CLEAN ENERGY SOLUTIONS CENTER
ASSISTING COUNTRIES WITH CLEAN ENERGY POLICY

Greening the Grid: Implementing Wind and Solar Power Forecasting

—Transcript of a webinar offered by the Clean Energy Solutions Center on 19 April 2016—
For more information, see the clean energy policy trainings offered by the Solutions Center.

Webinar Panelists

Bri-Mathias Hodge National Renewable Energy Laboratory (NREL)
Jennifer Leisch United States Agency for International Development (USAID)

This Transcript

Because this transcript was created using transcription software, the content it contains might not represent precisely the audio content of the webinar. If you have questions about the content of the transcript, please contact us or refer to the actual webinar recording.

Stephanie Bechler

Hello everyone. I'm Stephanie Bechler with the National Renewable Energy Laboratory, and welcome to today's webinar which is hosted by the Clean Energy Solutions Center in partnership with USAID and NREL. Today's webinar is focusing on "Greening the Grid: Implementing Wind and Solar Power Forecasting."

One important note of mention before we begin our presentation is that the Clean Energy Solutions Center did not endorse or recommend specific products or services. Information provided in this webinar is featured in the Solutions Center's resource library as one of many best practices resources reviewed and selected by technical experts.

Before we begin, I'd like to go over some of the webinar features. For audio you have two options. You may either listen through your computer or over your telephone. If you choose to listen through the computer, please select the "mic and speakers" option in the audio pane. If you choose to dial in by phone, please select the "telephone" option and a box on the right side will display the telephone number and audio PIN you should use to dial it. If anyone's having technical difficulties, you may contact GoToWebinar’s Help Desk at 888-259-3826 for assistance.

If you would like to ask a question—and we encourage you do—you can use the "Questions" pane or you can type that in. If you're having difficulty viewing the material through the webinar portal, you'll find PDF copies of the presentation at cleanenergysolution.org/training, and you may follow along as our speakers present. Also an audio recording and presentations will be posted to the Solutions Center training page within a few weeks, and they'll also be added to the Solutions Center YouTube channel, where you can find other informative webinars as well as videos with thought leaders on clean energy policy topic.

Today's webinar agenda is centered around a presentation from our guest panelists. Bri-Mathias Hodge: He's been kind enough to join us to discuss considerations associated with advancing the use of wind and solar forecasts to more efficiently integrate variable renewable energy into the grid. Before the presentations, I'll provide a short overview of the Clean Energy Solutions Center Initiative; and then following the presentations we'll have Jennifer Leisch moderate a question and answer session where the panelist will address questions submitted by the audience, will have closing remarks, and then a brief survey.

This slide provides a bit of background in terms of how the Solutions Center came to be. The Solutions Center is one of 13 initiatives from the Clean Energy Ministerial that was launched in April of 2011 and is primarily led by Australia, the United States, and other CEM partners. Outcomes of this unique initiative include support of developing countries and emerging economies through enhancements of resources on policies relating to energy access, no-cost expert policy assistance, and peer-to-peer learning and training tools such as the webinar you are attending today.

The Solutions Center has four primary goals. It serves as a clearing house of clean energy policy resources. It also serves to share policy best practices, data, and analysis tools specific to clean energy policies and programs. It also delivers dynamic services that enable expert assistance, learning, and peer-to-peer sharing of experiences. And, finally, the center fosters dialog on emerging policy issues and innovation around the globe. Our primary audience is energy policy makers and analysts from governments and technical organizations in all countries; but we also strive to engage the private sector, NGOs, and civil society.

A marquee feature of the Solutions Center is no-cost expert policy assistance known as "Ask an Expert." The Ask an Expert program has established a broad team of over 30 experts from around the globe who are available to provide remote policy advice and analysis to all countries at no cost. For example, an area of renewable electricity policy, we are pleased to have Paul Komor, Energy Education Director at the Renewable and Sustainable Energy Institute. He serves as one of our experts.

If you have a need for policy assistance in renewable energy policy or electricity policy or any other clean energy sector, we encourage you to use this valuable service. Again, the assistance is provided free of charge. If you have any question for our experts, please submit it through our simple online form at cleanenergysolutions.org/expert. We also invite you to spread word of this service to those in your networks and organization.

And now I'd like to provide a brief introduction for today's panelist. First up is Bri-Mathias Hodge, a group manager for Power System Design and Studies here at NREL. He has extensive experience with grid integration of renewable energy and has been involved with a number of wind and solar power forecasting studies.

And moderating today's question and answer portion of the webinar we have Jennifer Leisch. Jennifer is a climate change mitigation specialist with USAID Office of Global Client Change. She supports the U.S. enhancing capacity for Low Emissions Development Strategies program and manages the USAID Greening the Grid partnership. And with those introductions, I would like to welcome Bri to the webinar.

Bri-Mathias Hodge Good morning here in Colorado, and likely good evening wherever you are. My name is Bri-Mathias Hodge and we already received a bit of an introduction there. But I'd like to talk today about an introduction to wind and solar power forecasting.

The learning objectives for this particular webinar are to recognize how wind and solar power forecasting can enhance power system operations with large penetrations of renewable energy, understand how wind and solar power forecasts are produced, distinguish some of the approaches that have been taken in order to implement forecasting systems and collect the necessary data for those systems, and then identify policy and other actions that can support the implementation of these forecasting systems.

First I'd like talk a little bit about some power system basics just so that we're all on the same page. The power system is likely the most complex machine ever created, and one of the reasons for that is because consumption and production must be instantaneously and continuously balanced. This is very important because there's very little storage in most power systems, and electricity is not really a storable commodity in the terms of things like hydrocarbons.

And so why is this important that they're balanced? Well, the grid operates at a certain frequency. Here in the U.S. that's 60 hertz; in some other countries that's 50 hertz. But maintaining system frequency is done by balancing demand and supply at all times. And system frequency is one of the fundamental drivers for power systems reliability, because if your frequency deviates too far from your normal set point you will have both generators and loads tripping off in order to protect themselves.

Now demand—which is load—has always been variable and uncertain. In the new sort of paradigm that we're moving toward within the power system, we're also having quite a bit more variable and uncertain supply, and that is wind and solar generation primarily. So I prefer the term variable and uncertain to intermittent, and the reason for that is both of those provide you with different information.

