An Engineering Assessment of Select Technical Issues Raised in the 700 MHz Proceeding

Prepared for

Free Press

Media Access Project

May 2007

Columbia Telecommunications Corp • 10613 Concord Street • Kensington, MD 20815
301.933.1488 • www.internetCTC.com

Table of Contents

Table of Figures and Tables

I. Introduction

This Report presents the results of an engineering evaluation of some of the issues raised by the Further Notice of Proposed Rule Making with respect to the 700 MHz auction currently under consideration before the Federal Communications Commission.

The Report was prepared in May 2007 by Columbia Telecommunications Corporation (CTC) at the request of Free Press, Media Access Project, and Consumers Union. The Report addresses some of the technical issues raised by the debate over how to allocate spectrum in the upper tier of the 700 MHz band to best serve the public interest and to make viable emergence of a "third pipe" broadband alternative to cable modem and digital subscriber line (DSL) services. Specifically, this Report:

- Describes how a broadband wireless network operates
- Describes how an "open access" broadband wireless network operates and discusses how such an open access network differs from a network limited to only one service provider
- Discusses how an open access network enables greater engineering efficiencies (and therefore, better use of scarce spectrum) than a single-provider network
- Discusses how allocation of larger blocks of 700 MHz spectrum is more likely to result in a technically competitive "third pipe" because it will enable greater speeds and lower costs per bit for both deployment and operations
- Briefly describes the flexibility of the open access architecture to allow for integration of smart radio technologies such as real-time spectrum auctions in the future as such technologies emerge

II. How a Broadband Wireless Network Operates

A broadband wireless network consists of three major parts – the radio access network (RAN), the backhaul network, and the core and interconnect network. Figure 1 provides a schematic representation of such a broadband wireless network.

