A Manager's Guide to The Internet and Telecommunications

Draft version last modified: May 31, 2010 – comments welcome john.gallaugher@bc.edu

INTRODUCTION

There's all sorts of hidden magic happening whenever you connect to the Internet. But what really makes it possible for you to connect to servers half way around the world in just a fraction of a second? Knowing this is flat-out fascinating stuff, but it's also critically important for today's manager to have at least a working knowledge of how the Internet functions.

That's because the Internet is a platform of possibilities and a business enabler. Understanding how the Internet and networking works can help you brainstorm new products and services and understand roadblocks that might limit turning your ideas into reality. Marketing professionals that know how the Internet reaches consumers have a better understanding of how technologies can be used to find and target customers. Finance firms that rely on trading speed to move billions in the blink of an eye need to master Internet infrastructure to avoid being swept aside by more nimble market movers. And knowing how the Internet works helps all managers understand where their firms are vulnerable. In most industries today, if your network goes down then you might as well shut your doors and go home, it's nearly impossible to get anything done if you can't get online. Managers who know the net are prepared to take the appropriate steps to secure their firms and keep their organization constantly connected.

INTERNET 101: UNDERSTANDING HOW THE INTERNET WORKS

LEARNING OBJECTIVES:
After studying this section you should be able to:
1. Describe how the technologies of the Internet combine to answer the questions: What are you looking for? Where is it? And how do we get there?
2. Interpret a URL, understand what hosts and domains are, describe how domain registration works, describe cybersquatting, and give examples of conditions that constitute a valid and invalid domain-related trademark dispute.
3. Describe certain aspects of the Internet infrastructure that are fault tolerant and supports load balancing.
4. Discuss the role of hosts, domains, IP addresses, and the DNS in making the Internet work.

The Internet is a network of networks – millions of them, actually. If the network at your university, your employer, or in your home has Internet access, it connects to an *Internet Service Provider* (ISP). Many (but not all) ISPs are big telecommunications companies, like Verizon, Comcast, and AT&T. These providers connect to one another, exchanging traffic, ensuring your messages can get to any other computer that's online and willing to communicate with you.

The Internet has no center and no one owns it. That's a good thing. The Internet was designed to be redundant and fault tolerant – meaning that if one network, connecting wire, or server stops working, everything else should keep on running. Rising from military research and work at educational institutions dating as far back as the 1960s, the Internet really took off in the 1990s, when graphical web browsing was invented, and much of the Internet's operating infrastructure was transitioned to be supported by private firms rather than government grants.
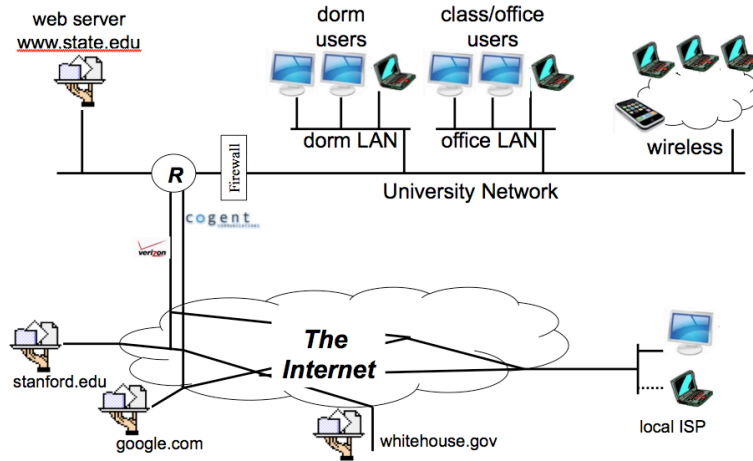

Figure X: The Internet is a network of networks, and these networks are connected together. In the diagram above, the "state.edu" campus network is connected to the other networks of the Internet via two ISPs, cogent and Verizon.

Enough history – let's see how it all works! If you want to communicate with another computer on the Internet then your computer needs to know the answer to three questions: What are you looking for? Where is it? And how do we get there? The computers and software that make up Internet infrastructure can help provide the answers. Let's look at how it all comes together.

The URL: "What Are You Looking For?"

When you type an address into a web browser (sometimes called a URL for uniform resource locator), you're telling your browser what you're looking for. Figure X describes how to read a typical URL.

Figure X: Anatomy of a Web Address. The URL displayed says "use the web (http) to find a host server named "www" in the "nytimes.com" network, look in the "tech" directory, and access the "indx.html" file.

The http:// you see at the start of most web addresses stands for hypertext transfer protocol. A *protocol* is a set of rules for communication – sort of like grammar and vocabulary in a language like English; The http protocol defines how web browser and web servers communicate and is designed to be hardware and operating system independent. It doesn't matter if messages come from a PC, a Mac, a huge mainframe, or a pocket-sized smartphone; if a device speaks to another using a common protocol then it will be heard and understood!

The Internet supports lots of different applications, and many of these applications use their own application transfer protocol to communicate with each other. The server that holds your e-mail uses something called *SMTP* or simple mail transfer protocol to exchange mail with other e-mail servers throughout the world. *FTP*, or file transfer protocol is used for – you guessed it – file transfer. FTP is how most web developers upload the web pages, graphics, and other files for their websites. Even the web uses different protocols. Surf to a bank website or the credit-card entry page at an Internet retailer and the http in front of your URL will probably change to https (the 's' is for secure). That means that communications between your browser and server will be encrypted for safe transmission. The beauty of the Internet infrastructure is that any savvy entrepreneur can create a new application that rides on top of the Internet.

Hosts and Domain Names

The next part of the URL in our diagram holds the host and domain name. Think of the domain name as the name of the network you're trying to connect to, and think of the host as the computer you're looking for on that network.

Many domains have lots of different hosts. For example, Yahoo's main website is served from the host named "www" (at the address www.yahoo.com), but Yahoo also runs other hosts including those named "finance" (finance.yahoo.com), "sports" (sports.yahoo.com), and "games" (games.yahoo.com).

SIDEBAR: Host and Domain Names: A Bit More Complex Than That

While it's useful to think of a host as a single compute, popular websites often have several computers that work together to share the load for incoming requests. Assigning several computers to a host name offers *load balancing* and *fault tolerance*, helping to ensure that all visits to a popular site like www.google.com won't overload a single computer, or that Google doesn't go down if just one computer fails.

It's also possible for a single computer to have several host names. This might be the case if a firm were hosting several web sites on a single piece of computing hardware.

Some domains are also further broken down into subdomains – many times to represent smaller networks or sub-groups within a larger organization. For example the address www.rhsmith.umd.edu is a University of Maryland address with a host "www" located in the subdomain rhsmith for the Robert H. Smith School of Business. International URLs might also include a second-level domain classification scheme. British URLs use this scheme, with the BBC carrying the commercial (.co) designation www.bbc.co.uk, and the University of Oxford carrying the academic (.ac) designation www.ox.ac.uk.

> You can actually go 127 levels deep in assigning subdomains, but that wouldn't make it easy on those who have to type in a URL that big.

Most websites are configured to load a default host, so you can often eliminate the host name if you want to go to the most popular host on a site (the default host is almost always named www). Another tip - most browsers will automatically add the "http://" for you, too.

Host and domain names are case-insensitive, so you can use a combination of upper and lower case letters and you'll still get to your destination.

SIDEBAR: I Want My Own Domain

You can stake your domain name claim in cyberspace by going through a firm called a domain name registrar. You don't really buy a domain name, you simply pay a registrar for the right to use that name, with the rights renewable over time. While some registrars simply register domain names, others act as *web hosting services* that are able to run your web site on their Internet-connected servers for a fee.

Registrars throughout the world are accredited by *ICANN* (the Internet Corporation for Assigning Names and Numbers), a non-profit governance and standards setting body. Each registrar may be granted the ability to register domain names in one or more of the net's generic top level domains (gTLDs), such as ".com", ".net", or ".org". There are dozens of registrars that can register .com domain names, the most popular gTLD.

Some generic top level domain names, like ".com" have no restrictions on use, others limit registration. For example, .edu is restricted to U.S. accredited post-secondary institutions. ICANN has also announced plans to allow organizations to sponsor their own top level domains (such as .berlin, or .coke).

There are also separate agencies that handle the over 250 different two character country code top level domains (e.g. .uk for the United Kingdom, .jp for Japan). Servers or organizations generally don't need to be housed within a country in order to use a country code as part of its domain name, leading to a number of creatively named web sites. The URL shortening site bit.ly uses Libya's ".ly" top level domain; many physicians are partial to Moldova's code (.md); and the tiny pacific island nation of Tuvulu might not have a single broadcast TV station, but that doesn't stop it from licensing its country code to firms that want a .tv domain name[1]. Recent standards also allow domain names in languages that use non-Latin alphabets such as Arabic and Russian.

Domain name registration is handled on a first-come, first-served basis and all registrars share registration data to ensure that no two firms gain rights to the same name. Startups often sport wacky names, in part because so many domains with common words and phrases are already registered to others. While some domain names are held by legitimate businesses, others are registered by investors hoping to resell a name's rights.

Trade in generic domain names can be lucrative. For example, the "Insure.com" domain was sold to QuinStreet for $16 million in Fall 2009[2]. But knowingly registering a domain name in order to profit from someone else's firm name or trademark is known as *cybersquatting* and that's illegal. The U.S. has passed anti-cybersquatting legislation, and ICANN has a mechanism for domain name dispute resolution that can reach across boarders. Try to extort money by holding a domain name that's identical to (or in some cases, even similar to) a well-known trademark holder and you could be stripped of your domain name and even fined.

