UNITED STATES OF AMERICA

FEDERAL COMMUNICATIONS COMMISSION

NATIONAL BROADBAND PLAN

THE BROADBAND AVAILABILITY GAP WORKSHOP

TECHNICAL PAPER NO. 1

Washington, D.C.

Thursday, May 6, 2010

PARTICIPANTS:

Panelists:

ROBERT CURTIS

ROHIT DIXIT

KEVIN KING

TOM KOUTSKY

BJ NEAL

STEVEN ROSENBERG

JOSEPH SOBAN

Speakers:

PHIL BELLARIA

DAVE BLESSING

TOM BROWN

JOE GILLAN

REBEKAH GOODHEART

CHUCK HELLER

BOB LOBE

CHRIS MILLER

MICHAEL PELCOVITS

LARRY SPIVEK

LYNN STANTON

SHARON STROVER

P R O C E E D I N G S

(3:02 p.m.)

MR. CURTIS: Thanks to everybody for coming. I'm Rob Curtis. I'm the director of the Deployment Team for the National Broadband Plan.

Before we get going, a couple of brief announcements. First, as I'm sure you all know, it's been a pretty busy news day at the FCC, but the purpose of this event is to focus on the data that applies to the broadband plan and not the Title 1-Title 2 question from earlier today. Inquiries on that matter should be addressed to Jen Howard. Statements regarding that issue are on the web.

Second, for those who are online, questions can be e-mailed to leo.fitzpatrick@FCC.gov and on twitter to #bbplan.

I guess we'll start here. So today we want to take you through a quick run of the $24 billion availability gap. We'll start with a quick overview of the gap itself, and then we'll drill down into a few key areas. We're not going to be able to cover all the analyses today, so we've selected a few areas that we think are key to understanding the whole. The goal today is to take a first step towards making a reasoning transparent in hopes that we'll facilitate and induce suggestions for improvement, and we'll try to save a bit of time for questions at the end.

Before we get started, I wanted to take a few minutes to introduce you to the team. Back in July Blair asked me to build a team to address our national deployment issues. I had no idea what a talented team would be able to put together in a very short period of time. My deepest thanks to them all.

To my right is Steve Rosenberg, who's going to be leading a lot of the presentation today. And behind me back there in -- well, I'll go down in the order -- is Joseph Soban, B.J. Neal, Rohit Dixit, Kevin King, Rebekah Goodheart, skip the guy in the white shirt, let's go Bellaria, Tom Brown, and Tom Koutsky.

Without these guys' help, there's no way we could have gotten done what we did in, you know, the six to eight months of the time we had.

Finally, also, thanks to Jim Stedman and Mark Guttman, and all the folks at CostQuest who really made this possible.

So we were asked to provide an analysis to the most efficient and effective mechanisms for ensuring broadband access by all persons in the United States. Answering that question required us to solve two qualitatively different problems.

First, we needed to determine the number of the unserved as well as their proximity to current broadband infrastructure. This step required the creation of a baseline model using the best available infrastructure and availability data.

Second, we needed to determine the level of funding to induce operators to deploy ubiquitous broadband. This step required us to create an economic broadband model and make a number of assumptions about the evolution of the industry that would induce a deter entry. This is a journey and not a destination, and we look forward to your feedback and suggestions for input.

It's worth highlighting a few key principles on the right side of this page. First, we didn't build a cost model, we built an NPV, or an economic model. We included revenue.

Second, we focused on incremental economics. What this means is that, fundamentally, we assumed the extra cost and the extra revenue of deploying broadband on top of what already exists, and we assume the existence of current infrastructure.

Finally, we built a granular model to ensure that we accurately captured the economies of scale with the diseconomies of scale.

This map shows broadband availability at the county level where blue is high availability and red is low, calibrated to the 4:1 availability target. The total unserved by 4 down, 1 up across all counties is about 5 percent of the population, or seven million housing units and 14 million people. This chart disaggregates the $24 billion investment gap and represents a 20-year explicit model without a terminal value. A couple of things here are worth noting. First, this output is the result of dozens of discreet model runs in combinations of many assumptions about the way the market might evolve. We try to highlight a few of these key assumptions as we go through today's discussion.

This is a side note: It's not your typical cell model, and this may come up in the context of how we get input on it. It doesn't easily run on a PC.

Second, as you can see from the inclusion of revenue in the calculation, it's not a cost level; it's an economic model that disaggregates cost and revenue as well as initial from ongoing cost. One implication of this disaggregation is that the model is fairly insensitive to the discount rate or cost of money since the initial cost isn't discounted and only the spread between ongoing cost and revenue is discounted.

We affectionately call this our "hockey stick rush". The left axis is density, the right axis is gap, and the horizontal axis is percentiles of unserved census blocks. It clearly depicts the high correlation between density and cost, and that's a theme that will repeat itself many times today.

Also notice the steep slope of the cost curve which highlights the importance of a granular model that accurate captures diseconomies of scale.

This map shows the geographic distribution of density where darker shading corresponds to greater density. To state an obvious fact, the east is considerably more dense than the west, and try to keep this image firmly in mind as we now look at the next page.

There are two things going on in this map. First, look at the rolling darkness of the colors, and keep in mind that darker is more expensive. If you compare this chart with the previous page, notice the dark colors tend to be in low density geographies indicating that cost is driven by density.

Second, look at the different colors red versus blue. Red indicates that we estimate DSL is cheapest; blue indicates we estimate that wireless is cheapest. As we'll explain a bit later, what tends to drive the color difference is terrain. In flat areas, cell radii can be large, which drives down the cost of fixed wireless. These flat areas tend to be shaded blue for wireless.

In more mountain areas, cell radii can dramatically shrink driving up the cost of wireless and the cost of DSL down. So stepping back, DSL tends to be least expensive where density is low and terrain is hilly or mountainous. Wireless tends to be less expensive where density is higher and terrain is flatter.

Although we estimate that the area to be covered by DSL and fixed wireless is similar, we estimate that 90 percent of the unserved housing units could be most efficiently served with fixed wireless. The over-indexing to wireless is predominantly driven by higher population density in flat areas.

This chart shows the present cost of each technology cut by density, but, importantly, it doesn't include revenue, and it doesn't include the effects of terrain. Notice that fixed wireless and 12-kilifoot DSL are least expensive across the entire range of densities modeled.

In addition to considering the cost of a particular technology, we also took a longer-term view and considered the cost of upgrading the higher capacity networks if and when the need arose. What we found is that by building a future proof network immediately is likely more expensive than paying for future upgrades as the need arises, largely due to the time value of money. With one exception, the analysis behind this chart tried to account for the salvage value of the preceding network. The exception is the case of fixed wireless to FTTP, which is second from the left.

In one sense, there's no salvage value since little fixed wireless infrastructure would be part of a fiber to the premises built. On the other hand, there may be significant salvage value in that a mobile broadband network would continue to provide value after its life as a fixed network is past. This value is not captured in these calculations and may be significant.

This chart is a reprise of a chart we put out in September and demonstrates that the investment gap is highly dependent on the speed of broadband considered. The gap ranges from 1.5 megabits -- or for $15 billion for a 1.5 megabit target to $320 billion for a 100 megabit speed target. As Steve will explain shortly when he drills down into the base case, the $24 billion gap is a combination of the two 4-megabit-per-second cases that you see on this page.

Steve?

MR. ROSENBERG: Thank you. I want to

take as little bit of time to dive deeper into the most expensive areas, so, as Rob said, the areas with low density/high gap parts of the hockey stick shown here in a slightly different format -- let me orient you to this graph -- first off, each vertical line represents one county.

So what we've done is aggregate up the census blocks to the county level. There are roughly 2,800, a little bit more, counties represented on this page from lowest-gap county to highest-gap county. The vertical access you'll see goes up by an increment of 10 each step up the graph. The reason there is to show some of the differential at the lower end of the hockey stick, and you show it as you did before. You just get that same sharp increase at the end, and you can't see much of the variation.