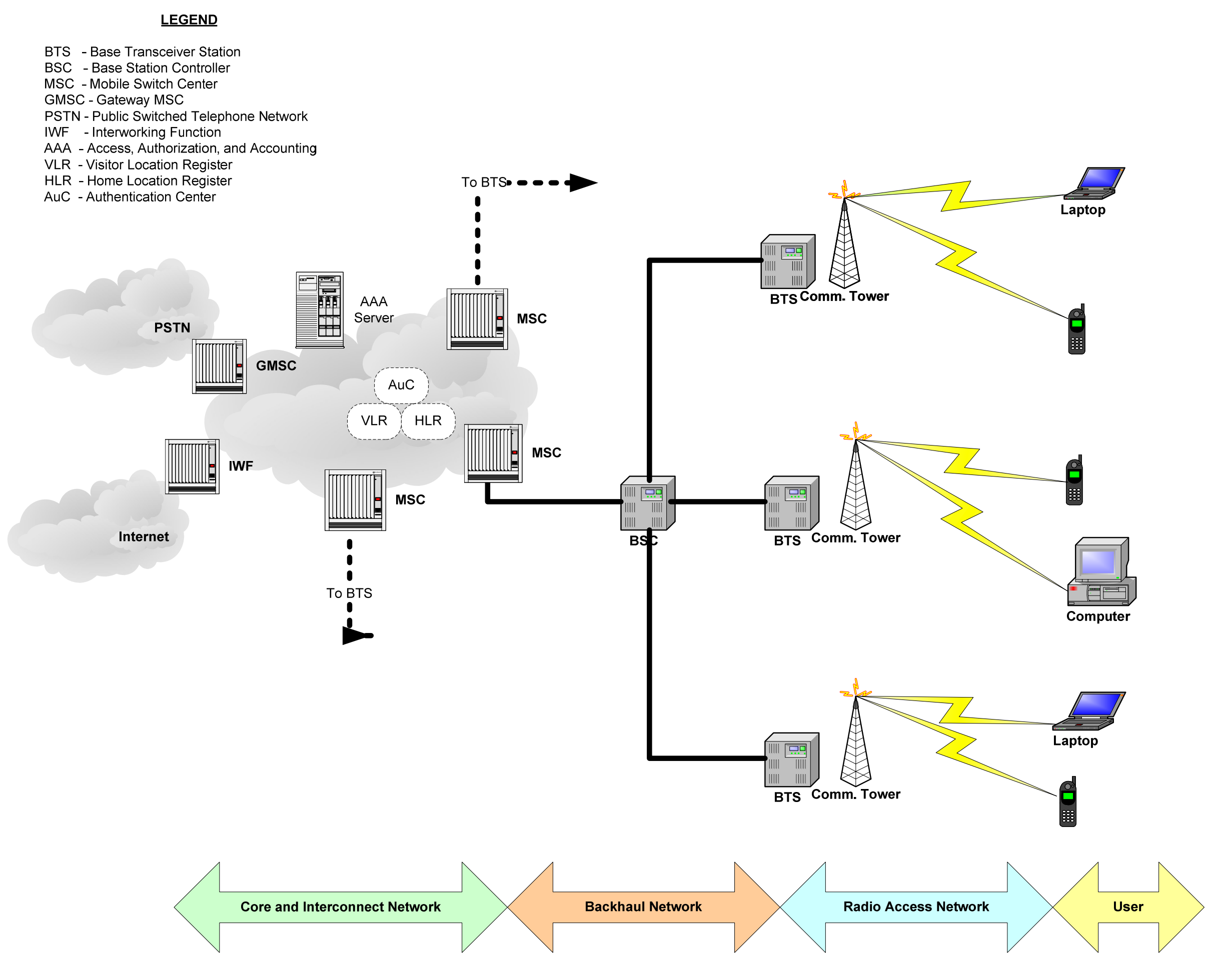

Figure 1: Key Parts of a Wireless Network

The user interacts with the network by using customer premises equipment (CPE) compatible with the network. Examples of CPE include mobile telephones, personal digital assistants (PDA), laptop computers with internal or external antennas, and fixed wireless CPEs installed at homes and businesses. Generally, the CPE is owned by the user or leased from the wireless operator. In order to connect with the network, the user must have an account or be otherwise authorized to connect. The network performs authentication of the user by matching the identity of the CPE (usually a serial number) with a particular account and that account's status.

The user's CPE makes a connection with the RAN segment through the radio interface between the CPE and network operator antenna. Antennas are located on monopoles, on rooftops, on utility poles, and in building facades. The spacing of the antennas depends on the design of the system and the capacity and RF coverage requirements of the area. More antennas are needed in busy areas, as well as in areas where there are physical obstructions, such as terrain and buildings. Typical antenna spacing on a cellular/PCS/broadband network providing broadband wireless services is 1.5 kilometers in a suburban area.

The base transceiver, located close to the network operator antenna, manages the radio communication, including power levels and available transmission speed over the link. The wireless transmissions to and from multiple base transceivers are then aggregated at a base station controller site, located regionally. The base station controllers are finally linked to a mobile switching center, where the central switching intelligence is located.

The connections between the antennas/base transceivers, the base station controller, and the mobile switching center together constitute the "backhaul" network. Generally, the backhaul requires fiber optic technology because of the high capacity needs of a broadband wireless network. Provisioning high-speed fiber optic services every mile is one of the most costly and logistically challenging aspects of building a broadband wireless network.

The service provider switching centers interconnect the wireless service provider with outside networks, including the Internet and the public switched telephone network (PSTN).

A. The radio access network

The radio access network provides wireless communications between user devices such as mobile phones, laptops and home modems on one end, and the network operator antennas on the other end.

The network provides service to fixed, nomadic, or mobile subscribers over an air interface. Multiple base transceiver stations (BTSs) are deployed over the service area. The BTS provides the radio equipment needed to establish communications over the air interface with web-enabled phones, broadband wireless-enabled laptop or desktop computers, home transceivers connected to an external antenna, and other broadband wireless devices, together known as customer premise equipment (CPE).

Some BTS can provide service using multiple spectral bands. BTS could also be multi-protocol and provide services simultaneously using multiple wireless communications standards (for example, CDMA, GSM, WiFi, and WiMax).

B. The backhaul network

The backhaul segment of the wireless network interconnects the network's wireless antennas with the network's mobile switching center.

The backhaul network consists of transmission and switching equipment that interconnects the BTS to the core of the network. The base station controller (BSC) interconnects multiple BTS and aggregates traffic. The mobile switch center (MSC) connects to the BSCs and provides the switching functionality and processing required to authenticate and set up connections to the rest of the parent network, including the public switched telephone network, the Internet, VoIP networks, and other wireless networks. Multiple-T1 and Ethernet links are typically used to interconnect MSC and BTS. As the need for backhaul capacity increases, higher capacities can be provisioned using T3, optical OC-N, and high-speed Ethernet technologies.

C. The core and interconnect network

This segment of the wireless network serves as the wireless network's backbone and also connects the network to other wired and wireless networks, including the Internet and the public switched telephone network.

The core and interconnect network consists of the interconnections between multiple MSCs, gateway MSCs, and the servers and databases such as the home location register (HLR), visitor location register (VLR), and authentication center (AuC) with Access, Authorization, and Accounting (AAA) servers that are integral to mobile network operations. The transmission links between MSCs are typically fiber optic.