So, for example, uncertainty is what you—this is where we really utilize the forecast. We can utilize forecast to reduce the uncertainty in the system; however, forecast cannot do anything for variability. You could know perfectly in advance how your wind or solar profile is going to change and you would still have to deal with that variability in the supply, and so it's important to distinguish between the two so that we can come up with effective solutions that address both of those.

Now on the right here, you see an example of this variability and uncertainty from a certain plant. This is a five megawatt, six panel plant in Gujarat, India. So the greed line you see is the clear sky output. So this is the diurnal pattern that solar irradiance would take over the course of the day just from the rising and setting of the sun. Now you can see there on the top figure, this is a day with quite a bit of variability, especially in the middle of the day. So you see power output going very low for this individual site and then rebounding very quickly.

In the bottom figure you see a morning that has almost perfect conditions, so it follows that clear sky pattern almost perfectly, but then you get variability occurring in the afternoon, perhaps due to some overcast weather. So this is just one example of the variability and uncertainty that we're increasing in the system from making not only demand variable and uncertain to supply.

Why is this important? Well, because the balancing of the power system has to occur at a number of different timescales. How the system generally operates is because there are a lot of thermal units which take quite a long time, maybe days or even a week to start up—those are generally your coal plants or nuclear plants—and what's known as the unit commitment process where at least one day ahead we say, "Tomorrow at these certain times we will have these generators turned on."

Now that utilizes a forecast itself, that utilizes a load forecast, and so that load forecast is not always perfect. Because of that, closer to the operating hour we have the economic dispatch. This is where you take units that are online and you adjust their set points up or down based on a revised load forecast.

What you may also do in that time period is start up some fast-starting unit. Maybe this is a natural gas turbine, for example. That gives you some of the flexibility you need just to meet the variability and uncertainly and look.

Now even with this, there's still variability at different time scales that we need to account for. In the seconds to minutes time scale, this is just the natural variability at those time scales and load. This curves because when I wake up in the morning, I don't turn on my lights at exactly the same time every day, or I turn them on at a slightly different time than my neighbors. And so, you have these small fluctuations on those sort of times scales that are known as regulation. So we hold regulation reserve to account for those.

Another type of reserve that we hold is known as load following or possibly flexibility reserve, and this has to do with, say, the morning ramp. Everyone tends to get up roughly around the same time, start turning on lights, cooking in the morning, and so this creates a very sharp increase in the demand. Now when you're going through this sharp increase in the demand, you need ramping capability, and one way to hold that is to have load following reserves ready to provide that.

So I keep talking here about flexibility, and that's because as we move into a new power system with high penetration of wind and solar power, one of the keys is having the flexibility to address that variability and uncertainty. Now there are a number of different options that you can see on this graph. On the X-axis we have a type of intervention. So this is labeled according to a couple different categories. We're going to focus mostly here on system operations which is on the left-hand side; and then you have a relative cost ranking of these.

So, quickly, there are a number of things that we can do to adjust our system operation in order to better accommodate variable renewable energy and to provide more flexibility in the system. I spoke there about flexibility reserves; that's the third point that you see all the way there on the left in the system operations. But what we're going to focus today on is renewable energy forecasting. And the reason for that is as you can see from this graph, while the cost for all of these different options are very system dependent and they are evolving over time, renewable energy forecasting is one of the cheapest and sort of easiest options that could be added to the system to help reduce the uncertainty associated with those generators.

So to dig a little bit more deeply there, changes in system operation and changes in how the variable renewable energy operates can provide low capital cost options; however, we have to remember that these are all occurring in systems operations, and so there may be institutional challenges associated with these. Here, flexibility reflects not only the physical system, but also these institutional frameworks—and we do need to keep that in mind.

Okay, so moving on from the basics of the power system, let's talk about how wind and solar power impact—sorry, rather forecasting how wind and solar power impacts power system operation.

So what is forecasting? Forecasting is prediction. I think everyone's probably familiar with the concept of load forecasting where predicting what the load will be the next day out or an hour ahead out. It's the same concept for wind and solar power.

As an example, we see here a forecast, or two different forecasts of wind power output in Texas. So this is the ERCOT system in the United States. What you see there over the course of one day in green is the actual wind power. So you could see it goes from roughly 12,000 megawatts down a bit to about 10,000 megawatts and back up, and we're sitting at about noon in this time point.

Now in the day ahead, we produced a forecast for this—that's the line that you see there in orange. As you can see, it followed the general trend so far fairly well, however, it was underestimating by about 2,000 megawatts. So this is a fairly large forecast error for this system. What you see in blue is a more recent updated short-term forecast, and you can see that that followed the actual wind power much more smoothly.

Now the key thing I want to emphasize here is with forecasting you can usually produce a better forecast with more recent information, and so this really speaks to the need to update your forecast quite frequently, but it also means that you need to have options for what you can do with your power system at those timeframes, because you can have the best forecast in the world, but if you cannot take an action within your system to utilize that new forecast information, then it doesn't really have much value.

I just want to show another example here about the importance of wind and solar forecasting. So what you see here on the right is a study we performed looking at a very high solar energy penetration in the ISO New England system within the United States. That's the market operator or the RTO for the New England northeastern portion of the United States.

In this figure on the left-hand side, you see the generation. So we have six different sorts of figures here, but they all correspond to the same axes. This occurs over a one-day period, and this is sort of our dispatch stack or a ranking of generators over that one-day period. In the top three figures, this is what happens in the day ahead; and then in the bottom three figures, this is what happens in the real time. So the day ahead utilizes a forecast, whereas the real time utilizes the actual output.

If we look in the top left-hand corner, this is where we have the current state-of-the-art forecast. Now this is a particularly bad forecast today, and what you see is there's a very small amount of solar energy predicted. In the real time, a lot of solar energy is actually being produced.

So what happens to the dispatch stack? You see that a number of your units, chiefly your gap combine cycle unit, are ramped down in order to accommodate that solar energy. The reason for this is solar and wind energy typically have roughly zero marginal cost, and so they're cheaper to operate on a marginal basis than a plant that requires a fuel source. Because of this, we want to incorporate as much of them as possible so as to operate the entire system as efficiently as possible.

