Net Neutrality Regulatory Proposals: Operational and Engineering Implications for Wireless Networks and the Consumers They Serve

January 14, 2010

http://www.rysavy.com

Table of Contents

Executive Summary

As the Federal Communications Commission considers applying "net neutrality" regulations to the wireless Internet, it has sought broad public comment on the implications of such rules for consumers and the networks which support them. The imposition of any such new rules will have a variety of impacts across the Internet ecosystem and particularly on wireless network operations and capabilities. Thus, it is critically important to the future of the mobile Internet that the FCC conduct a detailed and clear-eyed assessment of how the proposed new regulations will impact the engineering, operational, and technological requirements of current and next generation wireless networks.

It is the purpose of this paper to provide an engineering analysis of the impacts the proposed "net neutrality" provisions will have on the nation's wireless networks and consumers. The paper also assesses whether and how wireless networks may need to be adapted operationally to support any such new regulation, and the technological, operational, and engineering challenges in doing so.

Central to this assessment is the simple and unavoidable fact that wireless networks, compared to traditional wired broadband networks, are extremely constrained in capacity. Current broadband adoption trends, combined with trends in mobile computing, point to potential exhaustion of this capacity in the relatively near future. It is crucial that these networks implement functions that enable the efficient management of existing capacity and allow continued growth in innovative services, all while not foreclosing new business opportunities and revenue streams that are critical to ensuring network operators will have the money to invest in continued network expansion and upgrades.

Wireless innovation occurs at multiple levels. For example, digital signaling processing is the bedrock for radio innovations such as high-order modulation, new multiplexing methods such as orthogonal frequency division multiple access used in LTE and WiMAX technologies, and advanced antenna methods. These innovations are dynamic and ongoing, enabling network operators to extract more capacity and efficiency from spectrum and network infrastructure. The dynamic, integrated wireless ecosystem of innovation consists of innovation in the networks, the devices, applications, and Internet services all of which have collectively entered a virtuous cycle in which innovation and market growth continually propels new advances. This cycle, however, can only be maintained with careful management of available resources.

Today's most powerful wide-area wireless networks have significantly less capacity than wireline networks. This disparity is due to the constraints of physics: a single fiber-optic cable has more data capacity than the entire radio spectrum to 100 GHz. Providing access to more spectrum will certainly improve the situation, but spectrum by itself will not be a complete answer to the wireless network capacity dilemma. New technologies, such as 4G, will also help, but capacity will still be constrained. While today's wireless technologies are reaching the theoretical limits of the spectral efficiency that can be achieved, we have scarcely scratched the surface of growing consumer demand for data-intensive applications, which means the capacity issue is going to become exponentially more serious.

Mobile broadband traffic growth has been exponential and is expected to more than double every year for the next five years. Cloud computing, video-based entertainment and communications, and social networking are all applications that can consume significant amounts of data. In fact, just a handful of users running popular applications can easily consume the entire capacity of a local coverage area in today's 3G networks. And, these applications are becoming more bandwidth intensive. Thus, any added capacity will be readily exhausted by available applications, especially because many were developed for wireline networks where constraints are much less severe.

With exponential growth in mobile broadband usage, the only way to continue the network growth needed to support America's insatiable demand is to make more spectrum available and to implement network and traffic management mechanisms that carefully allocate spectrum and capacity to ensure good quality of service to all users. These dynamic, engineering-led efforts enable operators to provide a consumer experience that is satisfactory to the largest possible number of simultaneous users. They also enable ongoing innovation and spur the development of new classes of applications. The resulting application innovation and smooth network operation will allow the virtuous cycle of mobile innovation to continue.

Particularly in the wireless context, the regulatory equivalent of an 'all bits are created equal' edict would be tremendously disruptive because, quite simply, not all bits are equal. For example, in multimedia communications, the voice component is much more important for communications than the video component. Somebody's e-mail text is more important than a few background bits of a video stream. Network and traffic management, so long as it is done transparently and not anti-competitively, enables the best user experience for the greatest number of users. This is especially important in addressing the wireless network challenge in which a small number of users can adversely impact a much larger number of users. The proposed non-discrimination and network management proposals would block network management efforts that are essential to maintaining and enhancing the tools available to mitigate this risk for consumers.

Wireless network engineers need flexibility to manage finite capacity to provide the best service possible to the greatest number of consumers. The short term and long term approaches to maximizing capacity on wireless networks will continue to evolve and improve, but they would by undermined or eliminated at this defining juncture by some of the proposals contained in the FCC's current Open Internet NPRM.

Introduction

The rise of the wireless and mobile computing industries represents one of humanity's greatest achievement in terms of technology, social and cultural impacts, and economics. This rise has created and expanded economic opportunity for millions of Americans, and hundreds of millions of people around the world, has improved this country's productivity, and has had a profoundly beneficial effect on the health, education, and safety of society. The leading edge of this industry is called mobile broadband – the ability to deliver high-speed Internet connectivity to people regardless of location across this country and the globe using radio waves.

Due to fundamental constraints of physics, wireless service providers offer mobile broadband services across networks that are capacity constrained and that will likely continue to be constrained regardless of how much additional spectrum is made available to increase capacity. As explained herein, it is because of these fundamental technical constraints that operators of wireless networks require a heightened degree of technical flexibility and autonomy as they seek to support hundreds of thousands of bandwidth-intensive mobile applications and many millions of Internet-centric mobile devices, including non-telecom devices that use wireless networks as their preferred access point to the Internet.

