatenated Bose, Ray–Chaudhuri, Hocquenghem (BCH) and Low Density Parity Check (LDPC) encoding. This scheme allows operators to push data throughput closer to the Shannon limit at the expense of computational complexity; See Brady S. Volpe & Mike Collins, *It’s All About the FEC: Like a Box of Chocolates*, Broadband Library (May 26, 2018), available at <https://broadbandlibrary.com/fec>.

¹⁵ Errors in DOCSIS systems tend to occur in bursts. Error-correcting encodings are better at dealing with errors that are spread out over time, so operators can “interleave,” or mix up, symbols before they are sent. This increases the effectiveness of error correcting codes at the expense of more latency. See Cisco, *Understanding Data Throughput in a DOCSIS World*, available at <https://www.cisco.com/c/en/us/support/docs/broadband-cable/data-over-cable-service-interface-specifications-docsis/19220-data-thruput-docsis-world-19220.html>.

¹⁶ See Sundaresan White & B. Briscoe, *Low Latency DOCSIS: Technical Overview*, (Feb. 2019), available at <https://tools.ietf.org/id/draft-white-tsvwg-lld-00.html#LLD-white-paper>.

¹⁷ *Id.*

¹⁸ Alan Breznick, *Here Comes DOCSIS 4.0*, LightReading (May 22, 2018), available at <https://www.lightreading.com/cable/docsis/here-comes-docsis-40/d/d-id/743285> (Researchers have begun experimenting with using frequencies up to 3GHz for what will become DOCSIS 4.0, with the goal of having a full specification by the mid to late 2020s. Based on previous standard rollouts, we might expect to see widespread deployment of DOCSIS 4.0 3 to 5 years after that).

One big draw of DOCSIS is that cable companies can use existing infrastructure to continue delivering high-speed broadband. However, in order to serve cable customers with gigabit speeds and beyond, any remaining all-coax networks will need to be replaced with HFC networks and fiber nodes in HFC networks will have to be moved even closer to subscriber homes.¹⁹ Cable operators will need to increase their node counts by a factor of 10 or 20,²⁰ and the “last mile” will become closer to a “last meter.” In addition, it’s unclear whether the aging coax already in the ground will be able to support extended frequencies up to 3 GHz.²¹ Old coax may need to be decommissioned and replaced in order to take full advantage of DOCSIS 4.0.

To summarize: high-bandwidth broadband over coax is possible, but we are approaching the limits of what the technology can do. Current-generation DOCSIS technology suffers from relatively high latencies and huge discrepancies between upstream and downstream throughputs. Next-gen improvements to cable internet can mitigate these issues, but will require decommissioning miles of old coax and running fiber closer to subscriber homes. And while future versions of the technology will improve on the relatively high latencies of DOCSIS 3.1, high-throughput DOCSIS will continue to be subject to more latency than pure fiber.

5G and the Future of Wireless

Wireless broadband solves a fundamentally different problem than wireline technologies like cable and fiber. Wireline technologies deliver service to a fixed point, like a home or business. Wireless delivers data service to mobile devices through the air, and it’s the only way to offer flexible broadband service to large public areas. For the past two decades, wireless and wireline broadband technologies have coexisted harmoniously in the internet ecosystem. However, some industry representatives have suggested that the fifth generation of cellular broadband, known as 5G, will be able to compete directly with wireline broadband options or replace it altogether.²² This section will describe how wireless broadband works, and examine how it compares to wireline technologies as a last-mile link. It will argue that for the vast majority of users, wireline internet will remain the better option for fixed-point broadband.