Courts and dispute resolution authorities will sometimes allow a domain that uses the trademark of another organization if it is perceived to have legitimate, non-exploitive reasons for doing so. For example, the now defunct site verizonreallysucks.com was registered as a protest against the networking giant and was considered fair use since owners didn't try to extort money from the telecom giant[3]. However, the courts allowed the owner of the PETA trademark (the organization People for the Ethical Treatment of Animals) to claim the domain name peta.org from original registrant, who had been using that domain to host a site called "People Eating Tasty Animals"[4].

Trying to predict how authorities will rule can be difficult. The rocker Sting's name was thought to be too generic to deserve the rights to Sting.com, but Madonna was able to take back her domain name (for the record, Sting now owns Sting.com)[5]. Apple executive Jonathan Ive was denied the right to reclaim domain names incorporating his own name, but that had been registered by another party and without his consent. The publicity-shy design guru wasn't considered enough of a public figure to warrant protection[6]. And sometimes disputing parties can come to an agreement outside of court or ICANN's dispute resolution mechanisms. When Canadian teenager Michael Rowe registered a site for his part-time web design business, a firm south of the boarder took notice of his domain name - mikerowesoft.com. The two parties eventually settled in a deal that swapped the domain for an Xbox and a trip to the Microsoft Research Tech Fest[7].

Path and Filename

Look to the right of the top level domain and you might see a slash followed by either a path name, a file name, or both. If a web address has a path and file name, the path maps to a folder location where the file is stored on the server; and the file is the name of the file you're looking for.

Most web pages end in .html, indicating they are in *hypertext markup language*. While http helps browsers and servers communicate, html is the language used to create and format (render) web pages. A file doesn't need to be .html; web servers can serve up just about any type of file: Acrobat documents (ending in .pdf), PowerPoint documents (.pptx), Word docs (.docx), JPEG graphic images (.jpeg), and as we'll see in the security chapter, even malware programs that attack your PC. At some web addresses the file displays content for every visitor, and at others (like amazon.com), a file will contain programs that run on the web server to generate custom content just for you.

You don't always type a path or file name as part of a web address, but there's always a file lurking behind the scenes. A web address without a file name will load content from a default page. For example, when you visit 'google.com', Google automatically pulls up a page called "index.html", a file that contains the web page that displays the Google logo, the text entry field, the "Google Search" button, etc. You might not see it, but it's there.

Butterfingers beware! Path and file names are case sensitive – amazon.com/books is considered to be different from amazon.com/BOOKS. Mistype your capital letters after the domain name and you might get a 404 error (the very unfriendly webserver error code that means document not found).

IP Addresses and the Domain Name System: "Where is it? And how do we get there?"

The IP Address

If you want to communicate , then you need to have a way for people to find and reach you. Houses and businesses have street addresses, and telephones have phone numbers. Every device connected to the Internet has an identifying address, too – it's called an *IP (internet protocol) address*.

A device gets its IP address from whichever organization is currently connecting it to the Internet. Connect to a laptop at your university and your school will assign the laptop's IP address. Connect at a hotel, and the hotel's Internet Service Provider lends your laptop an IP address. Laptops and other end-user machines might get a different IP address each time they connect, but the IP addresses of servers rarely change. It's OK if you use different IP addressees during different online sessions because services like e-mail and Facebook identify you by your username and password. The IP address simply tells the computers that you're communicating with where they can find you right now. IP addresses can also be used to identify a user's physical location, to tailor search results, and to customize advertising. See the Google Chapter to learn more.

IP addresses are usually displayed as a string of four numbers between zero and 255, separated by three periods. Want to know which IP address your smartphone or computer is using? Visit a website like ip-adress.com (one 'd'), whatismyipaddress.com, or ipchicken.com.

SIDEBAR: THE INTERNET IS ALMOST FULL

If you do the math, four combinations of 0-255 gives you a little over 4 billion possible IP addresses. Four billion sounds like a lot, but the number of devices connecting to the Internet is exploding! Internet access is now baked into smart phones, tablets, televisions, DVD players, video game consoles, utility meters, thermostats, appliances, picture frames, and more. Another problem - a big chunk of existing addresses weren't allocated efficiently, and these can't be easily reclaimed from the corporations, universities, and other organizations that initially received them. All this means that we're running out of IP addresses. Experts differ on when ICANN will have no more numbers to dole out, but most believe that time will come by 2012, if not sooner[8].

There are some schemes to help delay the impact of this IP address drought. For example, a technique known as NAT (for network address translation) uses a gateway that allows multiple devices to share a single IP address. But NAT slows down Internet access and is complex, cumbersome, and expensive to administer[9].

The only long-term solution is to shift to a new IP scheme and fortunately one was developed more than a decade ago. IPv6 increases the possible address space from the 2^{32} (4,294,967,296) addresses used in the current system (called IPv4) to a new theoretical limit of 2^{128} addresses. 2^{128} is a really big number - bigger than 34 with 37 zeros after it.

But not all the news is good. Unfortunately, IPv6 isn't backward compatible with IPv4 and the transition to the new standard has been painfully slow. This gives us the equivalent of many islands of IPv6 in a sea of IPv4, with translation between the two schemes happening when these networks come together. While most modern hardware and operating systems providers now support IPv6, converting a network to IPv6 currently involves a lot of cost with little short-term benefit[10]. Upgrading may take years, and is likely to result in rollout problems. David Conrad, a general manager at the authority that doles out IP

---

addresses, has said "I suspect we are actually beyond a reasonable time frame where there won't be some disruption. It's just a question of how much"[11].

Some organizations have stepped up to try to hasten transition. Google, has made most of its services IPv6 accessible, the U.S. government has mandated IPv6 support for most agencies, China has spurred conversion within its borders, and Comcast and Verizon have major IPv6 rollouts under way. While the transition will be slow, when wide scale deployment does arrive, IPv6 will offer other benefits, including potentially improving the speed, reliability, and security of the Internet.

The DNS – The Internet's Phonebook

You can actually type an IP address of a website into a web browser and that page will show up. But that doesn't help users much because four sets of numbers are really hard to remember.

This is where the *domain name service* (or DNS) comes in. The domain name service is a distributed database that looks up the host and domain names that you enter, and returns the actual IP address for the computer that you want to communicate with. It's kind of like a big, hierarchical set of phone books capable of finding web servers, e-mail servers, and more. These "phone books" are called *nameservers* – and when they work together to create the DNS, they can get you anywhere you need to go online.
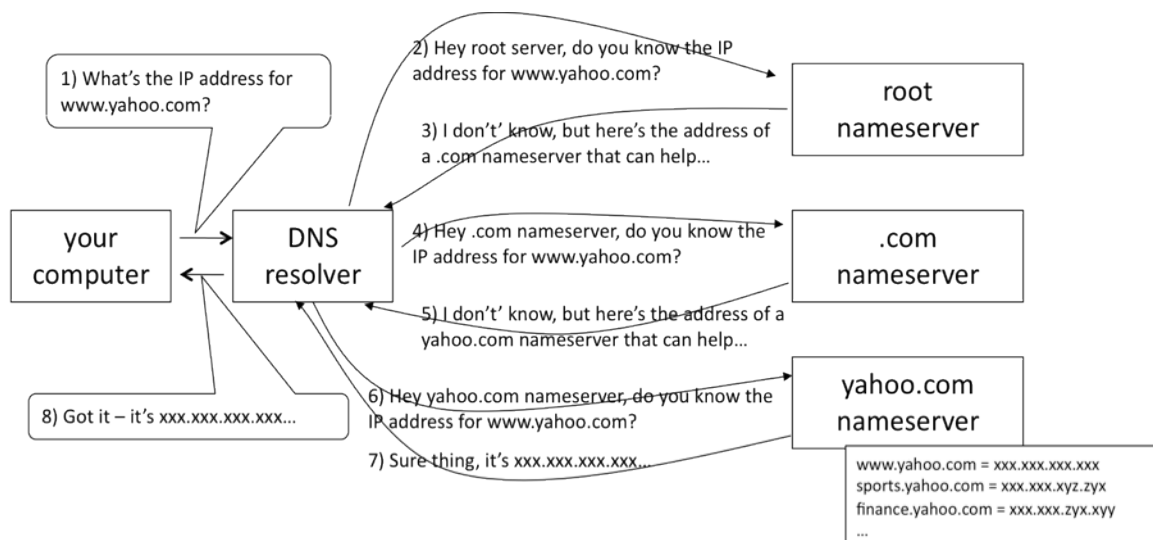

Figure X: When your computer needs to find the IP address for a host/domain name, it sends a message to a DNS resolver, which looks this up starting at the root. Once the lookup has taken place, that IP address can be saved in a holding space called a cache, to speed future lookups.

In order to get a sense of how the DNS works, let's imagine that you type www.yahoo.com into a web browser. Your computer doesn't know where to find that address, but when your computer connected to the network, it learned where to find a service on the network (called a DNS resolver). The DNS resolver can look things up in the "phone book" that is the DNS. The resolver doesn't know everything, but it does know where to start a lookup that will eventually give you the address you're looking for. If this is the first time anyone on that network has tried to find "www.yahoo.com", the resolver will contact one of

thirteen root nameservers. The roots act as a lookup starting place., they don't have one big list, but they can point you to a nameserver for the next level – that'd be one of the .com nameservers for our example. The .com nameserver can then find one of the yahoo.com nameservers. And a yahoo.com nameserver can respond back to the resolver with the IP address for www.yahoo.com, and the resolver passes that back to your computer. Once your computer knows that IP address, it's then ready to communicate directly with www.yahoo.com. The Yahoo.com nameservers includes IP addresses for all of Yahoo's public sites: www.yahoo.com, games.yahoo.com, sports.yahoo.com, finance.yahoo.com, and more.

The system also remembers what it's done so the next time you need an IP address of a host you've already looked up, your computer can pull this out of a storage space called a *cache*, avoiding all those nameserver visits. Caches are periodically refreshed, to ensure that you're always pointed to where you need to go.