III. Open Access Can Effectively Be Engineered Over a Broadband Wireless Network

The general organizational scheme described above is no different in an open access network. The only difference is that in an open access environment, multiple entities connect with this infrastructure at the switching centers, and subscribers to many service providers are able to connect through the same infrastructure.

From an engineering standpoint, wireless broadband networks are well-suited to open access arrangements. An open access infrastructure can be engineered and operated in many different ways using existing technologies and could offer such far-ranging services as voice, text and picture messaging, broadband access, and video.

Wireless open access is achieved by appropriately engineering the wireless network. As more users enter the network or as demand increases in particular geographic areas, the infrastructure service provider can increase coverage as needed so long as the infrastructure is built to be scalable. For example, the infrastructure owner can enhance the RF capacity of the existing network simply by adding or segmenting base stations, which would have the effect of enabling more potential users in a particular geographic area to connect at higher speeds to the network. If the network backhaul is built with fiber optics, upgrading network capacity to a given base station can be as simple as upgrading the network's electronics to higher-speed technologies.

Figures 2 and 3 below illustrate the similarities and differences between carrier-grade wireless networks for single providers (Figure 2) and multiple providers in an open access environment (Figure 3). Figure 2 shows the system-level configuration of a standard wireless broadband infrastructure that supports a single wireless service provider (WSP) over the proprietary network. The network infrastructure consisting of the RAN, backhaul, and core network are owned and operated by the spectrum licensee network operator. In addition, the same network operator owns and operates equipment supporting voice, data, and video services. All paying end-users of the wireless network are customers of the same spectrum licensee network operator. This general configuration is consistent with that used by wireless broadband carriers Verizon Wireless, Cingular, and Sprint/Nextel.

Figure 3 shows a system-level configuration of an open access wireless broadband infrastructure that supports multiple wireless service providers (WSPs). This configuration is described below. This configuration compares most closely to carrier broadband wireless architectures and is illustrated here to demonstrate the compatibility of carrier networks with open access. There are several other suitable architectures, including use of mesh and hotspot technologies and of WiMax.

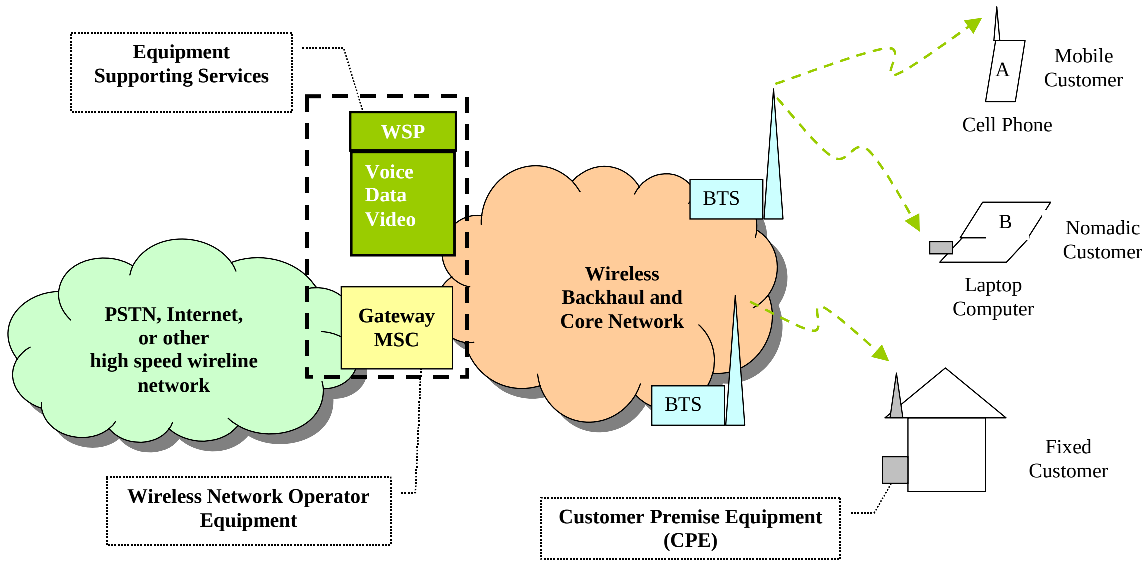

Figure 2: Generic Wireless Broadband Network, Supporting a Single Service Provider