What you see over on the left-hand side is some very low-gap counties down around and below $10,000 for the entire country indicating a small number of relatively inexpensive homes to serve with broadband all the way up to the right-hand side where you see very high-gap counties in excess of $200 million for some of them. This is indicative of a combination of very expensive housing units and a relatively large number of them.

You also notice here the color shading. The dark blue areas are areas where the initial investment is smaller than the gap. What that means is if you build this network through some combination of private and/or public financing, what you find is that that network will operate profitably. On the other hand, the light blue areas are ones where there is an ongoing gap, where the total gap exceed the initial CAPEX.

Even if you fully subsidize a network billed in that area, even if the network were handed over to a private operator free of charge, they could not operate that network profitably; it would need ongoing support in order to continue operations. You'll note it's not a one-to-one correlation, but, generally speaking, the higher-gap counties over on the right-hand side are lighter blue. So the larger the gap the more likely you are to end up in an area where you need some level of ongoing support.

Worth pointing out, by the way, this is for unserved housing units only, so this is not representative of the country as a whole. Generally speaking, areas that have service are profitable to operate in, and you in fact won't have any gap whatsoever.

Looking now at another map, this one is focusing on the number of unserved. We're going to continue talking about the most expensive, but I want to orient you again to a slightly different map. This starts off here with the number of unserved in each county, so we've taken the percents chart of unserved that Rob showed at the very beginning, the red to blue color bar, and converted this now to the absolute number of unserved. And what you see is on the coasts darker shading representing more unserved in each county.

What that means is that even in a county that has a high level of service, over 90 percent, in counties that are very high population you may end up with a relatively large number of unserved housing units there. So, for example, in Southern California, you'll see some dark blue shading in areas that are both according to our work and the work done by California State nothing effort a very high percentage of service, you see a relatively large number of unserved housing. What you have is a large county with a large population, and out West there, in fact, a large geographic area as well, so a large number of individuals.

If you look at this map, what you see is that there is a wide dispersal of unserved across the country, and an overlay onto that now a series of yellow dots. What these represent are the areas of highest gap. So what we do is we calculate the gap, and we'll talk a little bit more about this. We calculate the gap for every census block in the country, and order all of those census blocks of unserved housing. I think it's order of magnitude two million census blocks that have some -- that have unserved housing units in them.

We then take the -- take those, order then from highest gap to lowest gap, and then figure out where are the 250,000 most expensive housing units in terms of having the highest gap. And that's what's represented on this chart. Each of those yellow dots represents a census block with at least one housing unit. In total, there are 250,000 housing units spread out about 75,000 census blocks. There are 75,000 dots there.

If you look -- if you recall that hockey stick, this represents about 3-1/2 percent of all unserved represented by these roughly 70,000 to 90,000 dots. If you then look at the gap, removing those 250,000 housing units from the equation, you get a very different picture. Instead of a gap of $23.5 billion your gap is down to about $10.2 billion. What that means is that those 250,000 housing units spread largely across the Midwest, the middle part of the country, particularly the upper Midwest, represent about $13 out of the $24 billion gap. If you are able to serve those through a technology that does not have the same cost structure as terrestrial-based technologies, potentially with satellite, as suggested here, the gap could be reduced to about $10 billion.

So that's an overview of the answer. If you have had a chance to read the paper -- and I'm guessing that not many people have made it through the whole thing -- what you'll see is that this is basically going through chapter 2 -- chapter 1, I guess it is now -- of the paper, just reviewing some of the results. What we're going to do now is talk a little bit about the approach, about how we build the baseline and some of the decisions we made in building the financial base case.

So we start off, as Rob said, needing to understand how many housing units there are in the country that lack access to broadband service, in this case at 4 megabits down 1 megabit up. It's worth pointing out and being very explicit that the data that exists are far from ideal. Over time we hope to improve that situation. We are going to be happy to incorporate more data as we get it and incorporate it going forward, but we had to build a baseline based on the data that we have available right now to us.

So what did we do? First is we took a fundamentally different approach to different technologies to build the baseline. So let me walk through that starting with 100 percent of the country. We find that about 90 percent of the country have DOCSIS available. That is based on a commercial data source where we looked at the franchise boundaries for cable companies that offer two-way cable service. So this isn't just cable television, this is two-way cable service.

There are some risks here. Anecdotally, we hear that the commercial database does not necessarily include, or does not indicate any information, about when a cable franchise does not build out to the edge of the franchise boundary. So it is possible, certainly, that particularly with small and mid-size cable companies, that the cable franchise is not built out fully, and we may count some areas as served that do not currently have service.

At the same time, we've also heard that these cable databases do not include some of the smaller systems, and so we may tend to underestimate the availability of DOCSIS-based broadband in some areas.

Moving down now to the third row where only DSL is available, so I'm going to focus on the area where DSL is available. We do an analysis of where DSL is available nationwide, but we're really focusing in on those homes, those housing units that had DSL available, but are outside of a cable footprint because those that are within the cable footprint we already consider having broadband service available.

As we look at the Telco footprint, the data here are particularly difficult to work with. There is no commercial dataset analogous to the one available for cable and, in fact, there are very few data sources at all. What we did is based on a number of state mapping efforts that were completed prior to our starting this work. So there's a large number of state mapping efforts underway right now through NTIA as part of the BDIA work. Obviously, those data are not available yet. We hope to get those in the near future and update our work as appropriate.

What we did do is we looked at the states that had completed mapping efforts, and we put a pretty high bar on the state data that we wanted to incorporate. So we limited ourselves to states that had broken out broadband by technology. As I said, we were treating cable and DSL differently. We wanted to be able to tell those apart.

We limited ourselves to states that conducted analysis on a fine enough geographic level to differentiate service levels. So if a state had, for example, done an analysis of DSL availability at the county level, we felt like that was too coarse to really give us information about where service was actually available within the county footprint.

We also eliminated states where coarse assumptions had been made about the availability of DSL. So, for example, particularly in some of the early state efforts there were assumptions made about, say, drawing an 18-foot circle around the location of known DSLAMS. We felt like that that would not be indicative of real service availability, so we excluded those states. And we also did not look at states that had only information about DSLAMS located in CLs. So we know that DSLAMS can be in RTs. Looking only at COs would tend to systematically understate availability or at least have that potential.

So what that means is that we have a limited number of states. The good news is that we have faith that state's -- this data we used -- are accurate and indicative of availability. The bad news is we're basing our work on a limited number of states.

Given the state data, what we then do is take all of the data for that state and do, run a massive regression analysis. So we take the availability information to census block level by speed where it's available. And we understand that one of the things that correlate most with broadband availability, we then run a regression to quantify that and apply that to the rest of the country. So it is our best way to take data in areas where we have it and apply it to areas where we don't.

Standard statistical analysis applied about making sure that the effort was predictive of sample data, however, as it says here on the page, there are a couple of risks associated with this approach. Number one, as with any statistical approach, you will tend to overestimate availability in some areas and underestimate it in others. In aggregate, we believe that those will wash out, that we have every reason to believe that the errors have a normal distribution, so to speak.

The other potential risk is that the data on which we base our analysis is not representative of the entire country, so if you fundamentally believe that the states for which we have data look very different from others, you may have reason to believe that the regression analysis is not going to be predictive.

Having said that, where we have seen complaints about lack of service and we look at our maps, we haven't looked at anything yet and said, well, we really missed the boat on that. But, quantitatively, we can't rule out that possibility. We need more and better data to really understand how predictive this is, and then when we have the better data we'll use that instead. So this, we recognize, is something that we had to do in the short term to really understand availability.

The results are what you see here. As I said, 90 percent of the country as DSLAMS-3 available, and so that's the large bar on the left-hand side. We code that here as 10 megabits per second. In any case, it is certainly above the 4-megabit-per-second target rate that we have for download speeds.

Then you see distributed across the rest the speeds that we estimate a network is capable of delivering based on distance from either known or forecast locations of DSLAMS. What you see is in some ways counterintuitive. A large BOLIS at 6 megabits, and then much smaller number for 4, 3, 1-1/2, and then the submegabit categories. But what that reflects is that the area out to several thousand feet -- I think for 6 megabits it's about 9,400 feet -- is much, much larger than the rings between 9,000 feet and 12,000 feet, and between 12,000 and 15,000 feet. So the areas of that innermost ring is by far the largest. Then when you overlay the fact that DSLAMS tend to be in areas that are of high density as opposed to out in sort of the suburbs or the exurbs of even small towns, this starts to make a bit more sense.