Distributing IP address lookup this way makes sense. It avoids having one huge, hard-to-maintain, ever-changing list. Firms add and remove hosts on their own networks by just updating entries in their own nameserver. And it allows host IP addresses to change easily, too. Moving your web server offsite to a hosting provider? Just update your nameserver with the new IP address at the hosting provider and the world will invisibly find that new IP address on the new network by using the same old, familiar host and domain name combination. The DNS is also fault tolerant – meaning that if one nameserver goes down, the rest of the service can function. There are exact copies at each level and the system is smart enough to move on to another nameserver if its first choice isn't responding.

SIDEBAR: But What if the DNS Gets Hacked?

A hacked DNS would be a disaster! Think about it. If bad guys could change which websites load when you type in a host and domain name, they could redirect you to imposter websites that look like a bank or e-commerce retailer, but are really setup to harvest passwords and credit card data.

This exact scenario played out when the DNS of NET Virtua, a Brazilian Internet Service Provider, was hacked via a technique called *DNS cache poisoning*. Cache poisoning exploits a hole in DNS software, redirecting users to sites they didn't request. The Brazilian DNS hack redirected NET Virtua users wishing to visit the Brazilian bank Bradesco to fraudulent websites that attempted to steal passwords and install malware. The hack impacted about one percent of the bank's customers before the attack was discovered[12].

The exploit showed the importance of paying attention to security updates. A few months earlier a group that *Wired* Magazine referred to as "A Secret Geek A-Team"[13] had developed a software update that corrected the DNS poisoning exploit used against NET Virtua, but administrators at the Brazilian Internet Service Provider failed to update their software so the hackers got in. An additional upgrade to DNS system, known as DNSSEC, promises to further limit the likelihood of cache poisoning, but it may take years for the new standards to be rolled out everywhere[14].

KEY TAKEAWAYS:

- The Internet is a network of networks. Internet service providers connect with one another to share traffic, enabling any Internet-connected device to communicate with any other.
- URLs may list the application protocol, host name, domain name, path, and file, in that order. Path and file names are case sensitive.
- A domain name represents an organization. Hosts are public services offered by that organization. Hosts are often thought of as a single computer, although many computers can operate under a single host name and many hosts can also be run off of a single computer.
- You don't buy a domain name, but can register it, paying for a renewable right to use that domain name. Domains need to be registered within a generic top level domain such as .com or .org, or within a two-character country code top level domain such as .uk, .ly, or .md.
- Registering a domain that uses someone else's trademark in an attempt to extract financial gain is considered cybersquatting. The U.S. and other nations have anticybersquatting laws, and ICANN has a dispute resolution system that can overturn domain name claims if a registrant is considered to be cybersquatting.
- Every device connected to the Internet has an IP address. These addresses are assigned by the organization that connects that user to the Internet. An IP address may be assigned temporarily, for use only during that online session.
- We're running out of IP addresses. The current scheme IPv4 is being replaced by a IPv6, a scheme that will give us many more addresses and additional feature benefits, but which is not backward compatible with the IPv4 standard. Transitioning to IPv6 will be costly and take time.
- The domain name system is a distributed, fault tolerant system that uses nameservesr to map host/domain name combinations to IP addresses.

QUESTIONS & EXERCISES:
1. Find the web page for your school's information systems department. What is the URL that gets you to this page? Label the host name, domain name, path, and file for this URL. Are there additional subdomains? If so, indicate them, as well.
2. Go to a registrar and see if someone has registered your first or last name as a domain name. If so, what's hosted at that domain? If not, would you consider reserving your domain name? Why or why not?
3. Investigate cases of domain name dispute resolution. Find a case that you find especially interesting. Who were the parties involved? How was the issue resolved? Do you agree with the decision?
4. Describe how the DNS is fault tolerant and promotes load balancing. Give examples of other types of information systems that might need to be fault tolerant and offer load balancing. Why?
5. Research DNS poisoning online. List a case other than the one mentioned in this chapter, where DNS poisoning took place. Which network was poisoned, who were the victims, and how did hackers exploit the poisoned system? Could this exploit have been stopped? How? Whose responsibility is it to stop these kinds of attacks?
6. Why is the switch from IPv4 to IPv6 so difficult? What key principles, discussed in prior chapters, are slowing migration to the new standard?

GETTING WHERE YOU'RE GOING

LEARNING OBJECTIVES:
After studying this section you should be able to:
1. Understand the layers that make up the Internet – application protocol, transmission control protocol, and internet protocol – and describe why each is important.
2. Discuss the benefits of Internet architecture in general, and TCP/IP in particular.
3. Name applications that should use TCP, and others that might use UDP.
4. Understand what a router does, and the role these devices play in networking.
5. Conduct a traceroute and discuss output, demonstrating how Internet interconnections work in getting messages from point to point.
6. Appreciate why mastery of Internet infrastructure is critical to modern finance, and be able to discuss the risks in automated trading systems.
7. Describe VoIP, and contrast circuit vs. packet switching, along with organizational benefits and limitations of each.

TCP/IP – The Internet's Secret Sauce

OK, we know how to read a web address, we know that every device connected to the net needs an IP address, and we know that the DNS can look at a web address and find the IP address of the machine that you want to talk to. But how does a web page, e-mail, or iTunes download actually get from a remote computer to your desktop?

For our next part of the Internet journey we'll learn about two additional protocols – TCP and IP. These protocols are often written as TCP/IP and pronounced by reading all five letters in a row 'T', 'C', 'P', 'I', 'P' (sometimes they're also referred to as the *Internet Protocol Suite*). TCP and IP are built into any device that a user would use to connect to the Internet – from handhelds to desktops to supercomputers – and together TCP/IP make internetworking happen.
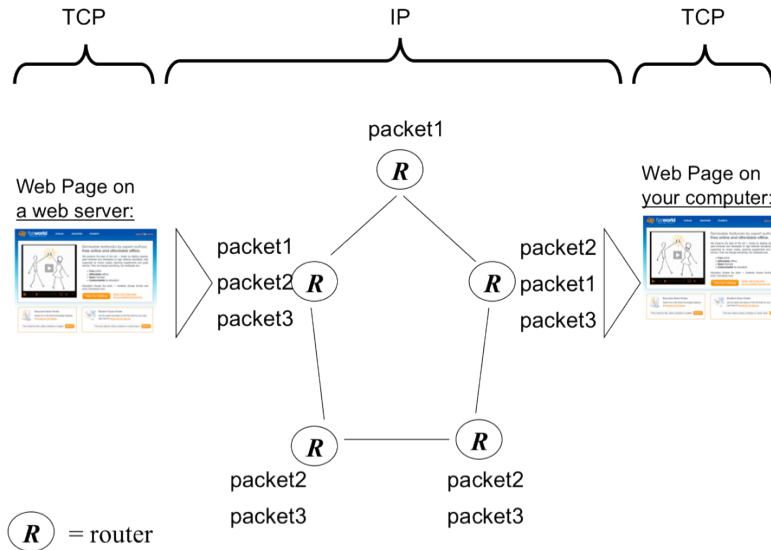
Figure X: TCP/IP in action. In this example, a server on the left sends a web page to the user on the right. The application (the web server) passes the contents of the page to TCP (which is built into the server's operating system). TCP slices the web page into packets. Then IP takes over, forwarding packets from router to router across the Internet until it arrives at the user's PC. Packets sometimes take different routes, and occasionally arrive out of order. TCP running on the receiving system on the right checks that all packets have arrived, requests a re-send of damaged or lost packets, puts them in the right order, and sends a perfect, exact copy of the web page to your browser.

TCP and IP operate below http and the other application transfer protocols mentioned earlier. TCP (*transmission control protocol*) works its magic at the start end endpoint of the trip – on both your computer and on the destination computer you're communicating with. Let's say a web server wants to send you a large web page. The web server application hands the web page it wants to send to its own version of TCP. TCP then slices the web page up into smaller chunks of data called packets (sometimes called datagrams). The packets are like little envelopes containing part of the entire transmission – they're labeled with a destination address (where it's going) and a source address (where it came from). Now we'll leave TCP for a second, because TCP on the web server then hands those packets off to the second half of our dynamic duo, IP.

It's the job of IP (*Internet Protocol*) to route the packets to their final destination, and those packets might have to travel over several networks to get to where they're going. The relay work is done via special computers called *routers*, and these routers speak to each other and to other computers using Internet Protocol (since routers are connected to the Internet, they have IP addresses, too. Some are even named). Every computer on the Internet is connected to a router, and all routers are connected to at least one (and usually more) other routers, linking up the networks that make up the Internet.

Routers don't have perfect, end-to-end information on all points in the Internet, but they do talk to each other all the time, so a router has a pretty good idea of where to send a packet to get it closer to where it needs to end up. This chatter between the routers also keeps the Internet decentralized and fault tolerant. Even if one path out of a router goes down (a networking cable gets cut, a router breaks, the power to a router goes out), as long as there's another connection out of that router, then your packet will get forwarded.

Networks fail, so good, fault-tolerant network design involves having alternate paths into and out of a network.

Once packets are received by the destination computer (your computer in our example), that machine's version of TCP kicks in. TCP checks to make sure that it has all of the packets it was sent, it makes sure that no packets were damaged or corrupted, requests replacement packets (if needed), and then puts the packets in the correct order, passing a perfect copy of your transmission up to the program you're communicating with (an e-mail server, web server, etc.).

This progression - application to TCP at the source (slice it up) to IP (for forwarding among routers) to TCP at the destination (put the transmission back together & make sure it's perfect) to application – takes place in both directions, starting at the server for messages coming to you, and starting on your computer when you're sending messages to another computer.