This paper, sponsored by Mobile Future, covers the key innovations in technology that are driving improvements in wireless network capacity, provides details on wireless network performance constraints, identifies trends in wireless data consumption and explains how these consumption demands threaten to overwhelm capacity. This paper also discusses the options currently available to wireless network operators to avoid network overload and the means by which the industry can stay abreast of capacity constraints, including managing what and how devices attach and interconnect with wireless networks. This paper provides the technical basis for understanding the unintended impact that certain of the current "net neutrality" regulatory proposals will have on the types of wireless services consumers want and that operators wish to offer.

Technology Innovation and Evolution

The wireless industry has fostered an environment that supports tremendous innovation. Radio technologies have become much more powerful, user devices have become more capable, applications more useful, and services more compelling. Today's mobile broadband is the intersection of wireless technology and the Internet, a combination of two white-hot technology areas that has produced a healthy ecosystem that encourages new technology investment at multiple levels including: signal processing, radio communications, Internet protocols, memory, handheld computing platforms, new emerging mobile platforms, applications, and Internet Web technologies.

Table 1 details these areas of innovation. While there are many specific technologies that could be listed, this table lists the key items that have revolutionized the industry.

Table 1: Wireless Industry Areas of Innovation

Area of Innovation	Technology	Relevance
Processing	Digital signal processors (DSPs)	Virtually all radio communications processing nowadays is done digitally, facilitating sophisticated radio innovations, as listed below.
	Central processing units (CPUs)	Powerful computers on devices enable rich user interfaces, multimedia rendition, and camera/video capabilities.
Radio	Higher-order modulation	With higher-order modulation, more data can be transmitted across a radio channel.
	New multiplexing methods such as Orthogonal Frequency Division Multiple Access (OFDMA)	OFDMA permits extremely high throughput rates that are difficult to achieve using alternate radio methods.
	Advanced Antenna Techniques such as Multiple Input Multiple Output (MIMO)	MIMO increases range, capacity and throughput rates.
Internet Protocols	All services handled in the Internet Protocol (IP) domain	Networks fully based on IP lower infrastructure costs and enable new types of services such as integrated voice/Web communications.
Memory	Solid state devices with large capacity	Tiny memory chips allow for storage of thousands of songs and photos, and hours of video.
Handheld Platforms	New operating systems for smartphones including new user interfaces	Highly competitive platform market now spans Android, BlackBerry, iPhone, mobile Linux, Palm Pre, Symbian, and Windows Mobile – providing users tremendous variety and usability.
Emerging Mobile Platforms	Between smartphones and laptops, there are now netbooks, smartbooks and Mobile Internet Devices.	Different usage models will favor different devices. A broad selection of device types expands the market.
Applications	Smartphones and other new platforms can run a wide range of applications.	More than 100,000 applications for the iPhone, along with tens of thousands for other platforms, demonstrate the momentum of app development.
Web Technologies	Powerful Web methods are being adapted for mobile devices.	Users can access more applications and content via their mobile Web browsers, further increasing the power of mobile platforms.

The outcome of this innovation in radio and computing has been exploding consumer consumption of powerful mobile-broadband technologies such as High Speed Packet Access (HSPA), Code Division Multiple Access 2000 Evolution Data Optimized (CDMA2000 EV-DO), Wireless Inter-Operability for Microwave Access (WiMAX), and Long Term Evolution (LTE). Accompanying these new network technologies are new families of Internet-centric mobile devices that avail themselves of the network capabilities. That more than fifty million Americans have adopted this technology is testament to how affordable the innovation has become for the average person.

The consumer-centric mobile broadband industry does not sit idly by. In the United States, even as the first digital cellular technology was being standardized based on an approach called Time Division Multiple Access (TDMA) in the 1990s, innovators proposed an alternate approach called Code Division Multiple Access (CDMA). The result today is global availability of cellular service based on TDMA in a technology called Global System for Mobile Communications (GSM), with simultaneous availability in much of the world of mobile broadband based on CDMA, namely HSPA and EV-DO.

Just as CDMA challenged the TDMA approach, now a new radio method is challenging all prior approaches: orthogonal frequency division multiple access (OFDMA). OFDMA is not only the basis of WiMAX as well as LTE, but the basis of what is being referred to as 4G. There are no true 4G technologies deployed today, but these are expected mid decade through subsequent versions of LTE and WiMAX.[1] This rapid pace of wireless innovation, with an entirely new wireless approach deployed globally every ten years at an investment level of many tens of billions of dollars, has only been possible due to fierce competition.

This unusually dynamic sector of the global Information and Communication Technologies (ICT) industry area now is the vanguard of the computing industry, especially as all phones are becoming "smart." The companies leveraging these mobile networking capabilities include both consumer and business application developers. Every major company in the computing industry today is now heavily engaged: Apple, Google, IBM, Microsoft, Oracle, SAP, Yahoo, and so forth.

We have entered a virtuous cycle, as shown in Figure 1, in which innovation and market growth feed on themselves. A crucial question for policymakers is how to foster this circle and expand its successes.

[1] Note: 1G was analog cellular, first available in the 1980s. 2G includes digital systems such as GSM and CDMA2000, made widely available in the 1990s. 3G refers to systems such as HSPA and EV-DO, broadly deployed this decade. True 4G systems will be those complying with recently published requirements from the International Telecommunications Union (ITU) in a project called International Mobile Telephone (IMT) Advanced. Technologies such as LTE and WiMAX at this time are not true 4G but will evolve to meet 4G requirements.