So that's sort of an overview of how we built the baseline and what the results are. I'm going to shift now and talk about the financial base case, and just to make sure to be super explicit about this, when we talk about "the baseline," we are talking about where service is currently available, where the infrastructure currently exists. Base case is the financial metric.

The base cases Rob alluded to was actually a very complex series of data runs, so I want to walk through that in a little bit of detail just so people can understand what assumptions we made. First off, one of the major complicating factors that we face is the likely presence of 4G out in the country in future years. It doesn't exist now except in a relatively small number of urban areas, not in the unserved footprint. It's currently only offered by Sprint and Clearwater through WiMAX, but there are commitments from a number of carriers to roll out 4G in the next few years. We wanted to make sure we accounted for that.

The reason is that we believe that while 4G, commercial 4G rollouts that are targeted towards mobile broadband and mobile telephone use, are not in and of themselves going to be sufficient to be a fixed replacement. In other words, commercial 4G build-outs are going to require incremental investment. We do believe that some number of people will be able to get a fixed replacement service with 4Gs. So if you think about a 4G build-out, you will have some number of people close enough to wireless bay stations to be able to have a high enough signal strength to get 4:1 service.

Absent incremental investment, we do not necessarily believe that that signal strength will be high enough everywhere to be a fixed replacement, and so we're going to calculate that cost. But the fact that there are a number of people, perhaps a majority of people -- we don't know -- who will have 4G service that is fixed replacement, we wanted to account for that both in the impact it will have on wireline deployments and an impact it will have on wireless. So let me walk through what I mean by that.

Start with the top rotor a the top of

the page looking at 12,000-foot loop DSL. As Rob mentioned, 12,000- foot loop DSL and wireless solutions are always the lowest cost, so we are focusing on them as part of the base case. For 12,000-foot loop DSL. In 4G areas, we are going to assume that there is a single competitor. That's the 4G operator.

As I said, we don't know that the 4G operator will be able to offer service to everyone in the footprint, however what we didn't want to have is a situation where we assumed no competition a 4G player offering service that actually took money away from the wireline provider and not accounting for that. What that would mean is that the wireline provider could have worse economics than we would model if we assumed zero competition. So we assumed one competitor in 4G areas. Outside of 4G areas we assume no competition.

What we are then able to do is calculate the gap for 12,000-foot loop DSL for every census block in the country that has unserved --

obviously, that they have unserved there is no doubt -- and then roll that up to the county level for DSL. So we have a county-level gap for the 12,000-foot loop DSL.

Let me move down now to the bottom of the page. I'm going to focus in on the eight-mile radius wireless, and I'll talk a bit about the radii in a second. A similar situation, a division between 4G areas and non-4G areas, here the issue is if it is a 4G area, if a carrier has deployed 4G service, they are going to have lower costs. They already have a certain amount of infrastructure in place that they can leverage for offering the fixed service. So what we want to do is account for that.

We also want to make sure -- as Rob mentioned, this is an incremental model -- we want to make sure that we are only accounting for incremental revenue. So in 4G areas for wireless, we only looked at fixed revenue, so the revenue for providing fixed service not mobile service, and we only allocate -- we only looked at 73 percent of the cost. That 73 percent -- we go into it in the paper -- is a cost based on the amount of traffic driven by fixed service. So what we do is we say fixed cost and fixed revenue in 4G areas.

Outside of the 4G footprint, we're assuming it's a Greenfield build, so all revenue, both mobile and fixed, and all costs is accounted for in the model.

We make that calculation for each of four radii: Eight-mile, five-mile, three-mile, and two-mile. The reason we do that -- Rob is going to talk about it in a minute -- is to account for terrain variation and making sure that we have an appropriate cell size for different geographies. But we do the calculation for each geography at each radius. We then choose a cell size and whether it's a 4G or non-4G area based on our data. And as I said, Rob will talk about the cell radius choice. And again, we're able to calculate a single wireless gap for each census block. So we choose the right case from among the eight wireless cases that we describe on the left-most part of the page.

We then again roll up the wireless gap in each census block to the county level. That brings you to the octagons -- I think it's octagons -- and sort of the right third of the page we now have a gap for both wireless and wireline, and we can look at both the lowest cost and second lowest cost technology to determine what the gap is in each county.

So as we go through this, we not only calculated the base case, so, as I said, there were just 10 data runs to describe what the base case is and how they're combined. We also looked at what we'll call sensitivity. So different assumptions and how they might impact the outcome, how they might impact the size of the gap. Those are each described in greater detail in the paper, but I wanted to present them all here to give you sort of an overview. I'm not going to walk through each of these.

There are several that we're going to

talk about in more detail coming up, and a couple that we've talked about a little bit, so, for example, the one, top one there, highlighted in gray focus on terrestrial solutions but estimate the potential impact of satellite. As I said earlier, moving from a $24 billion gap roughly to a $10 billion gap, means that there's a $13 billion effect of the accounting for the impact that satellite might have. So that's one we talked about already.

We talked a little bit about 4 megabit downstream and 1 megabit upstream in terms that if you were to pick different levels -- Rob talked about that in the table -- we'll talk a little bit more about that, how we make sure that we deliver 4:1 service.

And then a couple others that we are going to talk about across the next few pages is funding only one network, assuming the second-lowest gap and remaining technology-neutral. So let's dive into these, first off starting with funding only one network.

Obviously, it is possible to fund more than one network in these unserved areas, but it has a pretty big impact on the economics.

On the left-hand side, we talk about what the gap would be, if instead of funding only one network we choose to fund both a wireline and a wireless network. There are good arguments to be made about whether that is a worthwhile thing to do. We just wanted to be able to say this is how much it cost. and that's what's shown on the left-hand side. The gap on one network, as we said, is $23.5 billion. The incremental gap of a second network is an additional $10.7 billion for a total of 34.2.

What's going on here is when you build the second network, you incur more costs and you have the same revenue -- actually slightly less revenue given our assumptions about the effect of competition on revenue. So you're adding cost, you're not adding revenue, and so the gap goes up.

On the right-hand side another hypothetical. If for some reason wireless were excluded from the solution set for whatever reason, be it technology-based, policy-based, what have you, if you wanted to have two competitors both based on wired networks, what would that look like? One of those networks is relatively cheap. That's the gap for 12K, a 12,000-foot-loop DSL, with one competitor everywhere.

The second competitor would not be able to use the same loops that the 12,000-foot-loop provider is using. You only have one set of loops going into the house. If you have -- if you want a second facility as base competitor, they need to build out to the premises. You could use any number of technologies. He will use the fiber to the premises gap on top of that for one competitor, an incremental $67 billion for a total of $87 billion for two wired networks. So anywhere from a roughly $10 billion to $50 billion impact of having more than one network.

Second, assuming the gap of the second lowest-cost technology solutions, when I walked through the base case, I said that we got a gap at the county level. And then we can calculate what is the lowest cost-gap and second lowest-cost gap. So we see here the gap for wireless base network everywhere is about $13 billion. If you just assume we're going to put wireless everywhere, it's $13 billion. Pretty soon we're going to have 12,000-foot-loop DSL everywhere. That's $18.6 billion.

If you now optimize, if you say, well, I can choose the lowest cost technology in every county, so you pick wireless where it's lowest, DSL where it's lowest, the total cost is only $8 billion. The cost of the second lowest-gap technology, so sort of the loser in that head-to-head competition in every county, is up to $23.5 billion; that's the $24 billion gap that we talk about. And we do this, as we mention in the paper, to account for the effects of a market-base mechanism.

If you have two providers, you have about the same cost structure, about the same gap, they're probably going to be at about the same price. If you have two that have a differential, given perfect information the lowest gap provider can offer service or can bid, so to speak, depending on how the mechanism works, at that higher rate. And so that's the one that we take into account here.