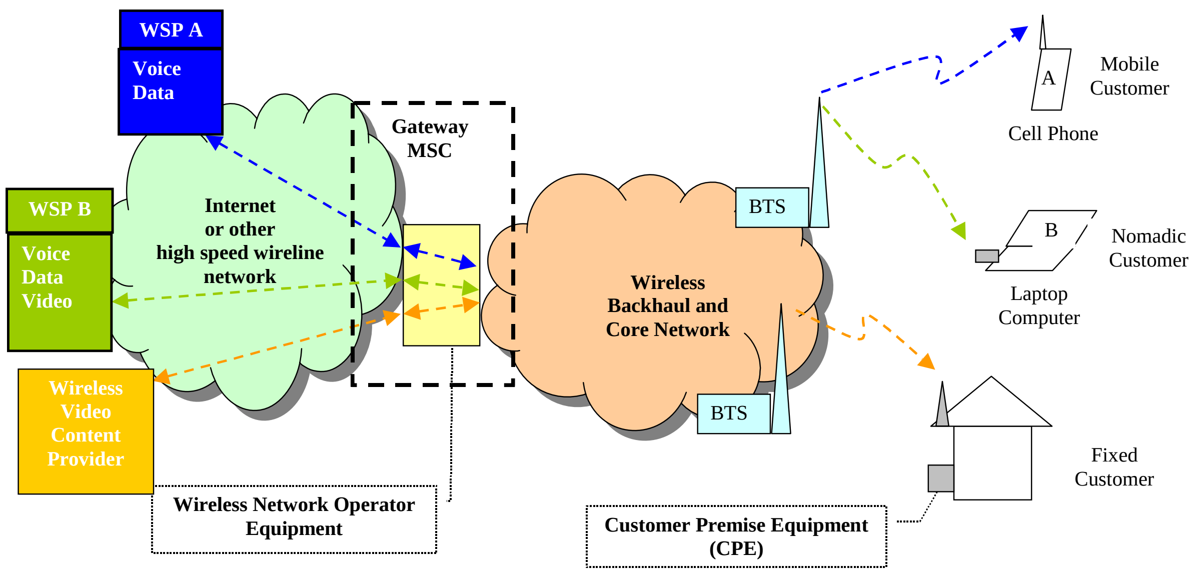

Figure 3: Generic Open Access, Supporting Multiple WSPs over a Single Network

In the example in Figure 3, each wireless end-user connects to the radio access network with one of three distinct service providers.[1] The wireless infrastructure operator owns the wireless network consisting of the spectrum license, antenna, the backhaul and core network, the BTS, BSC, MSC, GMSC, and access to the towers. The customer premise equipment (CPE) could be owned by the service provider (this equipment could include

---

[1] Three providers are used for purposes of illustration. There is no theoretical technical limit to the number of service providers.

the WiMax antenna and radio) or the end-user (this equipment could include a cell phone or broadband wireless card on a laptop).

The wireless infrastructure operator provides transmission services at wholesale rates to the various service providers. Service level agreements between the network operator and the service provider would dictate service attributes such as the number of users support, maximum bandwidth supported, and quality of service. The wireless service providers could offer a bundle of services including telephony, video, and Internet access. In this example, WSP A offers a triple-play bundle of services, WSP B provides only voice and data, and the third WSP is exclusively a video content provider. Subscribers could select one service provider to get all three services, or pick services from different service providers. In this example, subscribers A and B obtain services from WSP A and WSP B, respectively. Subscriber C, however, only obtains video services from the third provider.

Open access in this model compares closely to existing, operational open access networks in wired communications, including provision of services from multiple retail service providers on cable modem, DSL, and fiber optic networks in the United States, Europe, and Asia.

IV. Open Access Enables Greater Engineering Efficiencies than Allocation of Spectrum to Individual Providers

Open access conditions can make highly effective use of the available channels, if thoughtfully and thoroughly developed, enacted, and enforced. Open access makes highly efficient use of the spectrum relative to a band plan with multiple providers on multiple channels in the following ways:

- First, efficiency arises from the sharing, by multiple providers, of a single platform with a single set of antenna structures, base stations, backhaul, management systems, and RF designers.

- Second, efficiency arises from full use of all allocated spectrum on an open access platform. In contrast, a scheme with multiple bands assigned to many individual providers will result in a greater loss of spectrum use because of the need for guard bands and mitigation of RF interference among the many individual providers/bands.

- Third, in an open access environment, each provider can offer a higher theoretical maximum speed to their customers – their speeds are determined by the bandwidth of the *large* shared spectrum block, rather than the bandwidth of the *smaller* channel. In a well-designed network, multiple users would regularly approach a peak speed more than twice as fast as the speed they would reach if the spectrum were simply split between two or more providers (Figure 4). The technical workings of this efficiency are described in greater detail in Section V below.

- Fourth, there is no technical limit to the number of retail competitors on an open access channel, while the non-open access band plan limits competition to the number of blocks provided for commercial use.