SIDEBAR: UDP: TCP's Faster, Less Reliable Sibling

TCP is a perfectionist and that's what you want for web transmissions, e-mail, and application downloads. But sometimes we're willing to sacrifice perfection for speed. You'd make this sacrifice for streaming media applications like Windows Media Player, Real Player, Internet voice chat and video conferencing. Having to wait to make sure each packet is perfectly sent would otherwise lead to awkward pauses that interrupt real-time listening. It'd be better to just grab the packets as they come and play them, even if they have minor errors. Packets are small enough that if one packet doesn't arrive, you can ignore it and move on to the next without too much quality disruption. A protocol called UDP (User Datagram Protocol) does exactly this, working as a TCP stand-in when you've got the need for speed, and are willing to sacrifice quality. If you've ever watched a web video or had a web-based phone call and the quality got sketchy, it's probably because there were packet problems, but UDP kept on chugging, making the 'get it fast' instead of 'get it perfect' tradeoff.

SIDEBAR: VoIP: When Phone Calls are Just Another Internet Application

The increasing speed and reliability of the Internet means that applications such as Internet phone calls (referred to as *VoIP* or Voice over IP) are becoming more reliable. That doesn't just mean that Skype becomes a more viable alternative for consumer landline and mobile phone calls, it's also good news for many businesses, governments, and non-profits, too!

Many large organizations maintain two networks – one for data, another for *POTS* (plain old telephone service). Maintaining two networks is expensive, and while conventional phone calls are usually of a higher quality than their Internet counterparts, POTS equipment is also inefficient. Old phone systems use a technology called *circuit switching*. A 'circuit' is a dedicated connection between two entities. When you have a POTS phone call, a circuit is open, dedicating a specific amount of capacity between you and the party on the other end. You're using that 'circuit' regardless of whether you're talking, or not. Pause between words, put someone on hold, and the circuit is still in use. Anyone who has ever tried to make a phone call at a busy time (say early morning on Mother's Day, or at midnight on New Year's Eve) and received an "all circuits are busy" recording has experienced congestion on an inefficient circuit-switched phone network.

But unlike circuit-switched counterparts, Internet networks are packet switched networks and these can be more efficient. Since we can slice conversations up into packets, we can squeeze them into smaller

---

spaces. If there are pauses in a conversation or someone's on hold, applications don't hold up the network. And that creates an opportunity to use the network's available capacity for other users. The trade off is one that swaps circuit-switching quality of service (QoS) with packet switching's efficiency and cost savings. Try to have a VoIP call when there's too much other traffic on a portion of the network and your call quality will drop. But packet switching quality is getting much better. Networking standards are now offering special features, such as 'packet prioritization', that can allow voice packets to gain delivery priority over packets for applications like e-mail, where a slight delay is OK.

When voice is digitized, "telephone service" simply becomes another application that sits on top of the Internet, like the Web, e-mail, or FTP. VoIP calls between remote offices can save long distance charges. And when the phone system becomes a computer application, you can do a lot more. Well implemented VoIP systems allow users browser access to their voice mail inbox, one-click video conferencing and call forwarding, and other features, but you'll still have a phone number, just like with POTS.

What Connects the Routers and Computers?

Routers are connected together, either via cables or wirelessly. A cable connecting a computer in a home or office is probably copper (likely what's usually called an Ethernet cable), with transmissions sent through the copper via electricity. Long haul cables, those that carry lots of data over long distances, are usually fiber optic lines – glass lined cables that transmit light (light is faster and travels farther distances than electricity, but fiber optic networking equipment is more expensive than the copper/electricity kind). Wireless transmission can happen via Wi-Fi (for shorter distances), or cell phone tower or satellite over longer distances. But the beauty of the Internet Protocol Suite (TCP/IP) is that it doesn't matter what the actual transmission media is. As long as your routing equipment can connect any two networks, and as long as that equipment "speaks" IP, then you can be part of the Internet.

In reality, your messages likely transfer via lots of different transmission media to get to their final destination. If you use a laptop connected via Wi-Fi, then that wireless connection finds a base station, usually within about 300 feet. That base station is probably connected to a local area network (LAN) via a copper cable. And your firm or college may connect to fast, long-haul portions of the Internet via fiber optic cables provided by that firm's Internet Service Provider (ISP).

Most big organizations have multiple ISPs for redundancy, providing multiple paths in and out of a network. This is so that if a network connection provided by one firm goes down, say an errant backhoe cuts a cable, other connections can route around the problem.

In the U.S. (and in most deregulated telecommunications markets) Internet Service Providers come in all sizes, from smaller regional players to sprawling international firms. When different ISPs connect their networking equipment together to share traffic, it's called *peering*. Peering usually takes place at neutral sites called *Internet Exchange Points* (IXPs), although some firms also have private peering points. Carriers usually don't charge one another for peering. Instead, 'the money is made' in the ISP business by charging the end-points in a network – the customer organizations and end-users that an ISP connects

to the Internet. Competition among carriers helps keep prices down, quality high, and innovation moving forward.

SIDEBAR: Finance has a Need for Speed

When many folks think of Wall Street trading, they think of the open outcry pit at the New York Stock Exchange. But human traders are just too slow for many of the most active trading firms. Over half of all United States stock trades and a quarter of worldwide currency trades now happen via programs that make trading decisions without any human intervention.[15] There are many names for this automated, data-driven frontier of finance - algorithmic trading, black-box trading, or high-frequency trading. And while firms specializing in automated, high-frequency trading represent only about two percent of the trading firms operating in the United States, they account for about three quarters of all U.S. equity trading volume.[16]

Programmers lie at the heart of modern finance. "A geek who writes code — those guys are now the valuable guys" says the former head of markets systems at Fidelity Investments, and that rare breed of top programmer can make "tens of millions of dollars" developing these systems[17]. Such systems leverage data mining and other model-building techniques to crunch massive volumes of data and discover exploitable market patterns. Models are then run against real-time data and executed the instant a trading opportunity is detected (for more on how data is gathered and models are built, see the chapter on The Data Asset).

Winning with these systems means being quick – very quick. Suffer delay (what techies call *latency*) and you may have missed your opportunity to pounce on a signal or market imperfection. To cut latency, many trading firms are moving their servers out of their own data centers and into *colocation facilities*. These facilities act as storage places where a firm's servers get super-fast connections as close to the action as possible. And by renting space in a "*colo*", a firm gets someone else to manage the electrical and cooling issues, often providing more robust power backup and lower energy costs than a firm might get on its own.

Equinix, a major publicly-traded IXP and colocation firm with facilities worldwide, has added a growing number of high-frequency trading firms to a roster of customers that also include e-commerce, Internet, software, and telecom companies. In Northern New Jersey alone (the location of many of the servers where "Wall Street" trading takes place), Equinix hosts some 18 exchanges and trading platforms as well as the NYSE Secure Financial Transaction Infrastructure (SFTI) access node.

Less than a decade ago, 80 milliseconds was acceptably low latency, but now trading firms are pushing below 1 millisecond into microseconds[18]. So it's pretty clear that understanding how the Internet works, and how to best exploit it, is of fundamental and strategic importance to those seeking a career in finance. But also recognize that this kind of automated trading comes with risks. Systems that run on their own can move many billions in the blink of an eye, and the actions of one system may cascade, triggering actions by others.

The Spring 2010 "Flash Crash" resulted in a nearly 1,000-point freefall in the Dow Jones Industrial Index, it's biggest intra-day drop ever. Those black boxes can be mysterious – weeks after the May 6th event, experts were still parsing through trading records, trying to unearth how the flash crash happened[19]. Regulators and lawmakers recognize they now need to understand technology, telecommunications, and its broader impact on society so that they can create platforms that fuel growth without putting the economy at risk.

SIDEBAR: Watching the Packet Path via traceroute

Want to see how packets bounce from router to router as they travel around the Internet? Check out a tool called traceroute. Traceroute repeatedly sends a cluster of three packets starting at the first router connected to a computer, then the next, and so on, building out the path that packets take to their destination.

Traceroute is built into all major desktop operating systems (Windows, Macs, Linux), and several websites will run traceroute between locations (traceroute.org and visualroute.visualware.com are great places to explore).

The message below shows a traceroute performed between Irish firm VistaTEC and Boston College. At first it looks like a bunch of gibberish, but if we look closely, we can decipher what's going on.

```
Traceroute to www.bc.edu (136.167.2.220), 30 hops max, 38 byte packets
10  vlan120.switch.deg.vistatec.ie (85.159.16.65)  0.758 ms  1.141 ms  2.189 ms
2  ge0-1.router.deg.vistatec.ie (85.159.16.25)  4.109 ms  0.561 ms  0.485 ms
3  xe-0-1-0-119.dub20.ip4.tinet.net (77.67.66.213)  0.698 ms  0.734 ms  0.691 ms
4  xe-10-1-0.lon11.ip4.tinet.net (89.149.186.197)  11.290 ms  11.335 ms  11.300 ms
5  te7-6.mpd02.lon01.atlas.cogentco.com (130.117.15.49)  11.496 ms  11.197 ms  11.454 ms
6  te0-2-0-1.mpd21.jfk02.atlas.cogentco.com (66.28.4.189)  85.687 ms  85.627 ms  85.685 ms
7  te2-2.mpd01.bos01.atlas.cogentco.com (154.54.6.1)  233.730 ms  91.406 ms  91.368 ms
8  te4-2.ccr01.orh01.atlas.cogentco.com (66.28.4.222)  92.498 ms  92.615 ms  92.457 ms
9  38.104.218.10 (38.104.218.10)  94.491 ms  94.458 ms  94.253 ms
10  136.167.9.226 (136.167.9.226)  94.816 ms  94.475 ms  94.586 ms
```

The table above shows ten hops, starting at a domain in vistatec.ie and ending in 136.167.9.226 (the table doesn't say this, but all IP addresses starting with 136.167 are Boston College addresses). The three groups of numbers at the end of three lines shows the time (in milliseconds) of three packets sent out to test that hop of our journey. These numbers might be interesting for network administrators trying to diagnose speed issues, but we'll ignore them and focus on how packets get from point to point.