In this case, you ramped on your gas CC units, but then when the solar starts to fade we need a little bit more of some of our more expensive units—our gas and oil GT. So there is an economic associated with having this forecast be incorrect. If we go toward a 25 percent improvement and then a 50 percent improvement, we see not only is the system more prepared for all that solar energy that is coming onto the system, but this has a couple different impacts.

One, that very dark red on the top in the 0 percent real time case we see gets smaller as we improve our forecast. This is curtailment. Now this is essentially solar power that we could have produced that we are not producing just because the system cannot accommodate it. So this could be seen as a sort of lost opportunity cost because it does have a much lower marginal cost in the system.

Another important factor here, especially if you're in a utility that procures their solar power under PPA is often under a Power Purchase Agreement you are forced to pay for the generation that could have occurred. So essentially you're paying for that curtailed solar power even if you're not actually utilizing it, and improving your solar power forecast improves the economic efficiency in that manner. So that was just an example what the impact of solar power forecasting can be, and wind power forecasting would have similar impact.

But how do system operators typically use forecast? Now I'd like to stress here that there are a number of different market structures, even in the U.S., but also internationally for power system. Most of what I speak about today is taken from the view of the system operator, so we're looking sort of at the highest level. However, there are some nuances here about how these things would change if you were in a different situation—and I'd be happy to answer questions later about those.

From the system operator perspective, a couple of the different timeframes at which forecast may be utilized are the long-term. So let's say a week ahead. This is really where the forecasting results are quite frankly not as accurate. The reason for that is the atmosphere is chaotic.

It's very difficult to forecast that far in advance how weather systems will move over very large areas. And so what can be utilized here is perhaps not forecast of exactly how much generation you would have, but a rough forecast of, "Will this be a high generation time or a low generation time?" And those sort of forecasts can be used for resource and operations and maintenance planning. When we're getting to the day ahead, we can produce much more accurate forecasts, and these can be used for the unit commitment process; or if you're in a market environment, generators can bid directly into the market based on these forecasts.

When we're going down the timescale, it's also a good idea to have intra-day adjustments, and this is because we'll have better information closer to real time, and so we can adjust those forecasts to reflect that better information. The actions that can be taken in this system are, for example, reconfiguring peak plant schedule, or perhaps obtaining more reserves if we know that our forecast was going to be fairly far off.

As we go down in the hour ahead, this is often the sort of timescale that's used in economic dispatch, and so we should have a smaller uncertainty ban at this time, and we can used revised forecast to optimize the individual units which are generating and change their set points at the economic dispatch process. But we can also utilize it to understand if we need to trade with neighbors perhaps because we have a surplus of energy or if we need to buy energy from a different region in order to make up for a forecast and short fall. It can also be used to optimize transmission or relieve transmission congestion in those time scales.

At the intra-hour, these values are very important for control room operation, because operators can manually adjust units to reflect the changing information that they're seeing as well as assessing the sufficiency of the current reserve level. Here in the intra-hour ramping also becomes very important. Is there expected to a large wind or solar power ramp during this time period, and do we have sufficient capacity available and sufficient ramping capacity—whether that's up or down—in order to efficiently handle that ramp?

So here's just one. I mention ramps a bit, and so here's just one example about ramping. A ramp is a significant and sustained change in output from wind or solar power due to the resource. So the wind speed is changing or the irradiance is changing. Now, solar does follow a diurnal cycle, and so we want to consider ramp in solar the rising or setting of the sun. That's something that we know is going to happen is advance.

However, there is no standard definition of what a ramp is in renewable energy, just as there's no standard definition of a load ramp; and the reason for this is that a ramp—what would be a significant ramp—is very different in different systems. Let's say you're running a fairly small island system of 100 megawatts and you have 50 megawatts if wind power. If you had expected—if you had 50 megawatts of wind power and you went down to 25 megawatts, that would be very significant for your system. And so you would definitely want to be able to capture that in your ramp definition.

Now if you're ERCOT, which as we saw has more than 12 gigawatts of wind power, 25 megawatts in changes just isn't a significant difference for you. And so this is why it's really important that not only the magnitude, but timing and time scale of the ramp, the definitions that are utilized for those match your current system and the flexibility of your system.

In the figure, we show one day, a one-day example of a wind power forecast and actual. So the gray line there is the actual generation, and the first note we have there is during a time which in the blue line we had a ramp expected. So we were going from near 0 output up to about 75 megawatts. So we'll consider this a significant ramp in the forecast. However, this didn't actually occur in the actual generation. Well it's a false positive forecast, because the ramp was forecast but no ramp actually occurred.

In the later timeframe, about 2100 hours there, we have a ramp in the forecast. So we're going over the course of an hour from roughly 75 megawatts up to almost 200 megawatts. In the actual, we did see a significant increase in the generation and so we call this a true positive—a ramp was forecasted and a ramp did occur.

Now we also need to understand the dead band around our forecast. So we did not get the timing nor the magnitude of that is exactly right. And so depending on how large that is, that is important to define as one of your metrics along with what constitutes a ramps, what constitutes an accurate ramp forecast.

Okay. So let's talk about how forecasting can lead to better economic and operational system performance. So better forecasts improve the unit commitment and dispatch efficiency, so this means we can utilize the least expensive units more often, and our marginal units will have less mileage.

That means they'll be ramped up and ramped down less often, and they'll start up and shut down less often.

In time, we'll also develop confidence in the forecast, and this could allow us to reduce our reserve level, whether that's regulation reserve—that's that short sort of time period from seconds to minutes—or if it's our flexible or load-following type reserve in the minutes/hour time period. And as we showed before in one of the examples, it also leads to decrease curtailment of renewable energy generation, which can be very economically important, especially for vertically integrated utility, integrating renewables under Power Purchase Agreement.

One case study here that we'll talk about is the Xcel Energy Case Study. So Xcel Energy is a vertically integrated utility in the United States, and here I'm going to focus on what's known as their PSCo, Public Service of Colorado territory. So they're our local utility and so we do quite a bit of work with them. They are one of the leading utility wind providers in the U.S. and they also have quite a bit of solar power.