In an open access environment, there are a number of key technical arrangements to be managed in the relationship between the infrastructure provider (who manages the RF platform and backhaul network) and the retail service provider. From an engineering standpoint, none of these issues is problematic. The arrangements to be managed include:

1. Requirement for minimum level of RF coverage;
2. Demarcation and peering between infrastructure providers and retail providers at access points in multiple metropolitan areas;
3. Requirements for RF or capacity enhancement based on dropped communications or request of retail provider;
4. Ability for any standards-based hardware to have access to the platform; and
5. Ability for a customer with standards-based hardware to sign up with any available retail provider in that region.

Some of the more challenging technical matters include:

1. Developing and enforcing standards and rules for access to the network;
2. Setting enforceable performance standards for the infrastructure provider; and
3. Determining and prioritizing RF coverage areas and RF mitigation and scaling of capacity.

V. Allocation of a Large Channel is Most Likely to Result in a Technically Viable Third Pipe

From an engineering perspective, the goal of a "third pipe"—high speed performance at a competitive price—is best served by the largest channel, and the way the Upper 700 MHz Band can best deliver a third pipe is with availability of paired channels of 10 MHz or larger.

A. A large channel can better compete technically with existing wired services

In our opinion, 10 Mbps per customer is the current baseline for a viable third pipe, particularly in metropolitan areas where cable modem and DSL services are widely available.[2]

It is important to note, however, that by the time the 700 MHz networks are built and become operational over the next decade, the baseline will have risen in metropolitan areas where cable modem or fiber-to-the-premises service is available, perhaps as high as 20 or even 50 Mbps.[3]

A network built on an allocation of 5 or 5.5 MHz paired channels of spectrum will reach peak aggregate speeds[4] of ~12 Mbps, shared among users in a particular service area—speeds that are technically competitive in today's environment, but likely not that of 10 years from now.[5] Only allocations of 10 MHz or more paired channels can potentially provide peak speeds of ~24 Mbps or more shared capacity among users in a particular service area– speeds that will be able to compete with DSL and cable modem service in the future years when the 700 MHz service actually becomes available, assuming the spectral efficiency of existing and foreseen technologies.

With time, cable and DSL providers will continue to increase speeds by upgrading electronics and physical plant. The wireless providers will need to continue to upgrade their technologies to remain a competitive third pipe.

---

[2] Cable companies can provide more than 10 Mbps per customer using existing DOCSIS 2.0 technology and phone companies can offer speeds comparable to the current generation of cable modem through advanced DSL services (where they are available). Though this may be the baseline for technical competition in metropolitan and densely-populated areas, these networks and speeds are not available in many parts of the country, particularly rural and low-density areas.

[3] Cable modem service speed is expected to reach high as 100 Mbps with adoption of new technologies such as DOCSIS 3. Emerging fiber-to-the-premises networks are capable one Gbps, and even the long-awaited fiber-to-the-node networks (such as AT&T's U-Verse) may offer 10 Mbps or more if they are successfully built.

[4] "Peak speed" represents the maximum theoretical speed possible over a data network. Peak speed is a calculation equal to the spectrum bandwidth multiplied by the maximum spectral efficiency of the system. The peak speed, minus losses due to overhead, will be divided among the simultaneous users of the spectrum.

[5] This analysis assumes the use of existing CDMA broadband wireless technology, with placement of base stations at reasonable intervals analogous to network segmentation for PCS/cellular technologies.

Table 1: Peak Speeds Possible under Different Channel Allocations and With Other Technologies

	5.5 MHz Wireless Allocation	11 MHz Wireless Allocation	Cable Modems		ADSL
			DOCSIS 2.0	DOCSIS 3.0	
Upstream RF Bandwidth Available	5.5 MHz	11 MHz	Up to 6.0 MHz	3.2 MHz to 24 MHz or greater	200 KHz
Downstream RF Bandwidth Available	5.5 MHz	11 MHz	6.0 MHz	6 to 24 MHz or greater	2 MHz
Spectral Efficiency (Peak Speed / Spectrum Bandwidth)	2.4	2.4	4 to 7	4 to 7	6 to 12
Maximum Theoretical ("Peak") Downstream Speed	12 Mbps	24 Mbps	43 Mbps	160+ Mbps	24+ Mbps

B. A large channel will offer proportionally higher speeds than smaller channels because the theoretical maximum speed scales with the size of the channel

As a matter of engineering and physics, allocation of larger channels will enable proportionally higher peak speeds compared to smaller channels.