The next piece I want to talk about is what we call here 4:1 actual network. So we don't just want to talk about offering a 4-megabit-per-second down service, 1-megabit-per-second up; we want to make sure that a network is actually capable of delivering it. And for this we turn to some data on usage.

On the left-hand side here, what you see is a breakout of the percent of use by different categories of users. And what you see is the heaviest one percent of users account for up to a quarter of traffic on the network. What that means is that if you wanted to calculate the amount of usage, you need to kind of have a sense for how many -- what the traffic looks like across different categories of users.

We talk about that in terms of busy hour offered load, BHOL, and that's what we show on the right-hand side. If you average across all users on the network some of whom are not doing anything, some of whom are engaged in intense streaming, some of whom are doing web-page surfing, across all users what you see is an average use of between 111 and 21 -- we're down to 17 -- kilobits per second. What that means is the average user, the median user over on the far right is using between 17 and 21 kilobits per second on average. As you include more and more users, you actually move higher and higher in the average usage because you're pulling in the seven percent who use 65 percent of the total, the 1 percent who use the quarter. As you add more and more people, you get a higher busy hour offered load.

If you look just at the 90 percent usage figure, you get between 36 and 43 kilobits per second. So what you're saying is 90 percent of users in the busiest hour of the day use about 40 kilobits per second. You also see that that number grows over time, so what we are seeing is that average use is doubling roughly every three years, so this is a shift. This isn't sending twice as many e-mails every three years, it's a shift from e-mail to web surfing, from web surfing to rich web surfing, from low resident video to higher resident video, so from YouTube to Hulu, for example.

So we see growth in the busy hour offered load year on year. We forecast that out of assuming the same rate of growth that we've seen in the recent past, and instead of seeing about kilobits per second you see 160 kilobits per second several years out. So if you focus on delivering a network capable of serving up 4-megabits-per-second down, 1-up, you get 160 kilobits per second to serve 90 percent of users at the busy hour, and that is how we calculated the cost associated with our network.

So we didn't engineer a network here, obviously, but this is the step that we took to account for what engineering might look like in during the accounting, during the financial projections.

We also wanted to make sure that we took into account middle mile and backhaul issues. We know that this can be a pretty important part, so we need a series of conservative assumptions about backhaul that I want to walk through.

What we see on this page is a map of central office locations across the country, about 20,000 of them. These are ILEC central offices only. We focus on ILEC central offices because, honestly, the data is easier to get at, it's easier to know where the ILEC central offices are than where IFC pops are, where cable nodes are, all those things. So we focus on these, focusing only on the locations of ILEC facilities is a conservative point of view because we're omitting a lot of fiber that's out there from these other types of providers.

Each of the COs was coded according to the kinds of services it offers, so what you see here green dots indicate fiber services offered from the CO. We took that as a pretty strong indication that that CO is fiber fed.

We then coded yellow and red dots for COs that offered DS-1 and DSL services, respectively. In all likelihood those are fiber-based, fiber-served COs, but we're not 100 percent sure so we coded them in yellow and red.

And then blue dots represent COs where we do not have any information on the services offered that would indicate one way or the other. So not evidence of absence but an absence of evidence about whether there's fiber there.

What we get is 90 percent coverage, so we see 90 percent of COs have fiber using this analysis. That is probably very conservative. What we have seen filed in the record is that over 95 percent of COs have fiber. What we've heard, anecdotally, is that it's 97-98 percent. So this is almost certainly conservative, but what it shows you is that, again thinking back to the density map that Rob showed awhile ago, especially on the East Cost where there's high density, there's a lot of COs, and there's a lot of green COs in particular.

We then created a color of fiber route map, but I want to be clear on this: This is not a map of fiber routes; this is a calculation of how you might connect up all of the COs that we saw on the last page. So what we did is we understood the homing tables. We understood which COs maps to which regional tandems, and we took a least distant tree approach to connecting every CO to its local regional tandem. Then we ringed up the regional tandems in a way, again least distance routing.

So the actual routing may look nothing like this, but it's indicative of where fiber exists. We did this, number one, because we wanted to see what the map looked like, but we also used this as an input for the cost calculation, using this as a proxy for the cost of providing backhaul services using ILEC fiber. And again, there's a number of ways in which this is probably a conservative estimate of the amount of fiber that's available.

You'll not, for example, if you look at the state of Nevada, you'll note a lot of fiber that just isn't ringed up, and that comes from the fact that we didn't require that COs were joined in a ring, and in all likelihood that is part of what's going on out in the real world. You'll also notice again in Nevada there were a number of COs that, since we don't have information on, we don't fiber them up. So a conservative assumption about where fiber exists and therefore the cost of providing fiber for backhaul services.

Bob?

MR. CURTIS: I think it's back to me.

MR. ROSENBERG: Yep.

MR. CURTIS: So what we're going to do now is drill down relatively briefly on what we came up with as the two low-cost technologies. I'll run through wireless and then hand it back over to Steve to wrap us up, and he'll go through DSL.

So if you look at this page, you'll see that there's a $13 billion investment gap for fixed wireless, and that was the smallest gap of the different technologies that we modeled.

A couple of other things to call out on this page, notice that OpEx actually exceeds the sum of initial and ongoing CAPEX, predominantly doing of the cost of -- I'm sorry, ongoing CAPEX -- predominantly doing of the cost of tower releases. So on this page I'm going to walk you through how we thought about capacity in the wireless network. We took a great deal of care to ensure that the wireless network that we modeled could actually support the 4:1 target.

The primary cost here I ran a wireless network is the cost of the number and the cell sites required to provide both adequate signal density and capacity. So first we determined how many cells would be required to provide adequate signal density to hit the 4:1 target, which in turn required that we solve for cell radius. We assumed in this model a 700-megahertz spectrum, which does have excellent propagation characteristics.

Our link budget indicated that the limiting factor for cell radius was the 1-megabit-per-second uplink as opposed to the 4-megabit-per-second downlink. So we set our radius smaller than would be required to hit the 4-megabit downlink target.

As we indicated before and we'll dive deeper into in a few minutes, terrain is also a very important driver of cell radius. We took that into account. In addition to these coverage sites which we describe on the left-hand side of the page, we also thought about capacity, so we looked for cells that needed to be split to provide additional capacity. We split cells when the total usage on the cell exceeded its capacity.

Cell capacity, as we thought about it, is fundamentally a function of three things: 1, the usage in the busy hour which Steve described a few minutes ago; 2, spectral efficiency measured in bits per hertz; and 3, the amount of spectrum available. As we'll see in a few minutes, it turns out that providing adequate signal density for 1 megabit per second in the uplink results in low subscriber density per tower, which makes the waterless model relatively insensitive to changes in capacity demanded.

Using RF planning tools, we identified different cell radii for different terrain types. Assuming deployment in the 700-megahertz band and the signal density of about 140 db, as required by our link budget for the uplink, we selected a variety of cell towers ranging from eight to two miles radii depending -- ranging from 8 to 2 miles depending upon the terrain type.

Using the detailed terrain data, we got from, you know, a variety of data sources, we varied all the cell radii based on elevation variation. Yellow and green on this map are smaller radii while blue and dark blue are larger radii. Notice two things from this map:

First, cell radii are smallest where DSL is the cheaper alternative.

Second, you can clearly see the effect of the mountain ranges in the Appalachians on cell size and placement.

This chart shows two different measures of the wireless economics relative to terrain. Wireless, it turns out, is highly sensitive to terrain type which, of course unfortunately for us, is fixed and shows why wireless is not the most efficient technology in the West and in the mountains. But despite the sensitivity to terrain shown in this chart, we think our results are not actually terribly sensitive to terrain since we take the mountains and the plains as a given.

This chart indicates that because of the low subscriber density per tower, which we discussed just a few minutes ago, the fixed wireless gap is relatively insensitive to changes in capacity or demand. The chart shows the effect of altering cell capacity by increasing or decreasing spectrum allocation. The chart would look similar if we pulled either of the two other capacity levers, i.e., spectral efficiency or busy hour offered load.