In 2014 they had about 15 percent of their annual average energy supplied by wind, but it's important to note that even at 15 percent of average annual energy, they had times where they had instantaneous penetration of up to nearly 70 percent wind energy on their system. This all comes from having almost 6,000 megawatts of wind capacity installed.

Xcel was one of the early adopters of wind in the United States, and so they had the foresight to look into accurate wind power forecasting fairly early. In doing this, they partnered with two different national laboratories based here in Colorado to develop a state-of-the-art forecasting model which has been maintained by a third party now. Over the time period 2009 to 2014, they were able to reduce their average forecast here from 16.8 percent down to 10.10 percent. Just these improvements saved rate payers in Colorado almost $50 million over that time period. Now that doesn't even begin to speak to the savings that occurred just because of the utilization of a forecast. I think one of the most important points I want people to take away from this is improving forecast has quite a bit of economic benefit, but the most benefit you'll see is going from not utilizing a forecast to utilizing any decent forecast.

Let's talk some more about some of those factors that influence the forecasting benefit. Now this is going to be different from system to system as everyone has different operational practices and market design. So as one example, unit commitment in some regions takes place essentially 24 hours ahead, and in some others that's about 48 hours ahead essentially. Because of this, you have much better forecasts at 24 hours than you do at 48 hours, and so that can influence the benefits that you see from the forecasting.

Other institutional drivers include the regulations and incentives or penalties associated with accurate forecasts, the timescales that are utilized and how accurate and reliable they are. Some regions are easier to forecast than others. Just as an example, we live here in Colorado which is a very mountainous state, which means very rough terrain; and complex terrain such as that is very difficult to resolve within the model they do the forecasting. And so you would expect slightly worse performance in a terrain environment such as Colorado than you would in an area that, say, has a reliable coastal sea breeze.

Another important factor is operator competence in the forecast, because you can have the best forecast in the world; but if your operators don't actually utilize it or they just turn off the forecast monitor it doesn't do you any good. And then the different ways in which the forecast is used is also a very important institutional driver.

Some of the physical drivers are that terrain surrounding the environment. This really influences the variability of your resource. Your network size is also important, and your generation resource mix is also important. Let's just say you were trying to integrate wind into a system that was entirely nuclear energy. Well nuclear generators typically do not like to startup and shutdown that often. They don't like to ramp-up and ramp-down. And so this would be much more difficult, and thus it would make your forecasting much more valuable in that case.

And then as you expect, the penetration level of wind and solar power really drive the value. If you're at one percent renewable energy penetration the value is not as large, and that's because the system's designed to handle that level of uncertainty just from demand. But as you go to higher and higher penetration level you'll start to see more value associated with it, because more of your generation comes from those resources. So it's important to remember here that forecasting is not a silver bullet; it reduces the uncertainty associated with your wind and solar power; however, it must be integrated with other flexibility solutions in order to have the full benefit.

Okay. So here on Slide 18 I'd just like talk a little bit more about how more frequent decisions can reduce uncertainty. In this figure, what you see is the net load—so that's essentially the load minus the wind and solar power—and we have different time periods there. Now what you can see in the green and the blue is sort of the uncertainty associated in those timeframes. The reason for this is fairly simple. If I look out the window now—and here in Colorado it's nice and sunny—

I can imagine that I have great solar power resource at this time of day, and I can feel reasonably confident that that's going to persist for the next five or ten minutes, because I don't see a cloud in the sky. However, I don't know what it's going to be like at noon; that could change dramatically at that time point. And so as we move out in time, our uncertainty sort of band has to increase just because there are a lot of different meteorological phenomena that could drive much different conditions than what I'm experiencing right now. And so more frequent updating of your forecast eventually leads to better forecasts and reduces the uncertainty associated with those forecasts. However, the key caveat here is you actually have to tie those updates to an action in the system in order to see the benefit.

So in North America, we have quite a few different utilities, whether there's vertically integrated utilities or market operators who have been utilizing wind power forecasting for quite some time now. This chart just provides an example of some of the ways in which those forecasts are utilized in those utilities. You can see that almost all of the utilities use their forecast for intra-day unit commitment, and quite a few use it for forward unit commitment.

Many also use it to establish the amount of reserves that they need to hold, or to manage their gas storage or their hydro units. Some also utilize it for transmission congestion or for outage planning as well. So this just gives another example of some of the ways that utilities in the U.S. have historically utilized their forecast.

Okay. So we talked about the power system and how forecasts are utilized in the power system. But how are forecasts actually produced? The basis for most forecasting is numerical weather prediction model. So these are models that are run on very large supercomputers by national meteorological agencies. In the U.S., this NOAA, the National Oceanic and Atmospheric Administration. In Germany, it's the Deutscher Wetterdienst. In the U.K. it's the U.K. Meteorological Office, just as some examples.

Now these are the foundational forecasts that drive all of the weather prediction that you see. When you turn on the TV in the morning and your local weatherman says, "This is the weather conditions that you can expect today," it's likely that that information is on these foundational models.

The issue with these is that they're rather coarse, because they do have to be run on supercomputer. They can't have the geographic and temporal resolution that would be most useful for wind and solar plant data. This is why you need to utilize measured wind data, or measured solar or radiance data, and combine those together with these models either through downscaling through other numerical weather prediction models, or through statistical techniques to then get estimated wind speed or radiance forecasts, which are the basis for the power of prediction.

Let's dig a little bit more deeply here into those. I mention the numerical weather prediction model; this would be classified as a physical or dynamical method. So here at a large scale the weather data is taken in. This may be from something—like something at a 40 kilometer resolution. So the temperature, pressure, surface roughness are all put into the numerical weather prediction models to create terrain-specific weather conditions. These can be downscaled further, but it is very computationally expensive to do so.

Another method is to utilize statistical methods. Typically, physical methods work well at, say, the four to six-hour timeframe and above, and that's because it takes them quite a while to run. So to update, it takes maybe four hours. So anything under four hours you're utilizing old information. Statistical methods tend to outperform physical method when we're talking about those shorter time scales—so four hours and under, or even an hour an under.