At the start of each line is the name of the computer or router that is relaying packets for that leg of the journey. Sometimes routers are named, sometimes they're just IP addresses, and when they are named we can what network a packet is on by looking at the domain name. The first hop out of vistatec.ie (hop #3) shows a router at tinet.net – that must be the name of VistaTEC's Internet Service Provider.

Sometimes routers names suggest their locations (oftentimes they use the same three character abbreviations you'd see in airports). Look closely at the hosts in hops #3 through #7. The sub-domains dub20, lon11, lon01, jfk02, and bos01 suggest the packets are going from Dublin, then east to London, then west to New York City (JFK), then north to Boston. That's a long way to travel in a fraction of a second!

Hop #4 is at tinet.net, but hop #5 is at cogentco.com (look them up online and you'll find out that cogentco.com, like tienet.net, is also an ISP). That suggests that between those hops peering is taking place and traffic is handed off from carrier to carrier.

Hop #8 is still cogentco.com, but it's not clear who the unnamed router in #9, 38.104.218.10, belongs to. We can use the Internet to sleuth that out, too. Search the Internet for the phrase "IP Address Lookup" and you'll find a bunch of tools to track down the organization that 'owns' an IP address. Using the tool at whatismyip.com, I found that this number is registered to PSI Net, which is now part of cogentco.com. Routing paths, ISPs, and peering all revealed via traceroute. You've just performed a sort of network 'CAT Scan' and looked into the veins and arteries that make up a portion of the Internet. Pretty cool!

If you try out traceroute on your own, be aware that not all routers and networks are traceroute friendly. It's possible that as your trace hits some hops along the way (particularly at the start or end of your

journey), three "*" characters will show up at the end of each line instead of the numbers indicating packet speed. This indicates that traceroute has timed out on that hop. Some networks block traceroute because hackers have used the tool to probe a network to figure out how to attack an organization. Most of the time, though, the hops between the networks at either end of a traceroute (the steps involving all of the ISPs and their routers) are visible.

Traceroute can be a neat way to explore how the Internet works, and reinforce the topics we've just learned. Search for traceroute tools online or browse the Internet for details on how to use the traceroute built into your computer.

SIDEBAR: There's Another Internet?

If you're a student at a large research university there's a good chance that your school is part of Internet2. Internet2 is a research network created by a consortium of research, academic, industry, and government firms. These organizations have collectively set up a high-performance network running at speeds of up to 100Gbps, to support and experiment with demanding applications. Examples include high-quality video conferencing, high-reliability, high-bandwidth imaging for the medical field, and applications that share huge datasets among researchers.

If your university is an Internet2 member and you're communicating with another computer that's part of the Internet2 consortium, then your organization's routers are smart enough to route traffic through the super-fast Internet2 backbone. If that's the case, you're likely already using Internt2 without even knowing it!

KEY TAKEAWAYS:
- TCP/IP or the Internet Protocol Suite, helps get perfect copies of internet transmissions from one location to another. TCP works on the end, breaking up transmissions up into manageable packets at the start and putting them back together while checking quality at the end. IP works in the middle, routing packets to their destination.
- Routers are special computing devices that forward packets from one location to the next. Routers are typically connected with more than one outbound path, so that in case one path becomes unavailable, an alternate path can be used.
- UDP is a replacement for TCP, used when it makes sense to sacrifice packet quality for delivery speed. It's often used for media streaming.
- TCP/IP doesn't care about the transition media. This allows networks of different types – copper, fiber, wireless – to participate in the Internet.
- The plug and play nature of the Internet – the ability to swap in new applications, protocols, and media – give the network tremendous flexibility.
- Decentralization, fault tolerance, and redundancy help keep the network open and reliable.
- VoIP allows voice and phone systems to become an application travelling over the Internet. This is allowing many firms to save money on phone calls and through the elimination of old, inefficient circuit-switched networks. As Internet applications, VoIP phone systems can also have additional features that circuit switched networks lack. The primary limitation of many VoIP systems is quality of service.
- Many firms in the finance industry have developed automated trading models that analyze data and execute trades without human intervention. Speeds substantially less than one second may be vital to capitalizing on market opportunities, so firms are

increasingly moving equipment into collocation facilities that provide high speed connectivity to other trading systems.

QUESTIONS & EXERCISES:
1. How can the Internet consist of networks of such physically different transmission media – cable, fiber, and wireless?
2. What is the difference between TCP and UDP. Why would you use one vs. the other?
3. Would you recommend a VoIP phone system to your firm or University? Why or why not? What are the advantages? What are the disadvantages? Can you think of possible concerns or benefits not mentioned in this section? Research these concerns online and share your finding with your instructor.
4. What are the risks in the kinds of automated trading systems described in this section? Conduct research and find an example of where these systems have caused problems for firms and/or the broader market. What can be done to prevent such problems? Whose responsibility is this?
5. Search the Internet for a traceroute tool, or look online to figure out how to use the traceroute command built into your PC. Run a three or more traceroutes to different firms at different locations around the world. List the number of ISPs that show up in the trace. Circle the areas where peering occurs. Do some of the 'hops' time out with "*" values returned? If so, why do you think that happened?
6. Find out if your school or employer is an Internet2 member. If it is, run traceroutes to schools that are and are not members of Internet2. What differences do you see in the results?

LAST MILE: FASTER SPEED, BROADER ACCESS

LEARNING OBJECTIVES:
After studying this section you should be able to:
1. Understand the last mile problem, and be able to discuss the pros and cons of various broadband technologies including DSL, cable, fiber, and various wireless offerings.
2. Describe 3G and 4G systems, listing major technologies and their backers.
3. Understand the issue of net neutrality and put forth arguments supporting or criticizing the concept.

The Internet Backbone, is made of fiber optic lines that carry data traffic over long distances. Those lines are pretty speedy. In fact, several backbone providers, including AT&T and Verizon, are rolling out infrastructure with 100 Gbps transmission speeds (that's enough to transmit a 2 hour high-definition movie in about 8 seconds)[20]. But when considering overall network speed remember *Amdahl's Law* – a system's speed is determined by its slowest component[21]. More often than not the bottleneck isn't the backbone, it's the so-called *last-mile*, the connections that customers use to get online.

High-speed last mile technologies are often referred to as *Broadband Internet Access* (or just *broadband*). What qualifies as broadband varies – in 2009 the FCC redefined broadband as having a minimum speed of 768 Kbps (roughly 14 times the speed of those

old 56Kbps modems). Other agencies worldwide may have different definitions. But one thing is clear, a new generation of bandwidth-demanding services requires more capacity. As we increasingly consume Internet services like HD streaming, real-time gaming, video conferencing, and music downloads, we are, in fact, becoming a bunch of voracious, bit-craving gluttons.

With the pivotal role the United States has played in the creation of the Internet, and in pioneering software, hardware, and telecommunications industries, you might expect the U.S. to lead the world in last-mile broadband access. Not even close. A recent study ranked the U.S. 26[th] in download speeds[22], while others have ranked the U.S. far behind in speed, availability, and price[23].

Sounds grim, but help is on the way. A range of technologies and firms are upgrading infrastructure and developing new systems that will increase capacity not just in the U.S., but worldwide. Here's an overview of some of the major technologies that can be used to speed the Internet's last mile.

SIDEBAR: Understanding Bandwidth

When folks talk about *bandwidth* they're referring to data transmission speeds. Bandwidth is often expressed in bits per second or bps. Prefixes letters associated with multiples of bps are the same as the prefixes we mentioned in the Moore's Law chapter when discussing storage capacity in bytes: Kbps = thousand bits (or kilobits)/second, Mbps = million bits (or megabits)/second, Gbps = billion bits (or gigabits)/second (or terabit), Tbps = trillion bits (or terabits)/second.

Remember, there are eight bits in a byte, and one byte is a single character. One mega*byte* is roughly equivalent to one digital book, 45 seconds of music, or 20 seconds of medium-quality video[24]. But you can't just divide by 8 to estimate how many bits you'll need to transfer. When a file or other transmission is sliced into packets (usually of no more than about 1500 bytes), there's some overhead added. Those packets 'wrap' data chunks in an envelope surrounded by source and destination addressing and other important information.

Here are some rough demand requirements for streaming media. For streaming audio like Pandora you'd need at least 150Kbps for acceptable regular quality, and at least 300Kbps for high quality[25]. For streaming video (via Netflix), at a minimum you'd need 1.5 Mbps, but 3.0 Mbps will ensure decent video and audio. For what Netflix calls HD streaming, you'll need a minimum of 5 Mbps, but would likely want 8 Mbps or more to ensure the highest quality video and audio[26].

Cable Broadband

Roughly 90 percent of U.S. homes are serviced by a cable provider, each capable of using a thick copper wire to offer broadband access. That wire (called a *coaxial cable* or *coax*) has shielding that reduces electrical interference, allowing cable signals to travel longer distances without degrading and with less chance of interference than conventional telephone equipment.

One potential weakness of cable technology lies in the fact that most residential providers use a system that requires customers to share bandwidth with neighbors. If the guy next door is a BitTorrent-using bandwidth hog, your traffic could suffer[27].

Cable is fast and it's getting faster. Many cable firms are rolling out a new technology called Docsis 3.0. that offers speeds up to and exceeding 50 Mbps (previous high-end speeds were about 16Mbps and often much less than that). Cable firms are also creating so-called fiber/copper hybrids that run higher speed fiber optic lines into neighborhoods, then use lower-cost but still relatively high-speed copper infrastructure over short distances to homes.[28]. Those are fast networks, but also very expensive to build, since cable firms are laying entirely new lines into neighborhoods instead of leveraging the infrastructure that they've already got in place.