Notice that even if we reduce capacity by 50 percent by cutting spectrum in half in this case, we barely move the gap. The implication is that even if busy hour offered load doubled, the gap impact would be minimal. We believe these sensitivities are, however, somewhat unique to unserved areas largely due to the low density of those areas.

Finally, although the gap is insensitive to spectrum allocation in capacity, it's highly sensitive to spectrum band used. We estimate that moving from the 700-megahertz band to the PCS band would nearly double the fixed wireless gap. And there's a note not shown on this page, if you assume that instead of going to 2-by-10 megahertz you go down to 2-by-5, so, in fact, you get a 4X decrease in capacity, we estimate the gap only moves up to about $14.1 billion or about a billion dollars. So 4X change in capacity creates about a 10 percent sensitivity of the gap.

Steve is now going to take us through

DSL.

MR. ROSENBERG: Thanks. Similar graph now looking for DSL instead of wireless, the total gap, as I mentioned earlier for DSL, 12,000-foot loop DSL, is $18.6 billion. A couple of things to note here:

First, in contrast to wireless, here the initial CAPEX is much higher. Really, you're paying for loop shortening down to 12,000 feet, in some areas several miles of loop shortening as opposed to the tower release that you face with wireless.

Second, the assumption that you can get for one service implicit in here mentioned in the paper is that we're talking about 24 American wire gauge loops using a DSL 2-plus. We do not believe, for a number of reasons, that that is -- that one requires 24 wire gauges as opposed to 26, but, obviously, if you're at 24, you may, in fact, be able to get longer loops as with 24 wire gauge.

Next, I just want to put up a couple of graphs. These are not in the paper, that are taking data straight out of the paper. Just to array, if you wanted to serve just the unserved areas with different technologies, what that looks like, and so you see 15, 12, 5, 3K DSL and fiber depremises, increasing speed across the horizontal and gap to serve the unserved with a truer wire line on the vertical.

The curve fit is not anything terribly meaningful. I don't actually expect the costs to move along that curve, but it is indicative; you can see it turning over once you've shortened your loops to 3,000 feet you're most of the way there. You've done a great deal of the work to getting fiber depremises, the incremental cost between 3,000-foot loop DSL and fiber deprem is relatively small. So if you're talking about staying within the unserved footprint, really the biggest cost is getting up -- or reducing loop links down to about 5,000 feet.

On the other hand, if you now say, well, what would it cost to serve the entire country at these speeds, you get a very different looking curve. This goes back to the table that Rob showed earlier that talked about the gap at different speeds bringing in different numbers of housing units.

So what you see here is as you move up in speed, you move up in the number of unserved all the way up to, when you get to 100 megabits, you know, and you're talking about fiber deprem, in our analysis we didn't pull out areas that were already served by fiber depremises, so it's 130 million housing into the whole country. What you see is that because you're including more and more housing units, the gap actually starts to curve up the other way, much more expensive as you go up in speed because you're drawing in more and more housing units.

So I just want to take a moment to thank everybody, first off, all of you for joining us here in the Commission room. For those who are online, thank you for your attention. I also want to add my thanks to the team. It has been a real privilege, something I've really enjoyed maybe more than I thought I would, so thank all of you, those who are here and those who are not.

A couple of other comments. As you see on the page, there are some -- there's an aniline and paramount outstanding. I don't think they have yet come out in the Federal Register, but we'd love comments and thought in response to those pieces and any future ones that come out on these topics.

And then, you know, when Rob and I started right from the very beginning back in August, we kept saying please give us more data. I don't know if people took us seriously then. I think you can look at this now and understand that when we say that we mean it. So to the extent that you have data that can help our analysis, to the extent that you have information that you feel will improve the work that the Commission is able to do, please make it available to us. I think that we will be able to do good things with it and come to a better more fact-base analysis.

I think we want to open things up to

questions. If you are in the meeting room, I think just go up to the microphone there. Please announce who you are and your affiliation just so we can get that labeled appropriately, and then we're also taking questions online as well.

MR. GILLAN: And since I can't embarrass myself --

MR. ROSENBERG: Whoa.

MR. GILLAN: -- Joe Gillan affiliated with nobody. Can you get a slide 7 -- I think it's 7, it might have been 8?

MR. CURTIS: Eight? Seven? Eight.

MR. GILLAN: That one. That one, yes. The red is DSL, the blue --

MR. CURTIS: Blue's wireless.

MR. GILLAN: -- is wireless.

MR. CURTIS: Yes.

MR. GILLAN: Okay. All right, this is the question where I might embarrass myself: As I understand your analysis, that shows that DSL is basically cheaper and the last wireless is cheaper on the East. But because you're assuming a market-based disbursement mechanism, you effectively calculate the $23.5 billion which I won't round to 24 since $500 million still means something today. Since you model that usually taking the second-most or the second less expensive technology, so in effect, is it true that what you end up with is actually the reverse to that? That you ended up effectively modeling the cost of wireless in the West and DSL in the East?

MR. ROSENBERG: I would say in terms of calculating the cost that's correct; in terms of showing which is less expensive, the map that we have up shows that. I think the big key to remember is that we are not kicking the technology winners or losers. We're just saying we believe both are capable; we believe this one is cheaper in this area and the other one is cheaper in the other area. But what you said is correct.

MR. GILLAN: I wasn't drawing an inference. I was just trying to make sure I understood the core -- the relationship. Thanks.

MR. CURTIS: Yep.

MR. ROSENBERG: I have a question here in from e-mail. Are there any plans to release data on served and unserved census blocks for DSL and cable technologies? The answer is yes and no. We are, I think today we're working with the internal IT group to post the output of the model runs that we have done on the FCC website. All of the data that we are making available is at the county level. The reason for that is, No. 1, just a practical one, putting up several million rows in a database gets a little bit tough to deal with. But also we have denounced this, based in part on commercial data, and we are not able to just release, prerelease, so we cannot release things at the census block level.

Going forward, we will try and figure out the best we can to make sure that all of the data that can be released will be. Negotiating data rights is appropriate.

MR. HELLER: Hi. I'm Chuck Heller from Wilkinson Barker and Knauer. My question, actually, deals with a decision to exclude satellite from the base case calculation, and the capacity issue I understood there's a sentence that appears in sort of similar but slightly different forms in the technical paper in a number of places and also in the plan about the impact of the disbursement mechanism. And I have a guess about what that means, but it was a little -- I was a little unclear on exactly what that meant. I was hoping you could explain it.

MR. CURTIS: Yes, so it derives in part from the capacity limitation. There's a question if you've got a technology that can serve a subset of the whole, what the optimal way to introduce that technology to the entire solution of the problem, do you use it to serve some part of a census block? Do you use it to serve all of a particular census block and none of some other census block? And until, you know -- well, I should say after a lot of deliberation on this topic, what we determined was that until the decision gets made as to how disbursement is going to happen, it's not possible to identify the optimal use of satellite.

So we excluded it from the base case but tried to go to great pains to explain, you know, as Steve did, that it's efficient and useful in a lot of -- as a large part of the solution to this problem. It's just hard to pinpoint exactly how to use it optimally, so that's why it's not in the base case.

MS. STANTON: Lynn Stanton, TR Daily. Could you just explain why you -- the $22.5 billion reflects the second best? That just kind of went totally past me, why you wouldn't figure out the cheapest way to do it rather than the second cheapest way to do it.

MR. CURTIS: Do you want to try that, or do you want me to try it?

MR. ROSENBERG: Yeah, I'll try that one.

MR. CURTIS: Okay.

MR. ROSENBERG: You know, let's use --

MR. CURTIS: I'll pile on.

MR. ROSENBERG: Yeah. Let's just take a

simplified example, so imagine that Rob and I are operating competing companies with competing technologies. I can do something for $10, he can do it for $100. Let me also, for simplicity, assume a disbursement mechanism, don't know if this is what it's going to look like but imagine a world where the government is auctioning off support levels. I could do it for $10, but the government is saying, well, who wants to do it for $10? Nobody wants to do it for $11. Maybe it goes, you know, from the top down, from the bottom up, don't know. But if I have perfect information that I can do it for $10, but it's going to take him a hundred, I would want to do it for $99.99. So I want to hold out for as much as I can.