Now I'd like to talk to you about a couple of baseline forecasts. So any forecast is better than no forecast as long as it is as good as these baseline forecasts. So some of the baselines that we use in short-term forecasting—so that's up to maybe four hours ahead—we utilize persistence forecasting. Persistence forecasting is as simple as you can get. You can see an example of it there for one day in the right-hand figure. What you do here is you take the actual generation at the current time point—and let's say we're producing our head forecast—I'm going to assume that in one hour whatever wind generation I have now will be the wind generation I have then.

Now this is a very simple, very naïve approach; however, for short-term time periods for the reasons discussed earlier, it works fairly well. Statistical techniques can improve on this; but if your forecast cannot be persistent, then there's no reason to use it in those short timeframes.

In the longer timeframes—day ahead, week ahead—we often use a similar sort of persistence type method which is known as a climatological forecast. So this is essentially the average over some period of time that we have from historical data. What this is is in the month of April, I find that I have more—in Colorado that tends to be a windy month, for example, so I would have a higher level of generation in April than I would in, perhaps, July. And so utilizing the monthly average there would be a climatological forecast.

Once you create these forecasts—and there are a number of different methods than can be utilized to create these—one way to improve your accuracy is to utilize ensemble forecasting. This aggregates the results from multiple different forecasts; and the reason why this improves your results is because at any one time, one of my, let's say, ten forecasts may do very well or may do very poorly. But by averaging all of those, we can tend to reduce the error associated with the overall forecast.

Okay. So just a little bit more here on the time scales and methods that are utilized. In the intra-hour timeframe—so that's 5-60 minutes head—typically we use either statistical techniques or persistence type methods. In the short-term 1-6 hours ahead, here there's a blend of statistical techniques and numerical weather prediction model. In the day ahead or medium-term timeframe, it's often a numerical weather prediction model that has model output statistic corrections for systematic biases. And in the long-term forecast we use those climatological forecasts or the numerical weather prediction model. That's just your sort of normal deterministic forecast.

For actual decision support, which is what ramp forecasting fall under, you want to continuously update that using both numerical weather prediction models and statistical techniques. And if we're talking about load forecasting, load forecasting can be done in those same timeframes. Generally speaking, this is done statistically; however, I don't want to downplay the role that numerical weather prediction models even play here.

So just as one example, I've worked with ISO in New England on load forecasting before and they have to utilize for the day ahead a temperature forecast, because in the United States temperature really is the primary driver of load. And so if they had a perfect temperature forecast in the day ahead, they would have only about a one percent error in their load forecast, utilizing the uncertainty and temperature that increases their day ahead load forecasting error to two to three percent. So here, even those numerical weather prediction models provide you with that temperature forecast in the day ahead play an important role.

What else impacts the forecast quality? Well as I mentioned many times here, the quality of your data, your meteorological observation is very important. The more frequent, the more high-quality, and the more dense your observations are, the better your forecast will be. This is just because you're utilizing more information to create them.

There are a number of different things that you can tweak within your numerical weather prediction models whether that's how you assimilate data, what parameterization schemes you choose, or even the resolution of those models. The resolution is the ability to represent terrain features that could impact the resource—and I'll talk more about that here in a second.

But then another key point here is operational information about your generators is very important. If you have a 300 megawatt wind plant, but half of your capacity is offline for maintenance and you don't know that going into your forecast, you're obviously going to have a very large error in your forecast because you cannot physically produce that extra 150 megawatts of power no matter what the conditions are.

And then there's also impacts from your power conversion algorithms. The chief meteorological variable for solar power is solar irradiance, and the chief meteorological variable for wind power is wind speed. However, there are also secondary and tertiary effects from humidity, from temperature, from pressure; and these do impact your power conversion, and so that's an important point that needs to be considered as well.

Now let's move on to an example of how the rain and spatial resolution can impact your forecast. So if we look at this picture, this is sort of typical terrain for here in Colorado. You can see on the left-hand side you have sort of altitude. You have higher mountains in some places than you have in others.

Now let's imagine that you have a wind plant right at the valley there on that left-hand side. The wind speed in this area is going to be significantly influenced by that large peak you see there on the left-hand side. It's only a question of whether you're modeling and capturing that or not.

So if we use a 27 kilometer model, this entire sort of mountain range gets represented as just one sort of smooth curve there. So it looks to the model like it's rolling terrain. Now 27 kilometers is actually much better than many of the foundational numerical weather prediction models. So if you're trying to forecast for that notch there, that little valley, these models probably won't do an ideal job.

If we go down to a 9 kilometer resolution you start to get a little bit more of a representation of the peaks and valleys associated here. But as you can see, it's still not exact. That little notch that we're trying to forecast for, still the terrain in our model is a couple hundred meters higher than it would be in real life. Finally, if we go down to 3 kilometer resolution we can do a fairly accurate job representing the ups and downs of this roll of this mountainous terrain; and because of this, we can expect to have better forecasting accuracy for this region.

Now the obvious question here is, "Okay, well why doesn't everyone run it at one meter resolution?" Well as I mentioned, the numerical weather prediction models that are utilized here are very computationally expensive. Here in Colorado, NOAA and NCAR, which is the National Center for Atmospheric Research, have two or three of the largest supercomputers—of the top 100 supercomputers in the world—running these things all the time, going down to higher spatial resolution or higher temporal resolution requires more computing power, and so it's much more expensive. So there's obviously a tradeoff that needs to be made here in terms of resolution versus computational capability.

Okay. So let's take a step back and talk a little bit here about some of the considerations that are required for implementing forecasting systems and practice. There are different sort of paradigms that can be utilized here. A centralized forecasting paradigm—here I'm taking the example of a system operator performing this; or a decentralized forecasting paradigm here represented by each generator individually forecasting.

In a centralized forecasting paradigm, this enables the system operator to make use of good forecasting information, and unit commitment and economic dispatch. But it does require mechanisms to obtain the necessary data from the generators and encourage quality data being reported. It allows for greater consistency between individual wind or solar plants, and does reduce the uncertainty at the system level.

From a decentralized perspective, this is used by generators or off-takers who are making offers, and it can help a project operator or an owner/operator to optimize their operations and maintenance schedule. It also helps inform operators of potential transmission congestions, but the limited scope can decrease the utility to the system operator. This is really why a best practice is to combine these two approaches. The system operator has a centralized forecasting system that utilizes information provided from decentralized forecast by the generators in order to have both a backup system in place, but also more accurate forecasting because it has more information.