Figure 1: The Virtuous Cycle of Innovation

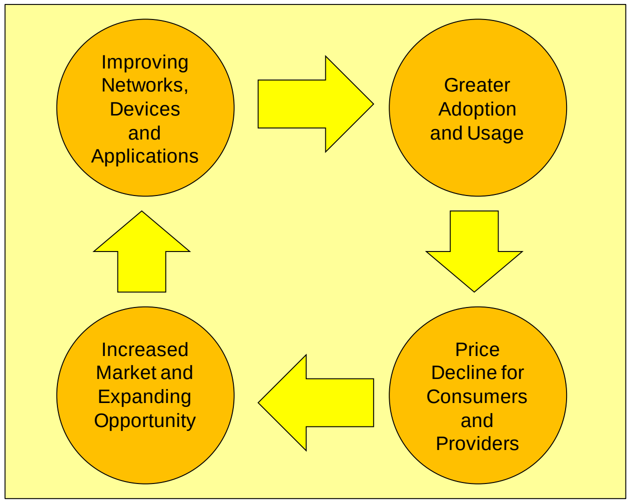

Wireless Network Capacity and Performance Constraints

Unfortunately, wireless capacity, the fuel that drives the mobile-broadband networks, is finite, especially when compared to other existing broadband technologies. Central to the understanding of why wireless capacity is finite is the concept of "spectral efficiency," which is the amount of data bandwidth obtainable in a certain amount of spectrum and which is measured in bits per second (bps) per Hertz (Hz), or bps/Hz. By looking at the amount of bandwidth available to operators, the technology used, and the spectral efficiency of that technology, one can readily calculate the amount of data capacity in a network in each cell sector.[2] The following table shows some examples of data capacity in different kinds of deployment scenarios.

Table 2: Wireless Capacity of Different Wireless Technology Deployments

Technology Used	Spectral Efficiency	Bandwidth Assumption	Downlink Sector Capacity
Evolution Data Optimized (EV-DO, as used by Sprint and Verizon)	0.75 bps/Hz	3 radio channels of 1.25 MHz (3 X 1.25 MHz forward + 3 X 1.25 MHz reverse)	2.81 million bits per second (Mbps)
High Speed Packet Access (HSPA, as used by AT&T and T-Mobile)	0.75 bps/Hz	1 radio channel of 5 MHz (5 MHz forward + 5 MHz reverse)	3.75 Mbps
HSPA	0.75 bps/Hz	2 radio channels of 5 MHz (10 MHz forward + 10 MHz reverse)	7.5 Mbps

[2] Most cells are divided into three sectors by using directional antennas on the cell tower. Each sector provides a relatively independent coverage area.

Technology Used	Spectral Efficiency	Bandwidth Assumption	Downlink Sector Capacity
WiMAX	1.0 bps/Hz	1 radio channel of 10 MHz (time division duplex with 2/3 of time for downlink)	6.7 Mbps
LTE	1.5 bps/Hz	1 radio channel of 10 MHz (10 MHz forward + 10 MHz reverse)	15 Mbps

The downlink sector capacity is not the amount of bandwidth that each subscriber has available to him or her. Rather, it is the shared bandwidth across all of the active subscribers in that cell sector. With about 240,000 cell sites in the U.S.[3] and approximately 280 million subscribers[4], this equates to over 1,000 subscribers per cell site. In some dense population areas such as New York or Chicago, the number of subscribers per cell site can be much higher. For the sake of argument, if we use 1,000, and divide by 3, we obtain 333 subscribers per cell sector. We can then assume different levels of mobile broadband penetration. Even if we take a modest penetration level of 10%, this equates to 33 users.[5]

Now, consider the most advanced cellular technology available in the newest deployment: forthcoming deployments of LTE. One major operator is deploying LTE using 10+10 MHz, which will give it a downlink capacity of 15 Mbps. If 33 mobile broadband subscribers in the same geographic area were to use LTE, they would have to share this capacity.

Contrast this, as shown in Figure 2, with a high-speed wireline Internet connection using services available today. This might be a 50 Mbps cable modem or fiber-to-the-home (FTTH) subscriber.[6] This entire capacity is available to that single subscriber. The difference is profound. Not only does the wireline channel deliver much higher throughput on average, but the potential to increase capacity is far greater. In contrast, radio spectrum is much more finite. It is also more expensive to add more capacity. Wireless networking company Camiant estimates that the cost of adding incremental wireless capacity is over seven times the cost of adding incremental wireline capacity.[7]

[3] Source: Wireless Week, "By The Numbers: July - August 2009," August 1, 2009.

[4] Source: CTIA, Wireless Quick Facts, Mid-Year Figures.

[5] Note: this is a highly simplified analysis. Detailed analysis would need to consider the distribution of subscribers, actual usage, specific geographical characteristics impacting radio communications, and so forth.

[6] Note: Even higher rates such as 100 Mbps are available in some areas. 50 Mbps is still a premium service, but will likely become common before long.

[7] Analysis provided as contribution to this paper.