¹⁹ Many providers have already begun reaching closer to homes with fiber to support the DOCSIS 3.1 rollout. In addition, proposed technologies like full duplex DOCSIS will require providers to upgrade their amplifiers or reach close enough with fiber to remove them altogether. Brian Santo, *Cable Nodes Becoming a Chokepoint*, LightReading (Dec. 5, 2016), available at <https://www.lightreading.com/cable/ccap-next-gen-nets/cable-nodes-becoming-a-choke-point/d/d-id/728754>; See also Daniel Frankle, *Cox Set to Take Fiber to the Node, Deploy DOCSIS 3.1*, FierceVideo (May 23, 2016), available at <https://www.fiercevideo.com/cable/cox-set-to-take-fiber-to-node-deploy-docsis-3-1>

²⁰ See *supra* 10

²¹ Philip Dampler, *Cable’s DOCSIS 4.0 – Symmetrical Broadband Coming*, Stop The Cap! (Jun. 25, 2019), available at <https://stopthecap.com/2019/06/25/cables-docsis-4-0-symmetrical-broadband-coming>.

²² See <https://www.lifewire.com/5g-internet-wifi-4156280> and <https://knowledge.wharton.upenn.edu/article/the-push-for-5g/>

Wireless broadband systems are significantly different from cable and other wireline systems. For one, wireless broadband doesn't need to be deployed to each customer; each wireless base station serves whoever happens to be in its vicinity. In addition, wireless signals degrade in power over distance much more quickly than wired signals. While a single cable headend can serve customers for many miles in every direction, cellular base stations in populated areas are typically placed no more than a mile apart.²³

Wireless internet deployments are also subject to constraints that wired systems are not. Low-frequency wireless signals, like AM/FM radio and broadcast TV, are able to pass through trees, buildings, and miles of open air without a problem. Higher-frequency bands have more bandwidth and generally carry more information. However, higher-frequency signals are also more susceptible to absorption and scattering, which limits how far they can be transmitted. While 2.4 GHz WiFi can pass through brick walls in a house, 5GHz WiFi has more trouble, and is often unable to reach across multiple rooms. The next generation of WiFi technology, known as WiGig, utilizes frequency bands as high as 60GHz.²⁴ At that frequency, signals are almost completely disrupted by walls and furniture, so 60GHz routers will work best for nearby, line-of-sight communication.

The current generation of cellular internet technologies is known as "4G" (for "4th generation"). 4G operates on frequencies between the 700 MHz and 2.6 GHz bands, which can serve devices up to a few hundred meters away in urban areas and up to a few miles away in rural areas. Technically, 4G systems are supposed to be capable of serving 1 Gb/s download speeds to low-mobility devices (like phones in the hands of pedestrians).²⁵ However, in the real world, most carriers offer speeds from 10 to 50 Mb/s down and 3 to 20 Mb/s up.²⁶ Tests of 4G networks in the US have measured latencies around 50ms, with the "air latency" link between the tower and the device accounting for a significant portion of that.²⁷

5G promises improvements over 4G in both throughput and latency. For long-distance links, 5G will use the same spectrum currently used by 4G, between 700 MHz and 4 GHz. Improvements to antennas and encoding technology will allow carriers to make better

²³ Bernard Prkić, *Understanding Small-Cell Wireless Backhaul*, ElectronicDesign (Apr. 3, 2014), available at <https://www.electronicdesign.com/communications/understanding-small-cell-wireless-backhaul> (In suburban areas, cell sites are typically installed 1-2 miles apart, while in urban areas, they may only be ¼ mile apart due to population density and to overcome interference caused by buildings).

²⁴ Wi-Fi Alliance, *Wi-Fi Certified WiGig: Multi-gigabit, Low Latency Connectivity*, available at <https://www.wi-fi.org/discover-wi-fi/wi-fi-certified-wigig>.

²⁵ International Telecommunications Union, *Requirements related to technical performance for IMT-Advanced radio interface(s)*, (2008), available at <http://www.itu.int/pub/R-REP-M.2134-2008/en>.

²⁶ 2019 tests found that Verizon, the fastest U.S. carrier, provides average speeds of 53 Mb/s down and 17.5 Mb/s up; Cricket, the slowest tested network, achieves 6.8 Mb/s down and 5.8 Mb/s up. See Tom's Guide, "Fastest Wireless Network 2019: It's Not Even Close." See Tom's Guide, *Fastest Wireless Network 2019: It's Not Even Close*, available at <https://www.tomsguide.com/us/best-mobile-network,review-2942.html>.