A very related topic here is who actually accrues the benefits from improved forecasting, and conversely, who bears the risk of a poor forecast? Now this is very systems specific. Even within the United States there are systems which are run by an independent system operator where wind and solar power participates in the market, and thus the onus is often placed on the plant owners or plant operators. There are other systems that are vertically integrated utilities in which the powers purchased under a PPA and the onus in that PPA is placed on the utility itself and not on the owner/operator. Ultimately, the idea here is that the consumers and the rate payers benefit for more efficient system operations from better forecasting. However, to what degree that occurs and who in-between benefits or bears the risk is something that is very, very systems specific.

A best practice here is to have those who bear the financial or risk, those are the ones who have the strongest interest in improving a wind or solar forecast. And so placing that onus upon the right people within the system is very important for efficient system operation.

But how do the forecasts actually impact decisions? Well it depends on how they're used: "Will they be used by transmission and distribution system operation?" Maybe that's in the unit commitment process, maybe that's in market operations or scheduling.

What time intervals are needed? "Do you need hourly forecasts with 15-minute resolution? Do you need day ahead forecast with hourly resolution? Do you need hour ahead forecast? Do you need 4-hour ahead forecasts?"

There are a number of different options here, and the key is tying this to system operations: "Do you want point forecasts which are deterministic estimates?" This is what we've spoken about so far. This is where you say, "In the next hour I'm going to have 50 megawatts in generation."

A probabilistic forecast estimates the uncertainty there. It says, "In the next hour I have a 50 percent likelihood of having 50 megawatts in generation, but I also have a 10 percent likelihood of have 20 megawatts, and I have a 40 percent likelihood of having 100 megawatts." So there's quite a bit of uncertainty provided in that forecast, and that tells you something that can be utilized to more efficiently operate your system.

A key component here about data is that monitoring and verification is really an essential component of forecasting. The old adage "garbage in, garbage out" is very important when talking about forecasting data. So the better information you have about your generators and the meteorological and power conditions at each of those, the better your forecasts are going to be.

And there are a number of reasons why we want to look at this for forecasting verification. One, it allows us to monitor the forecasting quality. So are the forecasts improving over time? Is the performance degrading? Those are very important considerations.

If we know this, it allows us to improve the forecast quality. It also allows us to compare different forecasting systems. So maybe we have multiple vendors who are providing forecasts, or maybe we're using a vendor in an in-house performing. If we can compare those two at a detailed resolution, we can figure out where each one works best and where each one doesn't work as well, and then we can improve our overall forecast.

This is also important for financial verification, especially if you're, say, a vertically integrated utility purchasing power under a PPA. One reason for this is perhaps you have to have curtailment at times. Generally speaking, in PPAs you have to compensate the generators for how much power they would have produced during that time of curtailment. Now if you find out—and I've seen this happen—every generator saying, "Every time they were curtailed they would have been producing maximum generation," and your forecasts say, "No, they'd only be producing half of their generation," that's an important financial difference that you need to be able to verify. And I guess that's an issue for your comptroller more than for the operators.

We've talked a lot here about forecasting accuracy, but how do you actually measure that? Well there are a number of different metrics that we utilize for measuring forecast errors. One of them is the mean bias error. This is most important for trying to improve your forecast, because if you have a mean bias error, this means that you're systematically under- or over-forecasting. This is something that can be fairly easily corrected, and you can improve your overall forecast.

Mean absolute error takes the average of the individual forecast—and so it takes the absolute value and then averages those. What this does is it measures the average accuracy of your forecast, but it doesn't allow positive and negative forecasts to cancel each other out.

Root mean square error does a similar thing by squaring your error return. This has the additional affect that it gives a relatively higher weight to large errors as opposed to small errors. An important facet to keep in mind here is that all the metrics are wrong. But some provide you with useful information. And so this is why it's useful to really dig into the data and utilize multiple metrics. Just because your mean bias error is zero doesn't mean you have a good forecast. Maybe you're alternating between having a positive and negative 100 megawatt error at every timeframe. That still gives you a zero mean bias error. But no one would say that that was a good forecast. The other forecast metrics can provide additional information that allows you, even when you break it down, to produce a better forecast in the long run.

An important consideration here is also that these error metrics are not static. If you're producing more generation in the summer than in the winter, it's likely that your error rates are higher at that time, just because there's more possibility to be wrong. And then extent of spatial or geographic aggregation can greatly influence these as well. It is much harder to get a low error metric on a single 10 megawatt wind plant than it is if you're balance over, say, the entire ERCOT system where you have multiple gigawatts of power output from wind. This is simply because you get a lot of spatial and temporal averaging that reduces the overall metrics.

Okay. So this is another way to look at your errors. So what we have here is a histogram essentially of your forecasting errors. So on the left-hand side that's your frequency, so that's how many times an error that falls in that then occurs. And on the X-axis you have your forecast error.

Now the system was designed to handle uncertainty, so errors that are centered around zero in that band that I say is covered by regulation, are relatively unimportant. Why is that? Well if you're running a five gigawatt system, an error of one megawatt is inconsequential. And if that happens a lot that's actually a good thing; your system's designed to handle that.

What really hurts you is those larger errors, the tails of the distribution where you either have to start up a new unit to account for it, which is very expensive, or you have to curtail wind or solar power which can also be very expensive. So focusing your efforts on improving those tails of distribution is a very good idea, because these are the ones that cost you the most money; and having metrics that reflect this is another important consideration.

We've talked quite a bit here about the data collection that's necessary here, and the most important thing to keep in mind here is it's much easier if you have these agreements in place to begin with to enforce them. Once a wind plant or a solar plant is built, it's much harder to get accurate data if it wasn't required at the beginning.

Now these can be accomplished in a number of different ways. It could be a policy mandate or it could be an interconnection requirement. It could be a part of your Power Purchase Agreement. You can set up penalties and rewards for establishing good data collection method., you could also have partnerships with meteorological agencies, and you can even buy this information from third-party vendors.