Now, we may complain as taxpayers that, you know, that's not what we would want to have happen, but with the market mechanism, if there's relatively few competitors, it may be hard to prevent that from happening.

MS. STANTON: Because you're still in a market-base mechanism for determining support.

MR. ROSENBERG: Correct.

MS. STANTON: Okay.

MR. ROSENBERG: Now it may also be the case, and it's been pointed out that if you have multiple, say, wireless providers, you know, we assume a single wireless provider is 700 megahertz. Well, if there were three wireless providers at 700 megahertz, then you would assume that that cost, the 700 megahertz, would prevail because one of the three would do it at that gap.

It's not clear that we have enough wireless providers with that kind of spectrum to be able to do it, and as soon as you get into, say, TCS spectrum, it starts to look a lot more expensive. And so that's why we did the interplay between wireless and wireline as opposed to two different wireless players. Ultimately, it may prove we may have to dig into that sort of wireless piece a little bit more deeply, but for this we wanted to make sure we had a number that reflected both sets of technology.

MR. CURTIS: So and -- sorry, I said I'd

pile on, and I will -- that's exactly right in terms of the technical answer. I think there's also a practical issue which is the reality is there were some 1,500 service providers in the country. They all have slightly different strategies on entry, exit, you know, what their business plan is, do they want to provide fixed wireless? Do they want to build out DSL in these areas? And short of knowing what the, you know, hurdle rates, the business plans simply to be more than say, you know, even if we make money we just don't want to play there. Other people may say, that's in our sweet spot, so we'd go there for less.

We wanted to make sure that we got enough that we were right and did the best we could to think about simplifying 150 -- 1,500 firm, you know, game theory problem in a way that we could realistically capture and model. And so the second cheapest has the virtue of both being consistent with a market mechanism and also ensuring that there is, you know, some reasonable way of accounting for a bunch of different business strategies, none of which were transparent to us as we were going through the process.

MS. STANTON: Thank you.

MR. ROSENBERG: That's fair.

MR. PELCOVITS: Thank you. Hi, Michael Pelcovits, Microeconomic Consulting and Research Associates. I had three questions, if you bear with me, but one you just answered, so I'm down to much less.

But maybe a little follow-up. If you're saying you didn't want to assume two wireless carriers bidding against each other, since we are -- that's really sort of within the nature of a market-type of mechanism -- if we assume wireless is cheapest, why would you not have two wireless carriers bidding against each other to be the one to put out 4G facilities?

MR. ROSENBERG: This gets back to the similar (inaudible) to say about the practical nature. We just do not have the data about where different wireless operators are providing serves are going to provide 4G at what spectrum bands, and the ability and the time we had to analyze that, so we made a simplifying assumption about a single 700 megahertz player and ran sensitivities off of it.

I think you're raising the right question, you know, if we have a market-base mechanism over time, and if we believed there's multiple wireless players, I think we have to look at that to understand what the impact on the ultimate gap will be.

MR. PELCOVITS: If you don't mind, another question. If you go to slide 20, I had a question about the, under the cost of the 4G where you're assigning 73 percent of the costs to fixed, if I understand that. If this is an incremental cost analysis, which I believe is one of the principles that you described initially, is it really true that, if you've already billed or you're already going to build the system for wireless, that the incremental cost to that in capacity for fixed would be much less than its whatever -- 73 percent based on its share of usage?

MR. ROSENBERG: So let me give a multipart answer to that. First off, I've been told by economists that any cost allocation is essentially an arbitrary exercise and, in fact, if my understanding of history is remotely correct, cost allocation battles within Commission proceedings have been somewhat toxic. So fully understand that that is an assumption that will bear some scrutiny.

The reason we did it that way, a couplefold: No. 1, it's certainly the case that without knowing the signal strength or signal density to which commercial 4G deployments are made, it's not clear that you won't have to add more towers for fixed service than you have for the mobile service. So one negative per second up, even with the fixed CPE, requires a pretty high signal density, and it may be the case that a single tall tower that you would build for a mobile deployment will have to have four smaller ones put around it to fill in the signal strength for wire -- for fixed service.

Having said all that, we really wanted to have a reasonable but conservative assumption. Ultimately, going back to the first answer, we recognize that short of maybe Ramsey pricing, it's a decision that reasonable people could differ on. And so fully acknowledge that the high cost attributed to fixed based on traffic is just that: It's high and if you were to do it based on revenue, you might get, you know, the exact opposite of that and much lower cost on fixed services.

MR. PELCOVITS: I mean, I guess it's just a question I have is, if you're doing an incremental analysis, then there is no need for a cost allocation at all.

MR. ROSENBERG: Well, the issue -- and this is something that we weren't able to do, and I'm not sure we would be able to do it with a lot more time. You know, what you really want to understand is, okay, how many towers are built for mobile service? Where are they? How many more towers do you need to supplement that for fixed service? Were they awarded their costs? That would be the ultimate measure.

We didn't have any opportunity to do that. I mean the data don't exist. We'd have been guessing multiple times to do that. I think that would probably be a non-allocated cost method of doing it, what are the actual incremental costs associated? Maybe we'd get there.

I mean one of the big questions in the NO, in the NOI that's outstanding is whether we should be looking at incremental costs or total costs. And if you want to look at total costs, that question kind of goes by the wayside.

MR. PELCOVITS: Okay. If you don't mind, one last question: I think, if I'm not mistaken, there were numbers in here of how small or what the size of the gap would be if you included satellite and there was also an estimate of what the gap would be if you used the lowest cost as opposed to the second lowest cost technology. Have you done an estimate of what it would be if you allowed for satellite and for only having, using, the lowest cost technology? I'll sit down.

MR. ROSENBERG: I don't think we did that. Honestly, the model runs were each model run took between 45 minutes and 12 hours, and there were 10 of those in the base case, and we ran multiple sensitivities. We weren't able to run sensitivities on multiple things at one time and cover the game board, so to speak. I don't think we did that.

My instinct says that the hockey stick will be recreated for any scenario, but that it won't be quite as steep. So my instinct is that, roughly speaking, you will probably save a comparable amounts in percentage basis with the cheaper technology, but I don't think we did that analysis, so I'm not sure.

MR. BLESSING: Hi. I'm Dave Blessing with Parrish Blessing. I think you guys did a really good job at a really hard task, and you acknowledge to your credit that there's a lot of uncertainty and a lot of error that's kind of built into that. And by "error" I'm --

MR. ROSENBERG: Uncertainty?

MR. BLESSING: I mean statistical error, okay? And that's what I'm wondering. Have you guys taken a look at developing kind of a confidence interval around your estimates, both overall by technology, you know, at the county level, or even within a county?

MR. ROSENBERG: So worth pointing out that the people who did the statistics were part of a contractor team, so we are not statistical experts, so if there are any in the room, please forgive the answer I'm going to give.

You know, there's a couple ways that you deal with error in regressions. One of them is to just look at the standard error. Well, even if you have only one state, you've got a couple hundred thousand census blocks as your source data, so the standard error is vanishingly small.

We don't think that that's really informative, so, you know, looks great, set that aside.

We are able to -- another standard technique is to take your source data set, divide it in half, create the regression with the half that you're using, apply it to the other half, see what your predictive power is. With that, I think that -- though keep in mind we did multiple different regressions. It wasn't a single regression, it was a different regression for each speed inside and outside of the cable footprint. So a total of I think it was 10 regressions.

The predictive power of one-half of the data by the other half was between 80 to 90 percent for those, so pretty good predictive power. The bigger question, I think, is, what does it mean -- yeah, when you apply data taken from, say, Pennsylvania and Alabama to Oregon, what does that look like? And until we get data that we have more confidence in on Oregon -- I'm picking on Oregon, I don't if they have a great state now, if it's coming today, so I don't want to single them out.

MS. STANTON: (inaudible), well, I think that now it's like Washington has a great map.