²⁷ See *supra* 2. Also Mehdi Daoudi, *There's No Avoiding Network Latency on 4G*, Catchpoint (Jan 15, 2014), available at <https://blog.catchpoint.com/2014/01/15/theres-no-avoiding-network-latency-on-4g> (A 2014 test found average pings on 4g networks to be around 55ms, compared to an average of 22ms on wireline broadband).

use of the same spectrum.²⁸ In terms of throughput, long-distance 5G may not be a massive step forward: tests of sub-6GHz 5G deployments have found it to be capable of a few hundred Mb/s in the best case, only slightly better than the most advanced 4G LTE systems.²⁹

In addition to re-using 4G spectrum, 5G will support “millimeter wave (mmWave)” frequencies at 26 GHz and above. Higher frequency channels are attractive because they offer more usable bandwidth, and can therefore support higher maximum throughputs. Using mmWave spectrum, 5G transmitters will be able to provide much better transfer speeds, maxing out between 1 and 10 Gb/s under optimal conditions. But since mmWave signals are so much higher frequency than traditional cellular signals, they suffer much greater absorption and scattering. Millimeter wave signals cannot pass through most walls, thick foliage, or even inclement weather without encountering significant interference. They also lose power much faster, even in clear conditions, than sub-6GHz signals.³⁰ That means mmWave won’t work well for outdoor-to-indoor communication. Early adopters of mmWave in US cities have reported needing to do the “5G shuffle”—physically dancing around 5G transmitters—in order to take advantage of gigabit speeds.³¹ As a result, mmWave transmitters will work more like WiFi, providing service to small, open areas, rather than drop-in replacements for 4G.

5G also promises to improve on the latency of 4G. While providers have promised air latencies between 1 and 4 ms, these numbers will only be available with mmWave spectrum.³² Real-world tests have found that the sub-6GHz 5G equipment being shipped today has air latencies between 9 and 12 ms, which is comparable to advanced 4G technology.³³

²⁸ One example of an improvement is “massive MIMO (Multiple Input Multiple Output)” technology. MIMO allows base stations to use multiple antennas to transmit over a greater portion of the available spectrum at once. See Qualcomm, *How 5G Massive MIMO Transforms Your Mobile Experiences*, OnQ Blog (Jun. 20, 2019), available at <https://www.qualcomm.com/news/onq/2019/06/20/how-5g-massive-mimo-transforms-your-mobile-experiences>.

²⁹ In a CNet experiment from July 2019, the best sub-6GHz deployment was from SK telecom in Seoul, which achieved peak download speeds of 618 Mb/s. In the US, the top tested deployment was in Dallas, where the Sprint 5G network achieved 484 Mb/s. See Jessica Dolcourt, *We Ran 5G Speed Tests on Verizon, AT&T, EE, and more: Here’s What We Found*, CNet (Jul. 3, 2019), available at <https://www.cnet.com/features/we-ran-5g-speed-tests-on-verizon-at-t-ee-and-more-heres-what-we-found>.

³⁰ See FCC Office of Engineering and Technology, *Millimeter Wave Propagation: Spectrum Management Implications* (1997), available at https://transition.fcc.gov/Bureaus/Engineering_Technology/Documents/bulletins/oet70/oet70a.pdf.

³¹ TechRadar journalists tested Verizon’s 5G deployment in Chicago in May 2019. They were able to achieve super-gigabit download speeds by physically moving around the mmWave transmitter. See Matt Swider, *5G Speed Test: 1.4 Gbps in Chicago, but Only if You Do the ‘5G Shuffle,’* Techradar (May 19, 2019), available at <https://www.techradar.com/news/5g-speed-test>.