Let me just talk quickly here about some of the differences between these. In terms of policy mandates—and you actually have something known as the Federal Energy Regulatory Commission. So this is the regulation authority for the power system, and in their Order 764 one of the aspects was that it required generators above a certain size, 20 megawatts, to provide meteorological data that could be useful for the transmission system operator. That's one option.

Another option in a vertically integrated utility is including this information in your Power Purchase Agreement. If you click on the link there on the right-hand side this takes you to the standard wind energy purchase agreement for Xcel Energy. Within that you can find the data requirements that they make their owner/operators provide to them.

What is some of that data? Well there are two different types; there's static data and there's dynamic data. If you have a wind plant, you need to know where it is in order to really produce a good forecast; I think that much is obvious. You also need to know how much capacity is installed there, and it'd be very useful to have some historical data so that you can train your forecasting algorithms on it.

This becomes a lot more difficult when we're talking about, say, distributed solar power. You may not know the location of every two kilowatt panel within your system; that's understandable, and that's a totally different problem. But it's something that we need to consider more when we're looking at these types of forecasts.

In terms of dynamic information, you always want to have the real time generation: "So how much power is that plant producing right now?" The availability data—so that essentially is my wind plant—if it's rated at 300 megawatts, do I have 300 megawatts of capacity online now, or is half of it out for maintenance? That's related to the park potential, which sort of combines forecast with the availability data; and then all the meteorological data such as wind speed, irradiance, temperature, pressure. Those, to a certain extent, are all very useful information that could be required in one of these PPAs.

So what are some of the options for procuring forecasts? Here, I'm just going to present two different options; and I'll preface this by saying there is no correct choice here, there are a lot of different options here, and there is room for combinations of both. One option is setting up an in-house forecasting team. So this is having staff meteorologists that can develop the power prediction model and gather all that meteorological and plant data to produce a forecast in-house. So this can allow for flexibility and for creating custom approaches that reflect your specific concern.

There is the problem that this does have significant cost. It will require a number of meteorologists on staff who are very well-trained in order to do this, and this will probably take quite a bit of computing power as well. You also have to have the expertise to maintain these forecasts. So, that's sort of one avenue that one can go.

The other avenue is there are companies—here we call them third-party vendors—who can provide this information. So they use their own proprietary models in order to estimate generation from wind or solar. These players still require wind and solar plant data in order to produce accurate forecasts. Here even an in-house meteorologist can still play an important role in reviewing the forecast and identifying critical periods.

Now just to give you a ballpark estimate; from a third-party vendor, you could buy a forecast for a single wind or solar project from anywhere from $200.00 to $2,000.00 USD per project per month. Those are just some rough numbers just so that you have an idea of the cost.

One thing I would like to stress here is that if you don't have any forecasting experience whatsoever within your utility or within your system operator, working with vendors can be a quick and inexpensive introductory way to get experience with forecasting. It could be a valuable first step upon which you could build to develop an in-house forecasting system if you are so inclined.

Okay. Let's talk about more of the policy side of things here and some of the actions that could be taken to support forecasting system. One example area is to update the interconnection standards and the Power Purchase Agreement to enable data gathering. So if these are not in place when a plant comes online it's going to be very difficult to get that data later on. So setting these things up before you get a large amount of renewable energy is a very good idea.

You can also start cooperating with your national meteorological institutes to either improve the underlying weather data that you'll need, or just to gain access to that data. You can facilitate the training of operators on the meteorology and how that influences the new generation that's coming online, how they can interpret the forecast and how they can work with vendors if you're choosing to have a vendor provide your forecasting. Now that may sound like an easy solution here, "Okay, we'll just hire a vendor to give us a forecast," but there's a lot of detailed information that goes into that that you need experience with before you can actually seamlessly integrate that into your system operation.

Another way is to support a vendor trial. The added benefit of this is to develop the IT interface that's going to be required between the forecast vendors and the users in the system. This, unfortunately, isn't as simple as just plugging this into your EMS system. What this does is it enables sort of the smooth communication between the forecasting vendor and your system operators. And why that's important is they can provide you with the best forecasts in the world, but if you don't get it until three or four hours later it's not doing you nearly as much good as it could.

Okay. So to sort of summarize here some of the key takeaways that I'd like people to get out of this are that forecasting can play an important role in facilitating integration of wind and solar power into the grid because it reduces uncertainty, and thus it improves the efficiency of power system operations on multiple time scales. However, that better information is only valuable when it leads to better decisions and real actions in the system. So understanding where forecasting can improve decision-making is the first step in considering how to implement forecasting.

Interpreting forecast is also a critical element of effective implementation, so these forecasts will be wrong at times. Understanding why they're wrong, how they're wrong, and what impact it will have on your system is almost as important as having the forecast in the first place. Both centralized and decentralized forecasting paradigms have different sort of values; and in general if you're a system operator, having a centralized forecasting system at least as a backup is very effective in reducing uncertainty at that system level. And a key thing is there is no one-size fits all approach for collecting data and/or procuring and monitoring forecasts. All of these issues are very systems specific, and so tailoring the forecasting program to your unique context and need is critical in order to make them as efficient as possible.

Okay. I'm going to hand things over to Jen now who will talk a little bit about Greening the Grid.

Jennifer Leisch Thank you, Bri. And I want to encourage everyone if you do have any questions—we've had a few come in—but if you have any additional questions, please feel free to type them in the "Questions" pane and we will try to get to them during the question and answer session.

What I do want to point out to everyone is that this webinar was brought to you in partnership with USAID and USAID's Greening the Grid project. And Greening the Grid is a project that provides a set of resources, tools, and best practices in integrating variable renewable energy into the power system. You can find more information on renewable energy forecasting, but also many other topics related to grid integration at greeningthegrid.org, which you can see here.

So in the meantime I want to take a few questions from the audience. And, Bri, maybe you can start with telling us—as we're discussing all of this forecasting and the importance of incorporating forecasting into the way systems are operated, is there a certain penetration of renewable energy on a system where a forecasting program really begins to bring value? Does it have to be over five percent, over ten percent? Is there a magic number?