Figure 2: Graphical Depiction of Wireline versus Wireless Capacity (Representative Scenario)

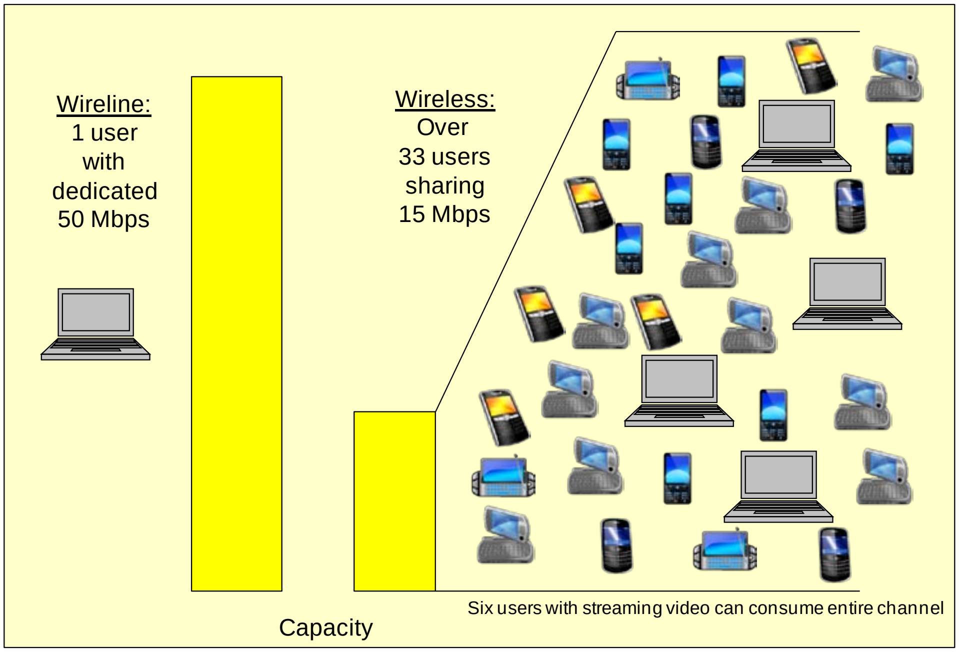

In the scenario just described, six wireless users watching video could consume the entire capacity of the LTE channel.[8]

The fundamental source of the discrepancy between wired and wireless capacity is that fiber-optic capacity is so much greater than RF capacity. In fact, one strand of fiber-optic cable has greater capacity than the entire RF spectrum. The state-of-the-art for fiber capacity (demonstrated capability) is over 10 Tera bits per second (10 Tbps, which is 10,000 billion bits per second). In comparison, if you had all the spectrum to 100 GHz (rather unlikely) and had huge spectral efficiency of 10 bps/Hz, you would still only have 1000 Gbps (1 Tbps), less than 10% of the achievable capacity of the fiber-optic cable, as depicted in Figure 3. And with the fiber-optic cable, you can add another one once you have consumed the first one.

[8] For example, Netflix offers free streaming video service to all of its subscribers. In high-definition mode, streaming rates are up to 3.8 Mbps.

Figure 3: Comparison of Wireless and Wireline Capacity

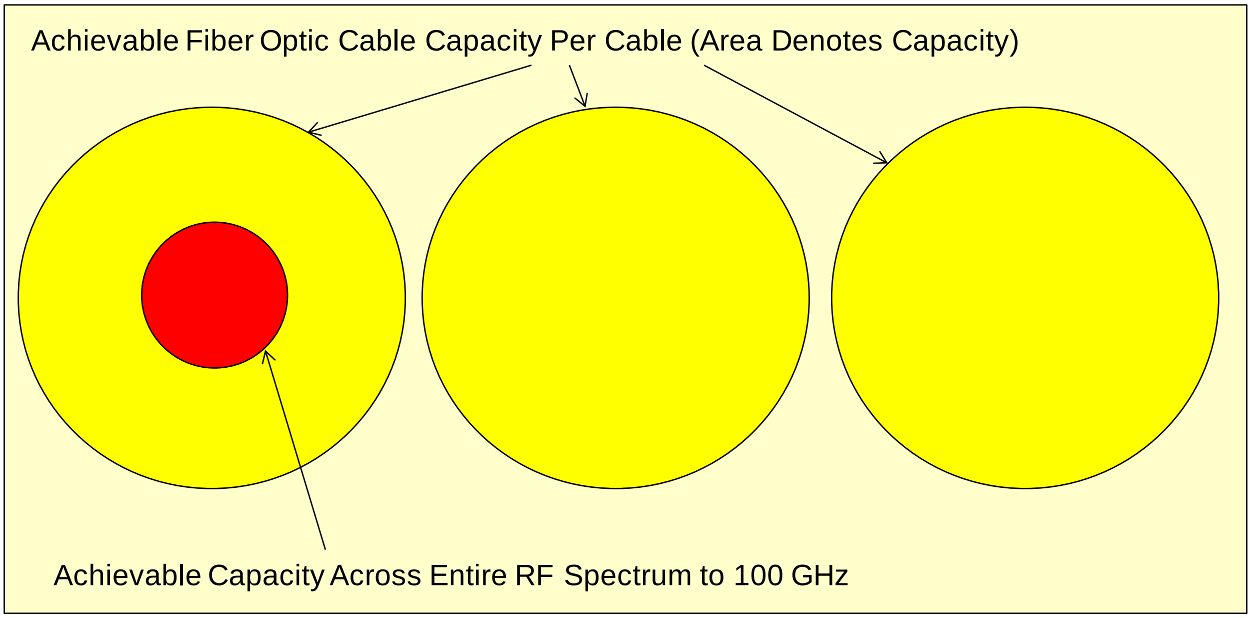

It is a looming shortage of capacity that is driving the wireless industry to seek access to more spectrum for commercial use. Some are calling for 800 MHz to 1 GHz of additional spectrum to be made available to commercial operators.