³² Ronan McLaughlin, *5G Low Latency Requirement*, Broadband Library (May, 25, 2019), available at <https://broadbandlibrary.com/5g-low-latency-requirements>.

³³ Jon Brodtkin, *AT&T’s 5G Trials Produce Gigabit Speeds and 9ms Latency*, ArsTechnica (Apr. 11, 2018), available at <https://arstechnica.com/information-technology/2018/04/atts-5g-trials-produce-gigabit-speeds-and-9ms-latency> (An AT&T test of mmWave 5G in Waco, Texas found “latency rates of 9-12 ms.” This likely refers to the air latency between the device and the tower, which matches up with Verizon’s 5G deployments

What about 6G and beyond? As time goes on, cell providers will likely find ways to squeeze more throughput out of the usable long-range frequencies below 6GHz. However, the bandwidth available at these frequencies is limited, and background noise will always be present. Cellular providers will soon run into the Shannon limit for wireless channels. Furthermore, as applications for mobile devices advance, they will likely demand higher sustained data rates than before, which will put greater strain on mobile networks. Since each cell tower has to serve all devices in an area using the same limited bandwidth, as more devices clamor for more data, the average available throughput will suffer. More base stations can be built to accommodate some of the increased demand, but the stations will still need to share a limited amount of spectrum. Speeds for everyone are likely to improve, but not as much as the lab-tested scenarios would suggest.

To summarize, 5G is a big step forward, but it is not a panacea. Millimeter-wave 5G will use more bandwidth to serve fewer devices in a smaller area, so it should be able to deliver true gigabit speeds. It should be able to deliver last-hop latencies that are comparable to, or even better than, fiber-to-the-home. However, mmWave deployments will require running fiber-optic cables to individual buildings in order to be useful.³⁴ In other words, the most exciting parts of 5G will supplement, rather than replace, fiber-to-the-home.

Fiber Today and in the 21st Century

Fiber-optic cables are long, extremely thin, and carefully crafted strands of glass that can “guide” beams of light from one end to the other. Fiber optics can carry light over hundreds of miles without allowing the light to scatter or disperse. Although the mode of transmission is different in fiber than in coax, the principle is the same: both fiber-optic and coaxial cables guide electromagnetic waves and protect them from interference in transit.

Fiber carries much higher-frequency signals than coax does. DOCSIS 3.1 uses frequencies up to 1.2 gigahertz, but common fiber-optic cables carry light in the infrared spectrum between 200 and 350 *terahertz*.³⁵ A typical fiber-optic cable has around 10,000 times more usable bandwidth than a typical coaxial cable. Furthermore, fiber-optic cables are

elsewhere); Wireless One, *Latency 30 ms at Verizon 5G* (Apr. 04, 2019), available at <http://wirelessone.news/10-r/1368-5g-latency-30-ms-at-verizon> (Real-world latency from device to server remains around 30ms).

³⁴ Gemalto, *Introducing 5G Networks - Characteristics and Usages*, available at <https://www.gemalto.com/brochures-site/download-site/Documents/tel-5G-networks-QandA.pdf> (Both the bandwidth and latency improvements that 5G promises assume fiber-optic links directly to base stations).

³⁵ Fiber optic cables carry light wavelengths between 850 and 1620 nm. Not all wavelength bands are viable due to absorption, and different protocols use different bands; PON protocols use wavelengths between 1400 and 1610 nm for transmission. See Alice Gui, *From O to L: The Evolution of Optical Wavelength Bands*, Cable Solutions (Oct. 13, 2015), available at <http://www.cables-solutions.com/from-o-to-l-the-evolution-of-optical-wavelength-bands.html>.

much less susceptible to interference and noise than coax or wireless channels. Beams of light do not interfere with other electromagnetic waves in the same way that radio-frequency signals do, so fiber isn't vulnerable to crosstalk or radio-frequency leakage like coax is. The main limiting factor for fiber is attenuation, or power lost over distance. Even modern fiber isn't perfectly transparent. Over the course of long distances, light is absorbed by tiny imperfections in the glass, causing the beam to become dimmer. Therefore, fiber cables spanning extremely long distances (like oceans) must have repeaters installed to periodically boost the signal.