Peak speed is a critical parameter in evaluating the performance of any data network. It is an important number because it can be quantified (and compared to other network peak speeds) more easily than average speeds, which will vary dramatically depending on location, time, usage pattern, application, and many other factors. Peak speed is a useful way to gauge the maximum capabilities of a network and therefore is one way to evaluate how best to allocate scarce spectrum. Absent authoritative figures on average speed, peak speed is one way to determine which channel allocation is most likely to produce a viable third-pipe -- one that may compete technically with DSL and cable modem service. Where we discuss peak speed here, we are not suggesting that these speeds will be available to all users at any or all times; rather, we are recommending that potential network performance can be gauged be comparing peak speeds because this speed will be available in the aggregate to all users.

From an engineering standpoint, the only way to increase peak speed available to any individual user is either to allocate more spectrum or to improve spectral efficiency. "Spectral efficiency" refers to the network's ability to use the available bandwidth to obtain the highest speed.[6]

A shared data network connection, as long as it is not saturated, provides an actual transmission rate greater than the total capacity divided by the number of users (mean speed) even if not quite at peak speed.[7] This is true because of the bursty[8] nature of data and Internet traffic. As a result, a 5 MHz channel block cannot achieve comparable speeds to a 10 MHz channel block. Hypothetically if a 5 MHz can provide its users with speeds of 7 or 8 Mbps, a 10 MHz provider can potentially provide speed of 14 to 16 Mbps in the same environment and same conditions, using the same technology.

Figure 4 illustrates how speeds scale up in proportion to increases in channel size, and scale down in proportion to decreases in channel size. Each illustration in the Figure shows a total allocation of 11 MHz—allocated in its entirety to one operator in the first illustration; and divided between two operators (5.5 MHz each) in the second.

---

[6] Spectral efficiency is the speed the system can offer using a certain amount of bandwidth; in other words, the data rate/speed in Mbps using a particular spectral bandwidth in MHz.

[7] This is true whether over wireless or cable modem, both shared media.

[8] "Bursty" is defined as the constantly changing bandwidth or bandwidth requirements that are typical of certain Internet applications and traffic. Bursty traffic is not constant in nature, but has a wide, dynamic range from very low to very high bandwidth utilization. For example, browsing a web site rich in multimedia content will require constantly changing bandwidth, varying in bandwidth need from 0 kbps while the user reads the page, to 30 kbps when the user downloads a new text page, to 5 Mbps or higher for a short time while downloading a large image or audio file. In contrast, streaming media (audio/video) typically does not require bursty bandwidth – it has very constant requirements (a 768 kbps videoconference, for example). Although streaming and multimedia applications are becoming common, the majority of Internet traffic remains bursty and variable in its capacity demand.

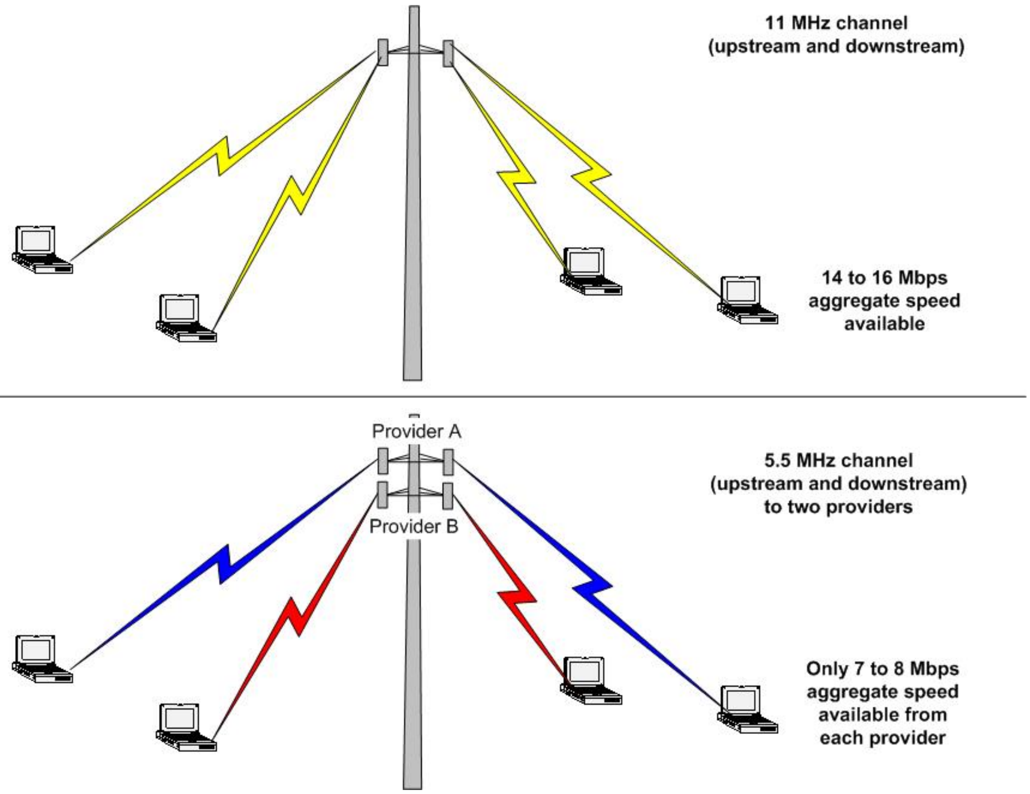

Figure 4: Transmission Speed Scales in Proportion to Increases (or Decreases) in Channel Size