Introducing 800 MHz to 1 GHz of additional spectrum into the U.S. mobile marketplace would result in almost a tripling of spectrum capable of supporting mobile communications; however, it is highly unlikely that this much new spectrum can be made available any time soon. Chetan Sharma, a wireless analyst, states that it takes seven to ten years to procure spectrum for wireless use.[9] Therefore, intermediate steps need to be taken to maximize capacity, while continuing to develop new, even more efficient technology solutions to extract more capacity from both existing and future spectrum allocations.

New technologies like LTE are more efficient than prior technologies such as HSPA, but not that much more so. Technology alone can only help so much since wireless technologies are approaching theoretical physics limits for spectral efficiency in what is known as the Shannon Bound. This limit dictates the maximum achievable spectral efficiency relative to noise. Today's most advanced wireless technologies such as WiMAX and LTE have effectively reached this limit.[10]

---

[9] Source: RCR Wireless, "Analyst Angle: Solutions for the Broadband World, November 4, 2009," http://www.rcrwireless.com/apps/pbcs.dll/article?AID=/20091104/OPINION/910309995/analyst-angle-solutions-for-the-broadband-world

[10] For a detailed explanation, refer to "HSPA to LTE-Advanced. 3GPP Broadband Evolution to IMT-Advanced (4G)," September 2009. See Figure 20 and associated discussion.

Data Consumption

There are a number of ways of quantifying mobile broadband data consumption. A good starting point is to realize that data usage is growing 100 times faster than voice, as evidenced in one major US operator's network.[11] Cisco projects mobile broadband traffic growing at over a 100% compound annual growth, as shown in Figure 4. Consistent with Cisco's study, Coda Research Consultancy anticipates a 40-fold increase in traffic in 2017 over 2009.[12]

Figure 4: Cisco Global Mobile Broadband Traffic Projection[13]

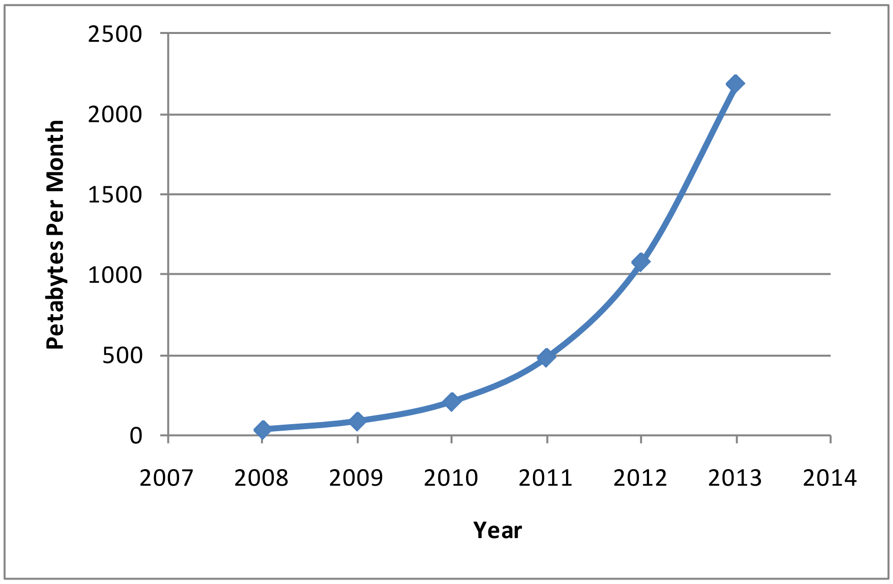

There are two aspects to this growth. One is that users across the board are increasing their wireless data consumption. The other is that usage across users is extremely uneven – a small percentage of users are consuming a disproportionate amount of traffic. On AT&T's network, just three percent of smartphone users consume 40% of network capacity.[14] Similarly, Cisco recently documented that in Q3

---

http://www.rysavy.com/Articles/2009_09_3G_Americas_RysavyResearch_HSPA-LTE_Advanced.pdf and accompanying text.

[11] Source: Rysavy Research, "EDGE, HSPA and LTE – Broadband Innovation," September 2008, Figure 4, http://www.rysavy.com/Articles/2008_09_Broadband_Innovation.pdf

[12] Source: Coda Research Consultancy, "Mobile Broadband and Portable Computers: Revenue, User and Traffic Forecasts 2009-2017," July 19, 2009, http://www.fiercebroadbandwireless.com/story/report-laptops-netbooks-drive-exponential-mobile-broadband-growth/2009-07-19

[13] Source: Cisco, "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update," January 29, 2009. One petabytes is one million gigabytes or 10^{15} bytes.

[14] Source: Associated Press, "AT&T: Tighter control of cell data usage ahead," December 9, 2009.

of 2009, the top 1% of Internet connections consumed over 20% of traffic and the top 10% consumed 60% of the traffic.[15]

Related to uneven usage is that there are multiple types of traffic on data networks. One is bursts of traffic of relatively short duration. Examples include sending or receiving an e-mail or downloading a Web page. Another is streaming, where data is sent continuously, such as listening to an Internet radio station or viewing a video. Streaming consumes far more data than bursty traffic, because it is continuous. While a Web page that a user reads might have some images on it, a video consists of 30 images per second, minute after minute. That is why viewing a video stream can easily consume 100 times as much data as reading more static Web content. An emerging type of usage pattern is from netbooks and smartbooks that rely heavily on cloud-computing methods wherein all of the applications are Internet based. Cloud computing demands constant and heavy network access. The type of applications that are running can thus have widely varying impacts on capacity, and network management schemes and service plans must be flexible enough to accommodate a wide range of usage patterns.