Today, fiber is often used to carry Internet signals through every part of the network *except* the last mile. We've already discussed how fiber carries data around the internet backbone, how it brings broadband from cable headends to curbside "nodes" in hybrid fiber-cable DOCSIS deployments, and how fiber will connect to base stations in 5G networks. When fiber-optic cables are used to deliver service directly to a subscriber's residence, it's known as "fiber-to-the-home" (FTTH). The most common FTTH architecture is the Passive Optical Network (PON), a design in which signal is driven down a single fiber and "split" using a series of passive lenses to serve individual subscribers. There are competing standards for last-mile fiber deployments, including the ITU-T's NG-PON2³⁶ and the IEEE's 10G-EPON,³⁷ but most of them use the same basic PON architecture.

We are nowhere near able to take advantage of fiber's full potential for last-mile connections. The huge amount of bandwidth available through fiber, and the minimal noise added during transmission, mean that the Shannon limit to fiber-optic channels tends to be extraordinarily high. In a lab setting, researchers have been able to achieve data rates upwards of 100 Tb/s over many kilometers in a single, standard fiber,³⁸ and it's likely that we'll see further improvements in the years to come. But transmitters and receivers capable of more than 1 Tb/s are still quite expensive. For now, they are only used in enterprise settings and the internet backbone.

A typical fiber-to-the-home deployment today has symmetrical upload and download speeds around 1 Gb/s, though currently adopted PON standards support symmetrical speeds up to 10Gb/s.³⁹ As technology continues to develop, better transmitters will become cheaper and more efficient, and providers will be able to upgrade existing fiber deployments without any changes to the fiber itself. Once fiber is laid, its capacity can be upgraded by orders of magnitude just by changing the transmitters at each end. Fiber-optic cables are typically designed for a lifetime of at least 25 years, though they

³⁶ Gigabit-capable Passive Optical Network (GPON) is the common name for the G.984 standard by the ITU-T, introduced in 2003. It has since been superseded by G.987, aka XG-PON, and by G.989, aka NG-PON2. See International Telecommunications Union, *40-Gigabit-capable passive optical networks (NG-PON2): General requirements*, available at <https://www.itu.int/rec/T-REC-G.989.1-201303-1>.

³⁷ Ethernet Passive Optical Network (EPON) was first standardized by the IEEE in 2004; updated versions of the standard that support 10 Gb/s, known as 10G-EPON, and beyond have since been standardized. See IEEE P802.3av Task Force, *10Gb/s Ethernet Passive Optical Network*, available at <http://www.ieee802.org/3/av>.

³⁸ Jeff Hecht, *Ultrafast Fibre Optics Set New Speed Record*, *NewScientist* (Apr. 19, 2011), available at <https://www.newscientist.com/article/mg21028095-500-ultrafast-fibre-optics-set-new-speed-record>.

³⁹ Both the ITU-T's NG-PON2 standard and the IEEE's 10G-EPON standard support symmetrical connections of 10 Gb/s or better, *supra* notes 36 and 37.

can, and frequently do, last much longer.⁴⁰ And as long as the cables themselves remain sound, FTTH connections are all but future-proof.

The fact that many PON architectures have fully symmetrical data speeds gives them a significant advantage over DOCSIS. As we discussed previously, DOCSIS 3.1 uses a small portion of spectrum for upstream traffic, and only allows for 1 Gb/s of upload throughput to be shared between all customers in a service group.⁴¹ Meanwhile, NG-PON2 allocates 4 different channels of 10Gb/s each for upstream data, yielding 40Gb/s of total upstream throughput to be shared among the customers on a network terminal.⁴² Latency is another area where fiber has a major advantage. In DOCSIS 3.1, upstream bandwidth allocation adds 2–8 ms of latency.⁴³ FTTH protocols need to address the upstream allocation problem too, but the excessive upstream bandwidth available in fiber-optic systems makes it easier to deal with. Testing has shown that dynamic bandwidth allocation in PON systems adds less than a millisecond of latency.⁴⁴