Spectral efficiency will continue to improve over time as wireless technology develops and improves, but improvements in actual user speeds will be limited by the many factors that reduce wireless performance, such as interference, terrain, building obstructions, and protocol overhead.

Providers will be able to serve maximum speeds to more customers over time by using advanced antenna technology and reusing spectrum by segmenting their systems, *but this will not increase the maximum theoretical speed available to any individual user—only more bandwidth and greater spectral efficiency can do that.*

C. A large channel enables engineering and network efficiencies, bringing down the cost per bit

A band plan with a large channel will minimize the capital cost per Hz allocated (and therefore the cost per bit) because a single infrastructure (with a single set of antennas, base stations, backhaul, management systems) can be used to deliver many bits per second to many end-users. This band plan will enable construction and deployment of just one network, engineered to allow many providers to use it, rather than many redundant networks, each engineered for just one provider. The efficiencies of the former are obvious, and make the capital and operations cost *per bit transmitted* significantly lower than the latter.

A technically-competitive third pipe will require construction of a network with high density of antennas and the capability for intensive spectral reuse, no matter what channel bandwidth is adopted. Construction of a network will therefore require considerable capital expenditure and effort.

To the extent that open access facilitates service provider competition without requiring each service provider to construct a separate network, it is providing the benefit of price and performance competition without imposing these monumental capital expenditures on each service provider who wishes to access the capacity.

If a service provider must build a new network to activate a separate channel band, the cost of the activation may be millions or tens of millions of dollars in a single metropolitan area. As discussed above, carrier broadband wireless architectures may require base stations every 1.5 kilometers. Individual base station costs vary widely depending on environment and the needs of a particular area, but are on the order of magnitude of $100,000, plus ongoing lease fees, and several months required to plan and obtain permitting. Backhaul costs are significant, with $50,000 to $150,000 required to build a mile of fiber optic cable, or thousands or tens of thousands of dollars per month required to lease comparable capacity from a service provider. Hotspot and mesh technologies may be less expensive on a per-base station basis, but would require many more antennas to operate.

D. A large channel can be more easily and efficiently engineered than can combining multiple smaller channels to achieve comparable functionality

As discussed above, the spectrum needed for communications speeds required for a technically competitive third pipe exceeds all but the largest channel blocks under consideration. Therefore an alternative may be for the service provider to provide service to users simultaneously over multiple smaller channels. The network would then allow an individual user to achieve a peak speed comparable to the speed achievable over a single large channel. For example, a service provider could potentially combine its spectrum in the 700 MHz, PCS, and AWS bands and use unlicensed spectrum.

However, there are technical complexities in a single customer simultaneously connecting over multiple bands widely separated over the spectrum, and this is not typically done by commercial broadband wireless providers. It is not realistic at the moment for a single user to simultaneously upload or download over multiple spectral bands using current technologies.

Some service providers currently use multiple bands to provide cellular and PCS services to users, and CPE equipment is capable of connecting over multiple bands. At any given moment, the individual user is only using capacity in one band for broadband data connectivity. Existing software and hardware would need to be enhanced to enable a user to simultaneously upload and download over two or more bands.

It is therefore difficult and costly to seamlessly create the same capability over multiple small channels as with the larger channel. Combining spectrum in this manner is possible if the same entity has licensed multiple smaller channels in the same geographic area or if two or more licensees arrange from a technical and legal perspective to enable users to simultaneously connect to the two entities.

Technically, combining communications over multiple bands requires that the network and end-user CPE be designed to enable a single user to *simultaneously* use two or more systems at once, which requires synchronization between 1) the end-user CPE, 2) multiple sets of base station equipment and 3) the service provider's (or multiple service providers') switching equipment.

The user CPE and all other CPE served by a particular base station would need to establish separate connections to base station equipment for each channel band in use. Base station components are optimized for the bandwidth and spectrum of the channel band in use, so separate transceivers are required for the separate bands.