Consumption trends show that available capacity could be exhausted before long, driven heavily by streaming and cloud-computing applications. With data capacity limited on current networks, with data usage growing exponentially, with spectrum relief years away, with networks already groaning under the weight of subscriber demands, there is only one way for the industry to move forward: dynamic, real-time allocation of existing capacity. Efforts to do just this will be severely hampered under the currently proposed net neutrality regulations as applied to wireless networks.

Quality of Service

Gaining access to more spectrum will be crucial in addressing demand, but by itself more spectrum will remain insufficient, unless it were possible to deploy a huge number of cell sites with sufficiently small coverage areas. But such a scenario is neither practical nor realistic. Cost, tower siting, zoning, and other regulatory and operational constraints make such a goal unrealistic.

In the absence of infinitely available spectrum, implementation of quality of service (QoS) architectures and related mechanisms will therefore be central to the management of the limited capacity of these networks. Without such management, a user's experience will be severely degraded, which will depress the demand for wireless broadband, thereby reversing the virtuous cycle of innovation and investment depicted above in Figure 1.

New technologies such as the Evolved Packet Core (EPC), which is the core network technology for LTE, hands operators an impressive array of new capabilities for managing QoS at the level that most American consumers have grown to expect in the voice realm. QoS is realized through traffic and queue management, a process for assigning packets from different users and applications to available network resources. QoS provides control over items such as guaranteed bit rates, the amount of delay, variations

[15] Source: Cisco, "Cisco Visual Networking Index: Usage Study," October 21, 2009.

in delay, and the amount of packet loss. QoS control is a basic requirement for certain types of applications that need specific quality levels such as voice and video. In wireline networks, QoS is often accomplished by simply over-provisioning bandwidth. In wireless networks, however, simply adding capacity to facilitate QoS is impossible. Instead, the network must control QoS through traffic management.

Technologies such as LTE have dynamic policy management as part of what is called the Policy and Charging Control (PCC) architecture. This architecture, which allows for dynamic QoS and other forms of traffic management, will allow wireless operators to easily customize and modify service plans to cater to fairly granular, yet varied, consumer preferences.

Key to mass market acceptance of wireless broadband and the realization of the technical capabilities of a 4G world will be the dynamic evolution of services, applications, devices, and service plans. At the core of this world, in which the customer can obtain a customized service plan, is the ability of wireless network operators to match demand to capacity in real time as shown in Figure 5.

Figure 5: Traffic Discrimination and Prioritization Enables New Services and Capacity Management

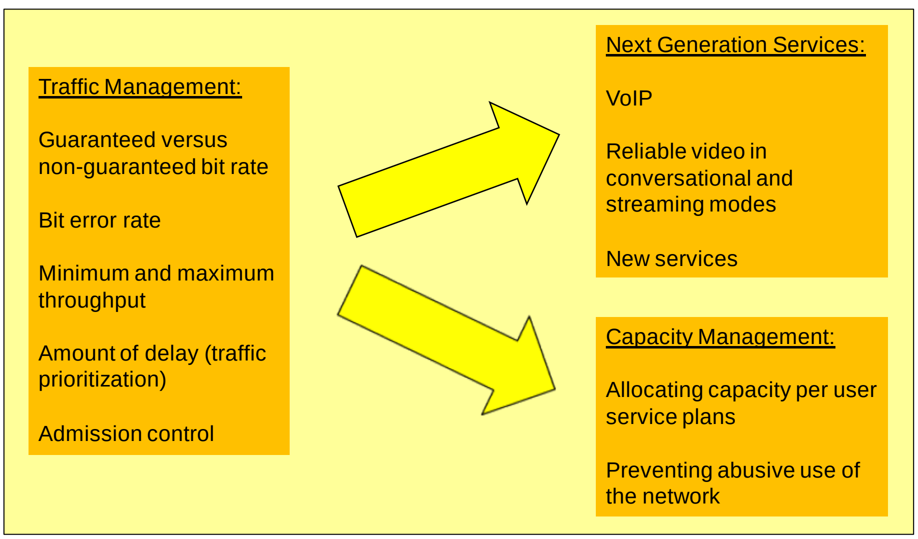

Beyond network management, these evolving QoS capabilities in networks will enable application developers to create new types of applications, since they will be able to leverage mechanisms such as amount of delay or minimum bit rates. A game developer, for example, would be assured that an interactive game would perform at satisfactory levels. Network neutrality, however, could block these mechanisms, since it effectively forces all applications and networks to operate on a “best efforts” basis.

Another capability that would benefit many applications is to be able to embed voice communications. A customer-support Web site, for example, might wish to integrate a voice channel. The voice path, however, will require different QoS parameters than other elements of the Web page. Denying this ability by requiring all bits to be treated equally undermines what could be new forms of interactions over mobile networks. Social networking applications, also a huge Internet growth area, would similarly benefit from QoS capabilities.

By denying application developers and network operators access to highly differentiated and creative services, the virtuous cycle described earlier will be disrupted because the range of possible applications and solutions will be diminished.

Industry Innovation in Response to Finite Capacity

As mobile-broadband constraints are better understood, application developers will architect their applications accordingly. Some applications already manage bandwidth carefully. Many, however, were developed or are being developed for fixed-line networks where capacity is much less of an issue. An ecosystem in which bandwidth is a well-understood and managed item will result in much more efficient application architectures. In contrast, an environment where net neutrality permits any application of the user's choice, *regardless of how inefficient or wasteful of bandwidth*, presents no incentive for application developers to consider a more efficient approach.