Furthermore, as described above, coax is more susceptible to noise than fiber, especially when carrying high-frequency signals. To overcome that noise, DOCSIS transmitters need to use ever-more complex error-correcting encoding schemes. Encoding and decoding symbols takes time at each end of the cable, and it limits how quickly data can travel. On the other hand, signals driven over fiber contain very little noise. GPON and other fiber protocols transmit data with less overhead for error correction.⁴⁵ As a result, total last-mile latency in GPON FTTH channels can be specified below 1.5 ms, even for links up to 20km.⁴⁶ In addition, because fiber-optic channels experience fewer dropped packets than coax channels do, they suffer from less jitter. Fiber provides a smoother, more real-time internet experience than any competing wireline technologies. This makes fiber the best choice for applications where responsiveness is critical, like voice-over-IP, video chat, remote-controlled robotics, and virtual reality.

In short, fiber is the superior medium for carrying fixed broadband by almost every metric: available bandwidth, SNR, theoretical capacity, real-world throughput, latency, and jitter. Furthermore, fiber cables can be installed now and upgraded for decades to come, while most existing coax infrastructure will likely need to be replaced within the

⁴⁰ David Stockton, *4 Factors That Influence How Long Your Fiber Network Will Last*, PPC Blog, available at <https://www.ppc-online.com/blog/4-factors-that-influence-how-long-your-fiber-network-will-last> (Cracks and other flaws in fiber optics, introduced during manufacturing or deployment, are exacerbated over time and can lead to failure after several years. For correctly installed tier-1 fibers, the probability of a given km of fiber failing on its own within 20–40 years is approximately 1 in 100,000. However, the most common cause of failure is construction or “dig-ups” that occur after the fiber has been laid. In lieu of these kinds of failures, fiber-optic deployments can last for many decades).

⁴¹ See *supra* 6

⁴² See *supra* 36

⁴³ See *supra* 16 for information about upstream allocation latency in DOCSIS.

⁴⁴ Pavel Sikora et al., *Efficiency Tests of DBA Algorithms in XG-PON*, MDPI Electronics 2019, 8, 762; available at <https://www.mdpi.com/2079-9292/8/7/762/pdf>.

⁴⁵ GPON systems have configurable error correction, and some systems may not require error-correcting encoding at all. See Calix Resource Center, available at <https://www.calix.com/content/calix/en/site-prod/library-html/systems-products/b-series/system-operation/b6-user-docs/release8-0/ug/index.htm?toc45430275.htm?52773.htm>.

⁴⁶ International Telecommunications Union, *Gigabit-capable passive optical networks (GPON): General characteristics*, available at <https://www.itu.int/rec/T-REC-G.984.1>.

next few years in order to keep up with consumer demand. While 5G is a promising upgrade over 4G, long-range wireless broadband cannot outperform fiber as a last-mile link to homes and businesses. In highly populated areas, mmWave 5G will be a supplement to, not a replacement for, fiber-to-the-home. In rural areas, attempting to install enough fiber to enough base stations to provide full mmWave coverage makes less sense than to simply run wireline service to each home.⁴⁷ And to top it off, future upgrades to both DOCSIS and wireless broadband will require laying many miles of new fiber. As a result, civic planners looking ahead should invest in last-mile fiber infrastructure today. Fiber-to-the-home is the best option to serve most Americans with high-speed, low-latency broadband now, and it will remain so for the foreseeable future.

⁴⁷ Jon Brodtkin, *Millimeter-wave 5G will never scale beyond dense urban areas, T-Mobile says*, ArsTechnica (Apr. 22, 2019), available at <https://arstechnica.com/information-technology/2019/04/millimeter-wave-5g-will-never-scale-beyond-dense-urban-areas-t-mobile-says/>