Base station equipment as currently deployed does not interact with the data stream at the level of data packets and users—that is the role of the service provider's switching equipment, which, as shown in Figure 1, is located regionally. The role of the base station equipment is to forward the data stream from the "air" to the service provider network, and vice versa, not to interact with the data or perform a combining or splitting role.

Alternatives include: 1) designing the network with a single user's multiple data streams traveling separately between the CPE and the regional switching equipment and 2) designing the base station equipment to become significantly more complex in order to read, combine, and separate information in a particular data stream between two or more channel bands.

In the first alternative, the data must arrive at the regional switching equipment close enough in time for them to be combined at the switching equipment. The requirement is particularly stringent if the communication is time sensitive (voice traffic, for example). The switching equipment must be able to recognize that both streams belong together and seamlessly combine them. As data is sent from the network to the user CPE, the network must do the opposite—it must effectively split the data at the switching equipment for apportionment between the multiple channels, forward them separately to the base station or base stations to the user CPE, and the user CPE must be able to synchronize and combine the two streams.

While Internet Protocol and other networking technologies are designed to enable communications streams to travel simultaneously over multiple paths, the problem is made more challenging with broadband wireless communications, because a user may be moving across base station service areas, the communications link quality may be variable or poor, and the communication may be time-sensitive (voice).

The complexity may be greater if the multiple channel bands are operated by separate service providers, requiring commercial arrangements as well as technical arrangements.

If the separate channel bands use separate technologies, for example, 700 MHz and satellite, the complexity is compounded. Satellite transmission latency imposes additional limitations on combining communications. The service provider would need to design its network to compensate for latency, otherwise it would be unable to effectively provide time sensitive communications as an effective third pipe. For example, the satellite link may be used for transmission in one direction and the terrestrial link for transmission in the other direction. Another alternative would be for the service provider to optimize the CPE and network software at the application level—long, continuous uploads or viewing of continuous multimedia streams may travel over the satellite link, while VoIP, short messages and browsing communications travel over the terrestrial link.

As propagation in the higher-frequency bands used for PCS, AWS, and unlicensed wireless (1.9 to 2.5 GHz) will not be as favorable as in the 700 MHz band, building penetration and terrain will affect separate bands differently. As a result, providers must ensure they have constructed base stations for high-frequency at higher density, including technologies such as indoor microcells, or the higher frequency bands may become unavailable in certain environments. As a result, the user may lose the use of one or more of the channel bands, resulting in degraded or dropped service.

E. Proposals 1 and 3 are most likely to result in a viable third pipe because they allocate the largest channels

Given the driving factor of channel size, a viable third pipe is most likely to be achieved under Proposals 1 and 3, which make available a C Band with 11 MHz in both the upstream and downstream directions.

To achieve efficient spectrum use and maximize capacity, the next best proposal is Proposal 2, which allocates a 6 MHz channel pair. Proposals 4 and 5 are least likely to serve the goal of a meaningful third pipe because 5.5 MHz channels are the largest available under those proposals.

VI. The Open Access Design Allows for Innovations Such as Real-Time Auctions as Those Technologies Emerge

One advantage of this open access architecture is its flexibility to incorporate new technologies as they emerge. The architecture does not preclude integration of innovative ideas, such as real-time auctions of bandwidth and other "smart" radio technologies, as these technologies develop and mature to carrier-grade levels.

Within the open access model, real-time auctions can be integrated in the future at two separate levels of the network: first, among retail service providers within an open access spectrum allotment, within the framework introduced in Section III; and second, between the service providers in the open access band and those in other spectral bands.

Given the policy goal of emergence of a third pipe broadband alternative, we would recommend that these emerging technologies be integrated in future iterations of the network. The open access architecture proposed above is based on currently existing technologies, deliberately selected so as not to slow the potential roll-out of technically competitive services as soon as is feasible.

Appendix A. About CTC

Columbia Telecommunications Corporation is a public interest communications consulting firm, specializing in business, policy, and engineering consulting services for public sector and non-profit clients. Since 1983, CTC has worked with the full range of existing and emerging communications technologies to provide services in strategic technology planning and deployment; communications network assessment and implementation; and project management.

During that time, CTC has provided communications engineering and other consulting services to such jurisdictions as San Francisco, Los Angeles, New York, Washington, DC, Seattle, Milwaukee, Cincinnati, Pittsburgh, Philadelphia, and San Jose—as well as numerous other communities. We have assisted many of these jurisdictions to plan, negotiate, and deploy state-of-the-art broadband networks – and to maximize public and community benefit from communications projects. As the technology and business models have evolved, our work has evolved to include numerous community broadband networks—both wired and wireless—throughout the country.

As a matter of policy and in order to provide clients with independent and unbiased advice, CTC declines any financial relationship with telecommunications and cable carriers.