The following is a list of ways that software applications can limit how much data they communicate:

- Compress data before sending it.
- Only send data that has changed, keeping local caches of prior data sent.
- Monitor the network type being used, and defer some operations for other connection types. (For example, defer operating-system updates to a Wi-Fi connection.)
- Provide means of partially viewing documents on an as-needed basis, rather than downloading entire large files.
- Resume operations from the point of failure rather than restarting them.

Service Plans and Management

In fixed-line networks, service plans have been relatively straight forward, with most ISPs nowadays offering different tiers (throughput rates) of service. ISPs can manage this approach fairly easily because they enjoy predictability – about where their subscribers are and their own network capabilities.

In contrast, wireless operators have to provide their customers with high quality of service in a much more dynamic scenario, including having to address in real-time: changes in the radio channel because signal quality can vary so widely across customers; unpredictable loading in each cell since users are mobile; and much lower overall capacity.

The current service and pricing plans for wireless broadband stem from service plans built around mobile voice in which every voice call presented the same network load factor as every other voice call, and where the necessary quality attributes to sustain communication are consistent across users and over time. These service plans are simply inapplicable to the wireless broadband world, because, in fact, not all bits are equal. The text portion of an e-mail message may be far more important to a user than the full resolution of graphical or video elements to that same user. In a network in which capacity is highly constrained, both users and operators need to work cooperatively to extract maximum connectivity benefit.

While proponents of net neutrality rules premise their calls for regulatory intervention on claims of "fairness," the current net neutrality proposals as applied to wireless networks ironically will accomplish the opposite: a small number of users will be able to adversely impact a much larger number of users. The optimized solution, one which matches best services against network capabilities, *is one that recognizes that not all bits are the same and manages the mobile traffic accordingly. To do otherwise will eviscerate the very consumer-centric capabilities that technologies like LTE make possible.* For example, consider an application, perhaps a social networking one that simultaneously communicates video and voice. If the network becomes congested, the approach consumers may prefer would be for the network manager to prioritize voice quality, especially as the video component consumes far more bandwidth than the voice component. Such an approach would allow the listener to continue to comprehend the speaker, albeit with a fuzzier picture. In contrast, an approach that mandates equal treatment of bits, would force the degradation of both voice and video elements, resulting in voice quality that would quickly become unintelligible.

With appropriate network and traffic management, operators will be able to offer service plans and the network mechanisms that enable communications that can dynamically respond to what the network can deliver in ways that most effectively meet consumers' needs and requirements.

Device Impact on the Network

Eliminating a carriers' ability to optimize devices attached to its network will directly affect the user's experience, the network performance, and the device performance. According to the proposed regulations, "Subject to reasonable network management, a provider of broadband Internet access service may not prevent any of its users from connecting to and using on its network the user's choice of lawful devices that do not harm the network." This approach might be feasible if devices simply worked or did not work. But in wireless technology, there is continual improvement in device capabilities which brings both improved and entirely new mobile devices to market practically every month or quarter. This benefits not just the users of the improved devices, but all users on the network.

An example is the Adaptive Multi-Rate (AMR) codec that was developed for GSM networks. A codec is the algorithm that takes a human voice signal, after it has been digitized, and encodes it in a much smaller number of bits using sophisticated compression techniques. This encoded stream is communicated over the air. Standard voice digitization occurs at 64 kilobits per second (kbps), but the encoded and compressed stream is closer to 8 kbps. AMR defines multiple voice encoding rates, each with a different level of error control. The AMR codec dynamically uses the most efficient coding rate based on the radio conditions at that time. Once implemented in a large percentage of devices operating in the network, voice capacity in the overall network doubles. This means twice as many users can make simultaneous phone calls.[16] The benefits of AMR are:

1. Lower service costs for users since the network is more efficient and the operator can offer more competitive plans.

2. A lower probability of blocked calls, because the network can handle more calls.

The operators can control migration to AMR by supplying handsets that implement the feature. If older devices do not support it, the operator can incentivize users to upgrade their equipment. But under the proposed net neutrality regulations, a carrier could not prevent a large percentage of users from attaching devices without AMR, which will preclude the benefits of increased capacity, even for users who had the latest equipment that did implement AMR. Thus, although the theory of "any device" sounds very democratic, in fact the sub-group of wireless subscribers who attach their own non-AMR equipment to the wireless network will be curtailing everyone else's experience and mitigating their own, because of the overall less efficient network they've created.

There are numerous other examples of advances in technology improving network capability. One is better device antenna technology such as being able to receive the signal from the tower on two antennas in what is called mobile receiver diversity. This technique provides a 50% improvement in network capacity once it is implemented across all devices. Users that bring devices to a wireless network without this feature would retard the capacity improvements spawned by mobile receiver diversity.

Other radio innovations such as Multiple Input Multiple Output (MIMO), equalization, and higher-order modulation – all of which boost overall network efficiency – are beginning to proliferate in the wireless broadband industry landscape. Maximum network capacity enhancements, however, only occur when these kinds of innovations are embedded across all devices connecting to the wireless network. With the wild west of devices envisioned in the NPRM, it is possible that more of these innovations will be stymied, leaving consumers without the truly cutting edge services and applications that Americans desire.

As these innovations are implemented, not only do individual users benefit from higher performance, but all users benefit from the more efficient network, as shown in Figure 6.

[16] For further information about AMR and other voice capacity enhancements, refer to white paper by Rysavy Research, "Voice Capacity Enhancements for GSM Evolution to UMTS," July 18, 2002.