DSL – Phone Company Copper

Digital Subscriber Line (*DSL*) technology uses the copper wire the phone company has already run into most homes. Even as customers worldwide are dropping their 'land line' phone numbers, the wires used to provide this infrastructure can still be used for broadband.

DSL speeds vary depending on the technology deployed. Worldwide speeds may range from 7Mbps to as much as 100Mbps (albeit over very short distances)[29]. The Achilles heel of the technology lies in the fact that DSL uses standard copper telephone wiring. These lines lack the shielding used by cable, so signals begin to degrade the further you are from the connecting equipment in telephone company offices. Speeds drop off significantly at less than two miles from a central office or DSL hub. Go four miles out and the technology is unusable. Some DSL providers are also using a hybrid fiber/copper system, but as with cable's copper/hybrids, this is expensive to build out.

The super-speedy DSL implementations that are popular in Europe and Asia work because foreign cities are densely populated, and so many high-value customers can be accessed over short distances. In South Korea, for example, half the population lives in apartments, and most of those customers live in and around Seoul. This density also impacts costs – since so many people live in apartments, foreign carriers run fewer lines to reach customers, digging up less ground or stringing wires across fewer telephone poles. Their U.S. counterparts by contrast need to reach a customer base sprawled across the suburbs, so U.S. firms have much higher infrastructure costs. [30]

There's another company with copper, electricity-carrying cables coming into your home – the electrical utility. BPL or broadband over power line technology has been available for years, however there are few deployments because it is considered to be pricer and less practical than alternatives[31].

Fiber – A Light-filled Glass Pipe to your Doorstep

FTTH or Fiber to the Home is the fastest last mile technology around. It also works over long distances. Verizon's FIOS technology boasts 50 Mbps download speeds, but has tested network upgrades that increase speeds by over 6 times that[32]. The problem with fiber is that unlike cable or DSL copper, fiber to the home networks weren't already in place. That means firms had to build their own fiber networks from scratch.

The cost of this build out can be enormous. Verizon, for example, has spent over $23 billion on its FTTH infrastructure. However, most experts think the upgrade was critical. Verizon has copper into millions of homes, but U.S. DSL is uncompetitive. Verizon's residential landline business was dying as users switch to mobile phone numbers, and while mobile is growing, Verizon Wireless is a joint venture with UK's Vodaphone, not a wholly owned firm. That means it shares wireless unit profits with its partner. With FIOS, Verizon now offers pay television, competing with cable's core product, it offers some of the fastest home broadband services anywhere, and it gets to keep everything it earns.

In 2010, Google also announced plans to bring fiber to the home. Google deems its effort an experiment – it's more interested in learning how developers and users take advantage ultra-high speed fiber to the home (what kinds of apps are created and used, how does usage and time spent online change), rather than becoming a nationwide ISP itself. Google says it will investigate ways to build and operate networks less expensively and plans to share findings with others. The Google network will be "open", allowing other service providers to use Google's infrastructure to resell services to consumers. The firm has pledged to bring speeds of 1Gbps at competitive prices to at least 50,000 and potentially as many as 500,000 homes. Over 1,100 U.S, communities applied to be part of the Google experimental fiber network[33].

Wireless

Mobile wireless service from cell phone access providers is delivered via cell towers. While these providers don't need to build out a residential wired infrastructure, they still need to secure space for cell towers, build them, connect them to a backbone network, and license the *spectrum*, or airwave frequency space, for transmission.

We need more bandwidth for mobile devices, too. AT&T now finds that the top three percent of it mobile network users gulp up 40 percent of the network's capacity (thanks, iPhone users!), and network strain will only increase as more people adopt smart phones. These users are streaming Major League Baseball games, exploring the planet with Google Earth, watching YouTube and Netflix, listening to streaming music through Pandora, and more. Get a bunch of iPhone users in a crowded space, like in a college football stadium on game day, and the result is a network-choking data traffic jam. AT&T estimates that it's not uncommon for 80 percent of game-day iPhone users to take out their phones and surf the web for stats, snap and upload photos, and more. But cell towers often can't handle the load.[34] If you've ever lost coverage in a crowd, you've witnessed mobile network congestion first-hand. Trying to have enough capacity to avoid congestion traffic jams will cost some serious coin In the midst of customer complaints, AT&T committed to spending $18 billion on network upgrades to address its wireless capacity problem[35].

Usage	Demand
Voice Calls	4MB/hr.
iPhone browsing	40-60 MB/hr.
Net Radio	60 MB/hr.
YouTube	200-400 MB/hr.

Table X: Avg. Demand Usage by function. Conventional mobile phones use an estimated 100MB/month, iPhones 560MB/month, iPads, upwards of 1GB/month[36]

We're in the midst of transitioning from third generation (*3G*) to fourth generation (4G) wireless networks. 3G systems offer access speeds usually less than 2Mbps (often a lot less)[37]. While variants of 3G wireless might employ an alphabet soup of technologies (EV-DO, UMTS, and HSDPA, among them), 3G standards can be narrowed down to two camps – those based on the dominant worldwide standard called *GSM* (for Global System for Mobile Communications), and the runner-up standards based on *CDMA* (code division multiplex access). Most of Europe and a good chunk of the rest of the world uses GSM. In the U.S., AT&T and T-Mobile use GSM-based 3G. Verizon Wireless and Sprint use the CDMA 3G standard. Typically, handsets designed for one network can't be used on networks supporting the other standard. CDMA has an additional limitation in not being able to use voice and data at the same time.

But 3G is being replaced by high-bandwidth 4G (fourth generation) mobile networks. 4G technologies also fall into two standards camps, LTE and WiMax.

LTE (which stands for *Long Term Evolution*) looks like the global winner, and in the U.S. every major wireless firm except for Sprint is betting on LTE victory. Bandwidth for the service rivals what we'd consider fast cable a few years back. Average speeds range from 5-12Mbps for downloads and 2-5Mbps for upload, although Verizon tests in Boston and Seattle showed download speeds as high as 50Mbps and upload speeds reaching 25Mbps[38].

Competing with LTE is *WiMax*, which stands for Worldwide Interoperability for Microwave Access. Don't confuse it with Wi-Fi. As with other 3G and 4G technologies Wi-Max needs cell towers and operators need to have licensed spectrum from their respective governments (often paying multi-billion dollar fees to do so). Average download/upload speeds should start out at 3-6Mbps and 1Mbps, respectively, although this may go much higher[39].

WiMax looks like a particularly attractive option for cable firms, offering them an opportunity to get into the mobile phone business and offer a 'quadruple play' of services: pay television, broadband Internet, home phone, and mobile. Comcast and Time Warner have both partnered with Clearwire (a firm majority owned by Sprint), to gain access to WiMax based 4G mobile.

4G could also rewrite the landscape for home broadband competition. If speeds increase, it may be possible for PCs, laptops, and set top boxes to connect to the Internet wirelessly via 4G, cutting into DSL, Cable, and Fiber markets.

Satellite Wireless

Wireless systems provided by earth-bound base stations like cell phone towers are referred to as *terrestrial wireless*, but it is possible to provide telecommunications services via satellite. Early services struggled due to a number of problems. The first residential satellite services were download-only, you still needed a modem or some other connection to send any messages from your PC to Internet. Many early systems also required large antennas and were quite expensive. Finally, some services were based on satellites in geosynchronous earth orbit. *GEO* satellites circle the earth in a fixed, or stationary orbit above a given spot on the globe, but in order to do so they're positioned out at a distance roughly the equivalent of the planet's circumference. That means signals travel the equivalent of an around the world trip to get to the satellite, then that same distance again to get to the user. The "last mile" became the last 44,000 miles at best. And if you used a service that also provided satellite upload as well as download, double that to about 88,000 miles. All that distance means higher latency (more delay). [40]

An effort named 03b Networks thinks it might have solved the challenges that plagued early pioneers. O3b has an impressive list of big-name backers that include HSBC bank, cable magnate John Malone, European aerospace firm SES, and Google.

The name 03b stands for the "Other 3 Billion", of the world's population who lack broadband Internet access, and the firm hopes to provide 'fiber-quality' wireless service to more than 150 countries, specifically targeting under-served portions of the developing world. These 'middle earth orbit' satellites will circle closer to the earth to reduce latency (only about 5,000 miles up, less than 1/4th the distance of GEO systems). To maintain the lower orbit, 03b's satellites orbit faster than the planet spins, but with plans to launch as many as 20 satellites, the system will constantly blanket regions served. If one satellite circles to the other side of the globe, another one will circle around to take its place, ensuring there's always an 03b "bird" overhead.

Only about three percent of the sub-Saharan African population are Internet users, vs. about 70 percent in the United States, but data rates in the few places served can cost as much as 100 times comparable systems in the industrialized world[41]. O3b hopes to change that equation and significantly lower access rates. O3b customers will be local telecommunication firms, not end users. The plan is for local firms to buy O3b's wholesale and then resell it to customers alongside rivals who can do the same thing, collectively providing more consumer access, higher quality, and lower prices through competition. O3b is a big, bold, and admittedly risky plan, but if it works, its impact could be tremendous.

Wi-Fi and other Hotspots

Many users access the Internet via Wi-Fi (which stands for *wireless fidelity*). Computer and mobile devices have Wi-Fi antennas built into their chipsets, but to connect to the Internet, a device needs to be within range of a *base station* or *hotspot*. Base station range is usually

around 300 feet (you might get a longer range outdoors and with special equipment; and less range indoors when signals need to pass through solid objects like walls, ceilings, and floors). Wi-Fi base stations used in the home are usually bought by end-users, then connected to a cable, DSL, or fiber provider.