MR. ROSENBERG: Washington is great now. So, you know, we have to get data whether it's from the NCIA, BDIA effort or, as recommended in the plan, revision supports 77 to really test that part which I think is where we stand a greater risk.

MR. SPIVEK: Hi, Larry Spivek from the Phoenix Center. First off, I want to congratulate you on a really great job well done. I think it's been really needed in the telecom debate for a long time, and you had a really difficult task. But I do feel compelled to ask about the elephant in the room.

You guys have come up with a $24 billion number, and to quote our friend Joe, that's a lot of money. And the big question is, is that the FCC has embarked as a major policy endeavor to impose price regulation in the form of open access, and this is going to have an effect on firm profitability and deployment costs, et cetera.

And so my question to you guys is, did you model this? Did you model the open network? Did you model that if you wasn't, you know, existing, would that number, $24 billion, be lower? Or if you didn't, would the number be higher? Or is it revenue neutral? And I think that's important to get out on the table. Thank you.

MR. CURTIS: So it feels like the disclaimer from earlier in the session snuck into the elephant in the room.

We did the best we could to model the status quo network, and, you know, that's -- that's how we think about it.

MR. SPEAKER: (inaudible) this is --

MR. PELCOVITS: It's a follow-up to your question, and it's a lot easier than the last one. On the question of error, estimation error -- and I just want to make sure I understand the report and your model -- as a practical matter, you have estimation error in -- is the $23.5 billion right or wrong?

But when you look at individual counties, you're not really even trying to necessarily predict with precision the investment gap in each individual county. You're saying at the total, once you aggregate across all these counties, the counties were an estimation technique to get at a $24 -- $23.5 billion number.

So when people think about how accurate or inaccurate this is, they really should be thinking about how accurate or inaccurate it is at the aggregate level, not at the street geographic unit. Is that correct?

MR. ROSENBERG: Yes. Your answer was better than mine, so, yes, thank you.

MR. CURTIS: That almost felt like a statement instead of a question, which we applaud.

MR. PELCOVITS: Yeah, thank you.

MR. CURTIS: Yes, it is.

MR. ROSENBERG: It is correct, and though we -- you can look at individual counties and the gaps there, what he said is absolutely correct. The higher level of aggregation the more likely any statistical errors in availability are likely, more likely, to wash out. We tend to try and talk about things at the country level for exactly that reason and do not encourage people to say, well, you've clearly got this one county wrong.

We hope that there's enough census blocks in any one county that that won't happen, but, you know, until we get better data we won't really know.

MR. CURTIS: Yes. I don't know that this is a direct answer to that question, but it seems like it should be said, in terms of the way we estimate the numbers well, and the way we think about the sensitivities. You know, there is a kind of sensitivity where you worry about did we get the cost of the DSLAM card right or wrong? Did we get our two right or wrong? Did we get a take-rate right or wrong? And when you aggregate all these things up with the number of variables that we have, I think in the aggregate those things largely tend to wash.

All right, if you look at the sensitivities on the page that Steve ran through awhile ago, what you see is that there are big swings not in terms of whether the model ran right or whether we got the amount of fiber that you need to put in the ground right. I don't think that that's where, you know, it makes sense for people to spend time thinking; it's sort of the bigger policy questions, one network or two? Our competition questions: Are we going to get two wireless competitors, or are we going to get a wireless and a DSL competitor?

Those are the kinds of places where you see the sensitivity and the sort of inaccuracy in the estimates come in, their market structure, their market dynamic, their, you know, competitive dynamic sorts of questions that are, I think and hope, is really where the action turns out to be.

MR. LOBE: I'm Bob Lobe with Rocco Lobe Sulzer Associates. I have had a whole bunch of questions but I'm trying to limit them. The first one gets back to this $23.5 billion. I was confused between whether it's county or census blocks and how you added, because in some places you added the negative and positives across the county, and sometimes you didn't. So does the $23.5 billion come from a sum of all the negative net present values at the census block level or at the county level?

MR. ROSENBERG: Great question. And it's confusing to us, too, as we make these charts up to make sure we get that right.

What we did is we defined counties, and I'm going to call it a market though markets will be defined by a disbursement mechanism, but just within this room right now let's call that a market. So we assume that a provider will provide at a county level, whether it turns out to be that or something else to be determined.

What we assume, then, is all that you basically sum within a given county, and so if there are positive NPV census blocks within that county, we assume that they will offset losses in negative census block -- negative NPV census blocks within that county.

So, basically, if I choose to serve Montgomery County, Maryland, just across the border here from Washington, I don't know is there's any unserved there, but take that as an example, if it turned out that in the denser parts of the county down by Washington, you were actually able to turn a profit in a couple of census blocks, we would assume that because that entire market is being bid out through the market base mechanism at once, that that would offset any losses of county more less dense areas.

We did not want to assume that if there was an NPD positive county in a state that that would offset losses in another county. In other words, that you might have different carriers in different counties.

So when we talk about the $23.5 billion gap, it is at the county level, is it NPD positive? If it is, then we assume that there is no gap there. We don't include it. If it's NPD negative, we have that gap and that, all of the NPD negative counties add up to give you the gap.

Does that make sense?

MR. LOBE: Sense?

MR. ROSENBERG: Not that --

MR. LOBE: I understood what he told me, okay. The second question, we were getting at incremental, and I was wondering when the wireline side, when you talked about the incremental costs of the wireline, did you run the cost promodel twice, once with it serving everybody and once with it serving only the served people? And that's the difference being the unserved people would be the incremental cost? Is that what -- how you got the incremental cost of DSL? Or did you get the incremental cost of DSL some other way?

MR. ROSENBERG: Some other way. What we did is basically we assumed that a twisted pair copper infrastructure existed, that COs existed, that COs that we knew were fiber fed were fiber fed. And then we modeled where things like roto terminals, DLCs would be. That we just assumed was there, no cost associated with that.

What we then do is, along road paths determine where would you put a DSLAM in order to shorten loops to 12,000 feet, and what would the cost of that DSLAM be building it, operating it, what would the cost of trench and fiber to be to connect that DSLAM back to the CL. So when we talk about an incremental model, it's what infrastructure do we have to build and operate, put in the ground and operate, in order to deliver the service?

So we do not get into the cost of building, maintaining twisted pair copper networks. We don't get into the cost of building or maintaining voice switches or central office equipment beyond central office located DSLAMS.

MR. LOBE: I realize that you're not rebuilding everything. You're assuming some of the places are there. But given that you assume where the wire center is and the poles are, et cetera, do you run the model twice, or how do you get the increment? How do you determine the difference?

MR. CURTIS: Let me take a try.

MR. ROSENBERG: Okay.

MR. CURTIS: So we have the baseline that tells us which households don't have broadband. So five million households. And let's say we're doing the DSL case. We assume there is wired infrastructure from a known CO to that house, and we, based on the baseline, make some estimate based on the regression of what kind of service they have. Maybe they have DSL that gives them 768 but not four. All right, so that implies a DSLAM at some distance from the home.

If you then think about that home, the increment is to say there's a DSLAM at 18,000 feet. The DSLAM needs to be at 12- or 10,000 feet. All right, so what is the cost of moving the DSLAM from 18,000 feet to 12,000 feet, and then what's the -- you know, what's the construction cost? The cost of the move, the cost of the blade, the cost of the chassis, that is the increment of providing the upgrade from 768 kilobits RS to 4 megabits RS.

Does that make sense?

MR. LOBE: Yes. And then my third question gets to your number of unserved houses East. When you went through the slide of talking about when you went from 10 to 5 percent of the DSL, and you ran regressions, you said that there would be a problem if the state you chose were not representative.

I'm wondering if you know that in Pennsylvania there's a Chapter 30 law that says that all the rural carriers already have DSL. So why would Pennsylvania be considered representative?

MR. ROSENBERG: Interesting. I don't know if the Pennsylvania data showed that everybody has service available, and it's worth pointing out that among the unserved only a third are in nationwide's -- in rural carriers' footprints. Also worth pointing out that

Pennsylvania only had data at one speed, I think it was 768, but it may be 1.5 megabits, I'm not sure.