Bri-Mathias Hodge That's a fantastic question. So my advice here is if you're expecting growing levels of penetration, the sooner you get started with forecasting, the more successful you will be and the more value it will bring. So if you have one percent, two percent penetration rate of renewable energy, depending on your system, this is the uncertainty and variability sort of within the bounds of what we expect for and demand. So you're not going to immediately see huge benefits from that.

But what you are going to get is you're going to get quite a bit of experience with a forecast, because this is not something that just can be switched on overnight; and it's much better to get that experience at one or two percent penetration than when it gets to, let's say, ten percent and you're thinking, "Uh-oh, I have real problems associated with this because we didn't have a forecast beforehand."

And so generally speaking, depending on your system, somewhere in that five to ten percent range by annual energy penetration is when you're really going to start noticing big differences in your system operation. And that's why it's important to get started with forecasting and develop that experience before you hit that point.

Jennifer Leisch So we talked about starting to develop that experience, and one thing that you mentioned was using a forecasting vendor initially, and I think probably a lot of folks on this webinar are maybe at that stage. So are there certain criteria for choosing a forecasting vendor? Or how, if I was a system operator, would I choose one vendor over another?

Bri-Mathias Hodge Sure. So one of the popular ways from utilities—and this definitely isn't the only way—is to set up a vendor trial. Well in this case, you would provide some historical data and you would complete a competition essentially among the forecasters for your system utilizing your data. Now this isn't as easy as it sounds. Even this requires quite a bit of upfront work in terms of setting up the IT infrastructure and really understanding what it is that you're looking for, because the more that you can provide in terms of information about what you need in your system, the better these forecasters can do.

You'll run into the problem here that not every forecasting company participates in trials. So the best thing that you could do is start talking to forecasting vendors, gain a degree of comfort with them, and then it's up to you whether you want to perform a trial or not. If you do perform a trial, I do want to say that this is a long-term plan, because in order to accurately see how the forecasters do, you want a long period of time. You don't want this trial to last a week or a month; you want a minimum of six months, more likely a year so you can see how they perform throughout all of the seasons.

Now there are a number of other issues or considerations associated with trials. "Do you do a blind trial?" That is no one can see who the different forecasters are. This has the advantage that perhaps you'd get more people participating, but it also has certain disadvantages as well.

There are other forecasting issues such as, "Do you provide some funding for each of these trial vendors?" If you do, more are likely to participate. But then that is, of course, an additional cost. And so there's a lot of work that goes into designing a trial, and that's something that I think really it's good to speak with some of the folks from your meteorological institute about when trying to design those.

Jennifer Leisch Thanks, Bri. We talked about how important it is for the vendors or whoever is doing the forecasting to have good information and good data; and it's not just weather data, but that's also data from the plant operators. So can you talk a little bit about what some of those barriers are to collecting that data that's needed a centralized forecasting program?

Bri-Mathias Hodge Sure. So, one of the things that we've historically seen within the wind and solar industries is that plant owner/operators treat their generation and their maintenance schedule as confidential information. They see this as an important competitive advantage. Now, if you're the system operator, that is information that you need in order to operate efficiently. If the plant is built and you do not have the mechanisms in place to receive that information, it's going to be very, very difficult for you to receive that information afterwards. And so that's why setting it up either through policy regulations, through interconnection standards or through your Power Purchase Agreement, it's absolutely critical, because without that information, your forecasting performance is going to suffer greatly.

Jennifer Leisch And maybe one last question. If anyone does have any further questions, please do type them in now. So one last question from the audience: "What are the benefits of centralized—can you explain a little bit more the benefits of centralized forecasting versus decentralized forecasting, or in other words forecasting by the plant operators?" Is there a time when one is better than the other, or one is structured differently than the other that would benefit the system?

Bri-Mathias Hodge Sure. Well let's take the example of participating in a market environment. In this environment you would probably expect to have decentralized forecasts from each of the individual plants. Now one of the reasons why this could be problematic if you don't also have a centralized forecasting system at the system operator level is that those individuals plants are going to bid strategically into the market. So they may intentional bias their forecast to receive greater financial outcome rather than just giving you the best estimate of the amount of power that you could have.

Now, obviously, this would influence your market, your locational marginal prices, and it could put you as a system operator into trouble with, for example, having transmission congestion at certain nodes within the system. And so, that's one of the reasons why it's always good to also have the backup centralized system so that you can utilize that good information that some of the owner/operators are providing. However, you also have another source of information with which to verify and validate that information so that you're not relying solely upon them. And they may be intentionally biasing things, or you know, there are less sinister sort of issues as well. There could be data communications issues. There could just be bad sensors, things of that nature; and if you have that backup centralized system, you're more likely to notice those issues sooner.

Jennifer Leisch Great. Thanks, Bri, and thanks audience for your attention and your great questions. It would be great if everyone can please stay on the webinar, because I'm going to pass it to Stephanie for the wrap-up and the survey. And I want to remind everyone to please check out greeningthegrid.org where you can find this presentation in this webinar as well as a lot of other resources on integrating ERE and sub-power system.

Stephanie Bechler Great. Thank you so much for that, and thank you again, Bri, for that excellent presentation. Before we wrap up here, we have a quick survey that we'd like to send you all. If you could please just click on your screen and answer the questions as they come up.

The first question should be visible on your screen right now: "The webinar content provided me with useful information and insight." Thank you.

Our next question: "The webinar's presenters were effective." Thank you very much.

Our third question: "Overall, the webinar met my expectations." Thank you.

Our fourth question: "Do you anticipate using the information presented in this webinar?" Thank you.

And finally our last question: "Do you anticipate applying the information presented to develop or revise policies or programs in your country of focus?" Great.

Thank you so much for answering our survey. On behalf of the Clean Energy Solutions Center I'd like to extend a thank you to all of our expert panelists and to our audience for their participation. I invite our attendees to check the Solutions Center website if you'd like to view the slides and listen to a recording of today's presentation, as well as any previously held webinars.

Additionally, you'll find information on upcoming webinars and other training events. We are also posting the webinar recordings to the Clean Energy Solutions Center YouTube channel; and please allow about a week for that to be posted. We also invite you to inform your colleagues and those in your network about the Solutions Center resources and services including no-cost policy support.

Have a great rest of your day and we hope to see you again on future Clean Energy Solutions Center events. This concludes our webinar.