And now a sort of mobile phone hotspot is being used to overcome limitations in those services, as well. Mobile providers can also be susceptible to poor coverage indoors. That's because the spectrum used by most mobile phone firms doesn't travel well through solid objects. Cell coverage is also often limited in the United States because of a lack of towers, and that's a result of the *NIMBY problem* (not in my backyard). No one wants an 80 ft. to 400 ft., unsightly tower clouding their local landscape, even if it will give their neighborhood better cell phone coverage[42]. To overcome reception and availability problems, mobile telecom services firms have begun offering fentocells. These devices are usually smaller than a box of cereal and can sell for $150 or less (some are free with specific service contracts). Plug a fentocell into a high speed Internet connection like an in-home cable or fiber service and you can get "5 bar" coverage in a roughly 5,000 square foot footprint[43]. That can be a great solution for someone who has an in-home, high-speed Internet connection, but wants to get phone and mobile data service in doors, too.

Net Neutrality – What's Fair?

Across the world, battle lines are being drawn regarding the topic of net neutrality. Net neutrality is the principle that all Internet traffic should be treated equally[44]. Sometimes access providers have wanted to offer varying (some say 'discriminatory') coverage, depending on the service used and bandwidth consumed. But where regulation stands is currently in flux. In a pivotal U.S. case, the FCC ordered Comcast to stop throttling (blocking or slowing down) subscriber access to the peer-to-peer file sharing service BitTorrent. BitTorrent users can consume a huge amount of bandwidth – the service is often used to transfer large files, both legitimate (like version of the Linux operating system), and pirated (high definition movies). Then in Spring 2010, a federal appeals court moved against the FCC's position, unanimously ruling that the agency did not have the legal authority to dictate terms to Comcast[45].

On one side of the debate are Internet service firms, with Google being one of the strongest net neutrality supporters. In an advocacy paper, Google states "Just as telephone companies are not permitted to tell consumers who they can call or what they can say, broadband carriers should not be allowed to use their market power to control activity online".[46] Many Internet firms also worry that if network providers move away from flat rate pricing, toward usage-based (or metered) schemes, this may limit innovation. Says Google's Vint Cerf (who is considered one of the "fathers of the Internet" for his work on the original Internet Protocol Suite) "You are less likely to try things out. No one wants a surprise bill at the end of the month"[47]. Metered billing may limit the use of everything from iTunes to Netflix, after all if you have to pay for per-bit bandwidth consumption as well as for the download service, then it's as if you're paying twice.

The counter-argument is that if firms are restricted from charging more for their investment in infrastructure and services, then they'll have little incentive to continue to make the kinds of multibillion dollar investments that innovations like 4G and fiber networks require. Telecom industry executives have railed against Google, Yahoo, Microsoft, and others, calling them free-riders who earn huge profits by piggybacking off of ISP networks, all while funneling no profits back to the firms that provide the infrastructure. Said one Verizon VP, "The network builders are spending a fortune constructing and maintaining the networks that Google intends to ride on with nothing but cheap servers… It is enjoying a free lunch that should, by any rational account, be the lunch of the facilities providers." [48] AT&T's previous CEO has suggested that Google, Yahoo, and other services firms should pay for "preferred access" to the firm's customers. And the CEO of Spain's Telefonica has said the firm is considering charging Google and other Internet service firms for network use[49].

ISPs also lament the relentlessly increasingly bandwidth demands placed on their networks. Back in 2007 YouTube streamed as much data in three months as the world's radio, cable, and broadcast TV channels combined stream in one year[50], and YouTube has only continued to grow since then. Should ISPs be required to support the strain of this kind of bandwidth hog? And what if this one application clogs network use for other traffic, such as e-mail or web surfing? Similarly, shouldn't firms have the right to prioritize some services in order to better serve customers? Some network providers argue that services like video chat and streaming audio should get priority over, say, e-mail which can afford slight delay without major impact. In that case, there's a pretty good argument that providers should be able to discriminate against services. But improving efficiency and throttling usage are two different things.

Internet service firms say they create demand for broadband business, broadband firms say Google and allies are ungrateful parasites that aren't sharing the wealth. The battle lines on the net neutrality frontier continue to be drawn, and the eventual outcome will impact consumers, investors, and will likely influence the continued expansion and innovation of the Internet.

<div align="center">•　•　•　•　•</div>

Hopefully this chaptered helped reveal the mysteries of the Internet. It's interesting to know how 'the cloud' works, but it can also be vital. As we've seen, the executive office in financial services firms considers mastery of the Internet infrastructure to be critically important to their competitive advantage. Media firms find the Internet both threatening and empowering. The advancement of last mile technologies and issues of net neutrality will expose threats and create opportunity. And a manager who knows how the Internet works will be in a better position to make decisions about how to keep the firm and its customers safe and secure, and be better prepared to brainstorm ideas for winning in a world where access is faster and cheaper, and firms, rivals, partners, and customers are more connected.

KEY TAKEAWAYS:

- The slowest part of the Internet is typically the last mile, not the backbone. While several technologies can offer broadband service over the last mile, the U.S. continues to rank below many other nations in terms of access speed, availability, and price.
- Cable firms and phone companies can leverage existing wiring for cable broadband and DSL service, respectively. Cable services are often criticized for shared bandwidth, DSL's primary limitation is that it only works within a short distance of telephone office equipment.
- Fiber to the home can be very fast, but very expensive to build out.
- An explosion of high-bandwidth mobile applications is straining 3G networks. 4G systems may alleviate congestion by increasing capacities to near-cable speeds. Fentocells are another technology that can improve service by providing a personal mobile phone 'hot spot' that can plug into in-home broadband access.
- The two major 3G standards (GSM-based and CDMA-based) will be replaced by two unrelated 4G standards (LTE and WiMax). GSM has been the dominant 3G technology worldwide. LTE looks like it will be the leading 4G technology.
- Satellite systems show promise in providing high-speed access to underserved parts of the world, but few have been successful so far.
- Net neutrality is the principle that all Internet traffic should be treated equally. Google and other firms say it is vital to maintain the openness of the Internet. Telecommunications firms say the should be able to limit access to services that overtax their networks, and some have suggested charging Google and other Internet firms for access to their customers.

QUESTIONS & EXERCISES:
1. Research online for the latest country rankings for broadband service. Where does the U.S. currently rank? Why?
2. Which broadband providers can service your home? Which would you choose? Why?
3. Research the status of Google's experimental fiber network. Report updated findings to your class. Why do you suppose Google would run this 'experiment'? What other Internet access experiments has the firm been involved in?
4. Show your understanding of the economics and competitive forces of the telecom industry – discuss why Verizon chose to go with fiber. Do you think this was a wise decision or not? Why? Feel free to do additional research to back up your argument.
5. Why have other nations enjoyed faster broadband speeds, greater availability, and lower prices?
6. The iPhone has been called both a blessing and a curse for AT&T. Why do you suppose this is so?
7. Investigate the status of mobile wireless offerings (3G and 4G). Which firm would you choose? Why? Which factors are most important in your decision?
8. Name the two dominant 3G standards. What are the differences between the two? Which firms in your nation support each standard?
9. Name the two dominant 4G standards. Which firms in your nation will support the respective standards.
10. Have you ever lost communication access – wirelessly or via wired connection? What caused the loss or outage?

11. What factors shape the profitability of the mobile wireless provider industry? How do these economics compare with the cable and wire line industry? Who are the major players and which would you invest in? Why?
12. Last mile providers often advertise very fast speeds, but users rarely see speeds as high as advertised rates. Search online to find a network speed test and try it from your home, office, mobile device, or dorm. How fast is the network? If you're able to test from home, what bandwidth rates does your ISP advertise? Does this differ from what you experienced? What could account for this discrepancy?
13. How can 4G technology help cable firms? Why might it hurt them?
14. What's the difference between LEO satellite systems and the type of system used by O3b? What are the pros and cons of these efforts? Conduct some additional research – what is the status of O3b and other satellite broadband efforts?
15. What advantages could broadband offer to underserved areas of the world? Is Internet access important for economic development? Why or why not?
16. Does your carrier offer a fentocell? Would you use one? Why or why not?
17. Be prepared to debate the issue of net neutrality in class . Prepare positions both supporting and opposing net neutrality. Which do you support and why?
18. Investigate the status of net neutrality laws in your nation and report your findings to your instructor. Do you agree with the stance currently taken by your government? Why or why not?

• • • • •

About This Work

The goal of this project is to have an impact. At my university, we've bucked the national trend, tripling our Information Systems majors in the three years since we've adopted a business-focused IS teaching approach, and I'm delighted to share this content with you. I hope that Flatworld's free online copies and low-cost print versions encourage wide adoption of this material, and I hope that you and your students enjoy it. Please tell others, and thanks!

Comments & feedback are most welcome! Contact Info:
- E-Mail: john.gallaugher@bc.edu
- Draft Chapters, Cases, Slides, and Podcasts: http://gallaugher.com/chapters
- For updates, supporting articles, and commentary, sign up for the Week in Geek at: http://www.gallaugher.com
- And follow on Twitter at http://twitter.com/gallaugher

About the Author

John Gallaugher is a member of the Dept. of Information Systems in Boston College's Carroll School of Management. Prof. Gallaugher teaches courses and conducts research at the intersection of technology and strategy. An award-winning teacher, he leads the School's TechTrek programs, has co-led the School's European and Asian field study program, and has consulted to and taught executive seminars for several organizations including Accenture, Alcoa, Brattle Group, ING Group, Patni Computer Systems, Partners Healthcare,

Staples, State Street, and the U.S. Information Agency. Prof. Gallaugher's research has been published in leading journals, including Harvard Business Review and MIS Quarterly. His comments on business and technology have appeared in many outlets including The New York Times, National Public Radio, WCVB-TV, The Seattle Times, the Associated Press, eWeek, The Daily Yomiuri (Japan) and The Nation (Thailand).