Figure 6: Importance of Performance Improvements in User Devices

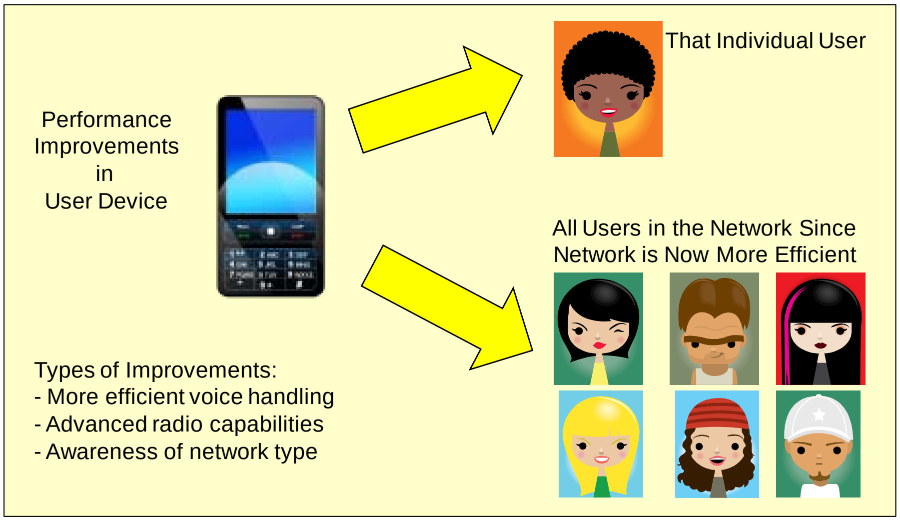

There are other improvements planned for wireless networks that will require coordinated and managed release by operators. The more users can attach their own sub-performing equipment, the less efficient the networks will be. Given the significant capacity constraints of wireless networks, it is essential to extract all inherent capability. Some examples of other planned improvements include:

- **Circuit Voice over Data Channels.** One way of boosting capacity in wireless networks is to run circuit-switched voice channels over the data channels in the radio link. This is a stepping-stone to all-IP networks, but unlike voice over IP (VoIP), the voice call runs through the existing voice infrastructure beyond the radio connection. The capacity improvement is a consequence of the data channels being more optimized and hence running at higher spectral efficiency than the voice channels.

- **VoIP Phones.** In current 3G, evolved 3G and eventually 4G phones, there are architectural approaches that enable voice to be carried extremely efficiently as VoIP. The eventual voice capacity using VoIP is higher than with current circuit-switched approaches.[17]

- **Voice Support in LTE.** In LTE, there are a number of ways that voice can be implemented. The most efficient way is the most complex, requiring additional infrastructure called the IP Multimedia Subsystem (IMS). Some initial LTE deployments will use a simpler approach called circuit-switched fallback. In this approach, voice calls use the underlying 2G or 3G networks in a circuit-switched mode, relegating LTE for just data. Over time, as operators enhance their network infrastructure, they will realize greatest economic efficiencies by handling all their traffic in the IP domain on their LTE networks.

[17] VoIP requires a coordinated release of multiple enabling technologies. Full efficiencies are not realized until supporting algorithms and techniques are implemented.

- **Software Interfaces.** Operators today have control over devices and can mandate that certain software interfaces be supported. Using these interfaces, application developers can make intelligent choices in their wireless applications. For example, if an application can make a query as to network type, it can automatically defer higher-volume transfers to a Wi-Fi connection. User-provided equipment would not necessarily support the software interfaces required for more intelligent application behavior.

In all these cases, if user equipment fails to implement the latest approaches, it will undermine the ability for the operator to offer the most competitive and efficient network. In worst-case scenarios, it will force operators to keep supporting obsolete and expensive modes of operation. By employing the most efficient radio innovations and protocols as soon as they are commercially viable, operators can maximize the capacity of their spectrum.

Conclusion

The success of the mobile-broadband industry is a highlight of the economy. But its very success threatens to overwhelm the capabilities of the networks. As consumer demand for broadband on wireless networks increases exponentially, American carriers continue to invest hundreds of billions of dollars to upgrade, enhance, and expand their networks to ensure consumer demand can be adequately met. The fundamental issue is that wireless networks have highly constrained capacity while at the same time user applications can consume almost any amount of bandwidth. This is why just a few users, with applications like YouTube or Slingbox, can consume the entire capacity of the channel. From both a technical and engineering point of view, operators need flexibility to manage the finite resources of their networks by using a wide range of approaches in order to provide service to all customers. All of these management and engineering approaches can be implemented in ways that promote competition, enhance, and do not harm, the consumer experience, and are fully transparent. These approaches, however, could either be undermined or eliminated by network neutrality regulation.

While this paper has addressed identifiable operational and engineering impacts of the imposition of network neutrality rules to wireless networks, equal consideration must be placed on unidentifiable effects. As discussed, the wireless sector continues to innovate at a furious pace. The very essence of much of the innovation, both at the radio level and at the application level, is the ability to respond in real time to changing operational conditions, usage patterns, technology deployments, consumer needs, and new network requirements. Think of the medium as both highly intelligent and highly adaptive. From an operational and engineering perspective, the imposition of the proposed net neutrality regulatory framework on the wireless environment is fraught with systemic uncertainty, and runs an extremely high risk of creating unintended engineering and operational consequences for both residential and enterprise consumers, and increasing the overall cost to subscribers of their wireless broadband services.