So it's a fair question, and as I said at the beginning when describing this, we would love better data. We would love not to have to do the regression. So if there's more data that we can incorporate, if we can improve the 477 collection, looking forward to doing that.

MR. HELLER: Chuck Heller from Wilkinson Barker again. So when you computed the gap, I know you included revenue. Obviously, this is not a cost problem, you included revenues as well. I know you include video revenues. Did you include also video programming costs? Yes?

MR. ROSENBERG: Yes. So where we had video, and we don't assume video revenue for either 12,000-foot DSL or wireless, but for 3,000-foot, 5,000-foot DSL, FTTP and cable, we do it. I'd have to check to confirm this, but I think it's a percent of video revenue to account for a programming cost.

MR. HELLER: And just very quickly one more, the last question. This model is the vision is this is just sort of for purposes of the Connect America Fund exercise, not a mobility fund exercise in terms of the --the rubric that the plan sets out?

MR. ROSENBERG: Yeah, I draw a slight distinction there, you know. This was for the purposes of the national broadband plan. I think it's to be determined exactly what the model looked like -- looks like for the Connect America Funds. But you're right, this is not getting into the mobility fund that was mentioned in the plan.

MR. HELLER: Thank you.

MR. BLESSING: Dave Blessing. And just to clarify what this gentleman had said, is it your intent that this model is to just estimate the broadband gap, the investment gap? Because one of the primary questions in the NOI is whether or not we need to replace the current USF model with a new model. And, you know, fairly or unfairly, this is the leading horse in that race.

So my question is, is this model -- do you envision this being able to make those distribution decisions down the road for USF funding under the broadband plan?

MR. ROSENBERG: I think we -- you know, No. 1, we need to be careful and not assume that we know what the answer is, given the NOIs and, you know, presumably a series of NPRMs and FNPRMs coming forward. I think that we envision that this kind of approach might be very useful. I think that understanding the cost in revenues at this kind of granular level in building it up could be very helpful for those kind of disbursements. But we don't know that that's the answer.

So in answer to your question, I think, yes, we envision it being possible, that this would be helpful, but we don't want to necessarily assume that that's the case.

MR. MILLER: Chris Miller at Verizon. Just a question on the unserved, identifying the unserved households. Do you think the model could be and, if so, how could be changed or modified to reflect those homes that are being served today but are only served because of existing subsidy streams?

MR. ROSENBERG: You know, yeah, given all the work that we've done, I have confidence that the model can be modified to do just about anything. It's a question of the pain involved in doing so and the amount of data we need to do it.

To me, the answer is yes. I think that gets into whether you would want to look at total cost or incremental cost. To me, what you're asking about is, is fundamentally a question about do we need to have a total cost model to accurately reflect the total support levels needed for deployment? I think that's a fair and open question, and we'd love to get feedback from everybody here and elsewhere on the NOI on that topic.

Any more questions? I see one.

MR. CURTIS: Yeah, we have a question from online, and it's, what assumptions changed after the model was complete and why? So that's an interesting question.

MR. ROSENBERG: You'll be volunteered to answer that one.

MR. CURTIS: Yeah, you know? After the model was complete. That's an interesting question in and of itself. I think that the way to think about that is the way I was responding to Joe's question a little while ago. The kind of the hard assumptions didn't change, as far as I can recall, almost at all. And by that I mean, you know, what is the cost of this, the cost of that, what did it -- you know, how much does it come out in cost to move fiber?

We spent a lot of time as a team thinking about the correct say to think about the strategy problems so that we could most accurately depict what we thought the real economic environment was, and that wasn't changing the assumptions, that was trying to come up with the most accurate, realistic way of describing and modeling the problem. And that was, I would say, the model was complete in the sense that it ran and gave us output, and then once we had a tool, we began to run scenarios and think about, you know, how to change it. And in that sense, you know, we continued to run the assumptions and the sensitivities that we had and have displayed for quite some time.

MR. ROSENBERG: So it looks like we may be winding down now.

MS. STROVER: Sharon Strover, Rural Utilities Service and University of Texas. I appreciate the extent to which you really try to encompass a lot of the heterogeneity in rural regions across the states. I'm wondering if you gave any consideration to cases that some people might consider more special or special: Alaska, Hawaii, American Samoa, and so forth. What was your thinking on those kinds of regions as you put your model together?

MR. ROSENBERG: Let's take Alaska. Alaska's first.

MR. CURTIS: I'm sure you had no idea.

At the risk of something like a broken record, I wish we have better data. So we did not have sufficient demographic data outside of the 50 states, and District of Columbia to model the DSL. So American Samoa, Puerto Rico, we just don't have enough data to do that well.

Alaska's another great example. There are, I think, central offices in Alaska that, according to the data that we relying on, have either no paved roads or no population. I'm pretty sure -- I could be wrong -- but it strikes me as odd that somebody would have built a CO if there's nobody living out there. But that's what the data say. So when we look at the wireline costs associated with Alaska, it's pretty low.

Now, because we're using second-lowest cost, not lowest cost, we don't think that that's going to be a massive swing on the gap, the $23.5 billion gap. But it's clearly indicative that we need to get better data in areas like that.

MR. ROSENBERG: Anyone else? Let me take (inaudible).

MR. CURTIS: That's a cherry-picked area as well, I assume.

MR. ROSENBERG: This question: How will you address cherry-picked areas where unserved areas exist inside a census blocks and tracks? It came in via Twitter.

We have -- the most granular unit of the model is the census block. For those who aren't as familiar with the census terminology as I know and, fortunately, am, census blocks are the smallest units that the Census Bureau supports. There's roughly 8-1/2 million census blocks in the country, roughly 6-1/2 million with population. Most of those are actually fairly small geographically. There's a relatively small number that are greater than, say, two square miles, which is a cutoff the NTIA drew in their agreement with carriers on the BDIA mapping piece.

So we do not pull apart census blocks. As it is, dealing with that much data is, in a model, is pretty challenging. Trying to get down below that to, say, 130 million housing units would be really tough computationally.

So, you know, within census blocks our assumption is that though everybody won't be necessarily clustered in one place within that block, we're seeing a certain homogeneity within the block and not assuming that carriers are pulling that apart. Outside of blocks, you know, block-to-block comparison, we're assuming only that the census blocks get aggregated up in this case to the county level where carriers are making a decision, and so a carrier whose serving would serve everybody within a county, some of whom may be profitable, some of whom might be unprofitable.

MR. CURTIS: Then I've got a question from Twitter: If you offer subsidies and large carriers say no to deployment, will you help establish nonprofit telegroup cooperatives?

I guess the first answer is, you know, we're sort of agnostic about who actually deploys. We simply want deployment. We did what we could to make sure that there was a business case -- that is the point -- and it's to build a business case for someone to enter, and, you know, I think we got that part of it right, and I guess it's worth pointing out that there are at least a few recommendations in the plan that address at the very least municipal entry in places where the markets and market mechanisms, you know, may not be efficient or otherwise fail.

So I think the bottom line is, gone out of our way to try to find a way to ensure entry by someone, large, small, brand new, or what have you. And that's kind of the goal.

MR. ROSENBERG: And that's all we've got from online.

MR. CURTIS: Anybody else? Great. Hey, thank you all for coming. We really appreciate it, and really do look forward to and encourage comments and feedback. It's a beginning, it's not an end. Thanks.

MR. ROSENBERG: Thank you, all.

(Whereupon, at 4:30 p.m., the PROCEEDINGS were adjourned.)

* * * * *

CERTIFICATE OF NOTARY PUBLIC

I, Carleton J. Anderson, III do hereby certify that the forgoing electronic file when originally transmitted was reduced to text at my direction; that said transcript is a true record of the proceedings therein referenced; that I am neither counsel for, related to, nor employed by any of the parties to the action in which these proceedings were taken; and, furthermore, that I am neither a relative or employee of any attorney or counsel employed by the parties hereto, nor financially or otherwise interested in the outcome of this action.

/s/Carleton J. Anderson, III

Notary Public in and for the
Commonwealth of Virginia
Commission No. 351998
Expires: November 30, 2012