This reading is available to faculty for non-commercial use. Enjoy! Edited versions of this chapter, plus additional chapters and cases, are available in Professor Gallaugher's book "Information Systems: A Manager's Guide to Harnessing Technology", published both free online and low-cost (less than $30) print version, by Flat World Knowledge (http://FlatWorldKnowledge.com).

DEFINITIONS:

URL (Uniform Resource Locator) – often used interchangeably with "web address", URLs identify resources on the Internet along with the application protocol need to retrieve it.

Protocol – enable communication by defining the format of data and rules for exchange.

http (hypertext transfer protocol) – application transfer protocol that allows web browsers and web servers to communicate with each other.

ICANN (Internet Corporation for Assigned Names and Numbers) – non-profit organization responsible for managing the Internet's domain and numbering systems.

Web Hosting Service – a firm that provides hardware and services to run the websites of others.

Cybersquatting – acquiring a domain name that refers to a firm, individual, product, or trademark, with the goal of exploiting it for financial gain. The practice is illegal in many nations, and ICANN has dispute resolution mechanisms that in some circumstances can strip cybersquatters of registered domains.

Load Balancing – distributing a computing or networking workload across multiple systems in order to avoid congestion and slow performance.

Fault Tolerant - systems that are capable of continuing operation even if a component fails.

HTML (hypertext markup language) – language used to compose web pages.

IP Address – unique identifier assigned to a network-connected device. All devices connected to the Internet are assigned a unique IP address.

NAT (network address translation) – a technique often used to conserve IP addresses by maps devices on a private network to single internet-connected device that acts on their behalf.

DNS (domain name system) – hierarchical system of nameservers that acts as the Internet's "phone book", mapping host/domain name combinations to IP addresses.

Cache - a temporary storage space used to speed computing tasks.

TCP (transmission control protocol) – works at both ends of most Internet communication to ensure a perfect copy of a message is sent.

IP (Internet protocol) – routing protocol that is in charge of forwarding packets on the Internet.

Packet (sometimes called a datagram) – a unit of data forwarded by a network. All Internet transmissions – URLs, web pages, e-mails, are divided into one or more packets.

Router – a computing device that connects networks and exchanges data between them.

UDP (user datagram protocol) – protocol that operates instead of TCP in applications where delivery speed is important and quality can be sacrificed.

Broadband – broadly refers to high speed Internet connections and is often applied to 'last-mile' technologies.

Bandwidth – network transmission speeds, typically expressed in some form of bits per second (bps).

Amdahl's Law – a system's speed is determined by its slowest component

Internet Backbone – high speed data lines provided by many firms all across the world which interconnect and collectively form the core of the Internet.

Last Mile Technologies – technologies that connect end users to the Internet. The last mile problem refers to the fact that these connections are usually the slowest part of the network.

Coaxial cable (or coax) – insulated copper cable commonly used by cable television providers.

DSL (digital subscriber line) – broadband technology that uses the wires of a local telephone network.

FTTH (fiber to the home) – broadband service provided via light-transmitting fiber optic cables.

Wireless spectrum – frequencies used for communication. Most mobile cell phone service have to license spectrum. Some technologies (such as Wi-Fi) use unlicensed, public spectrum.

REFERENCES

1 Maney, K., "Tuvalu's sinking, but its domain is on solid ground", *USA Today*, April 27, 2004.

2 Bosker, B., "The 11 Most Expensive Domain Names Ever", *The Huffington Post*, March 10, 2010.

3 Streitfeld, D., "Web Site Feuding Enters Constitutional Domain", *The Washington Post*, Sept. 11, 2000.

4 McCullagh, D., "Ethical Treatment of PETA Domain", Wired, Aug. 25, 2001.

5 Konrad, R., and Hansen, E., "Madonna.com embroiled in domain ownership spat", CNet, Aug. 21, 2000.

6 Morson, D., "Apple VP Ive loses domain name bid", *MacWorld*, May 12, 2009.

7 Kotadia, M., "MikeRoweSoft settles for an Xbox", *CNet*, Jan. 26, 2004.

8 Ward, M., "Internet Approaches Addressing Limit", *BBC News*, May 11, 2010.

9 Shankland, S., "Google Tries to Break IPv6 Logjam by Own Example", *CNet*, March 27, 2009.

10 Shankland, S., "Google Tries to Break IPv6 Logjam by Own Example", *CNet*, March 27, 2009.

11 Arnoldy, B., "IP Address Shortage To Limit Intenet Access", *USA Today*, Aug. 3, 2007.

12 Godin, D., "Cache-poisoning attack snares top Brazilian bank", *The Register*, April 22, 2009.

13 Davis, J., "Secret Geek A-Team Hacks Back, Defends Worldwide Web", *Wired*, Nov. 24, 2008.

14 Hutchinson, J., "ICANN, Verisign place last puzzle pieces in DNSSEC saga", *NetworkWorld*, May 2, 2010.

15 Timmons, H., "A London Hedge Fund That Opts for Engineers, Not M.B.A.'s", *The New York Times*, Aug. 18, 2006.

16 Iati, R., "The Real Story of Trading Software Espionage", *Advanced Trading*, July 10, 2009.

17 Berenson, A., "Arrest Over Software Illuminates Wall St. Secret", *The New York Times*, Aug. 23, 2009.

18 Schmerken, I., "High-Frequency Trading Shops Play the Colocation Game", *Advanced Trading*, Oct. 5, 2009.

19 Daimler, E., and Davis, G., "'Flash crash' proves diversity needed in market mechanisms", *The Pittsburgh Post-Gazette*, May 29, 2010.

20 Spangler, T., "Cisco Clarifies 100-Gig AT&T Backbone Claim", *Multichannel News*, March 9, 2010; and Zacks.com, "AT&T Tests 100 Gb Ethernet in Move Toward Faster Internet", *SeekingAlpha*, March 10, 2010.

21 Gilder, G., *Telecosm: How Infinite Bandwidth Will Revolutionize Our World*, Free Press, New York, NY, 2000.

22 Lawson, S., "US ranks 26th in new broadband index", ComputerWorld, May 25, 2010.

23 Hansell, S., "The Broadband Gap: Why Is Theirs Faster?", *The New York Times*, March 10, 2009.

24 Farzad, R., "The Truth About Bandwidth", *BusinessWeek*, Feb. 3, 2010

25 Pandora FAQ at Pandora.com.

26 LG Knoweldge Base, "Bandwidth needed for instant streaming", LGknowledgebase.com.

27 Thompson, R., "DSL Internet vs Cable Internet", *High Speed Internet Access Guide*, March 23, 2010.

[28] Hansell, S., "The Broadband Gap: Why Is Theirs Faster?", *The New York Times*, March 10, 2009.

[29] Hansell, S., "The Broadband Gap: Why Is Theirs Faster?", *The New York Times*, March 10, 2009.

[30] Hansell, S., "The Broadband Gap: Why Is Theirs Faster?", *The New York Times*, March 10, 2009.

[31] King, R., "Telecom Companies Scramble for Funding", *BusinessWeek*, Aug. 3, 2009.

[32] Higginbotham, S., "Verizon Tests 10 Gbps to the Home. Yeah, You'll Have to Share", *GigaOm*, Dec. 17, 2009.

[33] Ingersoll, M., and Kelly, J., "Think Big with a Gig: Our Experimental Fiber Network", The Google Blog, Feb. 2, 2010; and Rao, L, "The Final Tally: More Than 1100 Cities Apply For Google's Fiber Network", *TechCrunch*, March 27, 2010.

[34] Farzad, R., "AT&T's iPhone Mess", *BusinessWeek*, Feb. 3., 2010.

[35] Edwards, C., and Kharif, O., "Sprint's Bold Play on a 4G Network", *BusinessWeek*, March 30, 2010.

[36] Source: Farzad, R., "The Truth About Bandwidth", *BusinessWeek*, Feb. 3, 2010.

[37] German, K., "On Call: Welcome to 4G", *CNet*, March 9, 2010.

[38] German, K., "On Call: Welcome to 4G", *CNet*, March 9, 2010.

[39] Lee, N., "Sprint's 4G plans explained", *CNet*, May 19, 2010.

[40] Ou, G., "Why Satellite Service is So Slow", ZDNet, Feb. 23, 2008.

[41] Lamb, G., "O3b Networks: A far-out plan to deliver the Web", *Christian Science Monitor*, Sept. 24, 2008.

[42] Dechter, G., and Kharif, O., "How Craig McCaw Built a 4G Network on the Cheap", *BusinessWeek*, Mau 24. 2010.

[43] Mims, C., "A Personal Cell Phone Tower", *Technology Review*, April 7, 2010.

[44] Honan, M., "Inside Net Neutrality", *MacWorld*, Feb. 12, 2008.

[45] <no author> "What is Net Neutrality", *The Week*, April 7, 2010.

[46] *A Guide to Net Neutrality for Google Users*, Google, 2008. Available at: http://www.docstoc.com/docs/1064274/A-Guide-to-Net-Neutrality-for-Google-Users

[47] Jesdanun, M., "As the Internet Turns 40, Barriers Threaten Growth", *Technology Review*, Aug. 31, 2009.

[48] Mohammed, A., "Verizon Executive Calls for End to Google's 'Free Lunch'", *The Washington Post*, Feb. 7, 2006.

[49] Lunden, I., "Broadband Content Bits: Web Drama Investment, PPL Video Store, Telefonica To Charge?", *PaidContent:UK*, Feb. 11, 2010.

[50] Swanson, B., "The Coming Exoflood", *The Wall Street Journal*, Jan. 20, 2007.