ARE YOU SMART ENOUGH FOR THE SMART GRID?

A REPORT ON AN ENERGY EFFICIENCY WORKSHOP

LONG BEACH CALIFORNIA
MARCH 16, 2013

2013 PSMA - EPRI Workshop Committee

Ed Herbert *Chair*	*Dennis Symanski* *EPRI*
Dusty Becker *Emerson Networks*	*Micah Sweeney* *EPRI*
Joe Horzepa *Horizon Consultants Ltd.* *Executive Director PSMA*	*Brian Fortenberry* *EPRI*
Doug McIlvoy *Workshop Facilitator*	*Marek Samotyl* *EPRI*
Laurie House *Formatting Editor*	*Lori Wessely* *Stenographer*

EPRI Purpose

The Electric Power Research Institute (EPRI) is an independent nonprofit center for public interest energy and environmental research. In 1973, during the aftermath of the 1965 Northeast blackout, all sectors of the U.S. electricity industry - public, private, cooperative -voluntarily pooled their funds to launch EPRI. With over 30 years of proven success, EPRI is recognized as a world leader in creating technology and providing solutions to challenging scientific and technological solutions for members. EPRI members represent over 90% of the electricity generated in the United States. International participation represents over 15% of EPRI's total program. EPRI's members and funders represent investor-owned utilities, government utilities, competitive power producers, independent system operators, and government organizations involved in funding public benefit R&D programs. ***EPRI brings together members, participants, the Institute's scientists and engineers, and other leading experts to work collaboratively on solutions to the challenges of electric power.*** These solutions span nearly every area of electricity generation, delivery, and use; including health, safety, and environment.

EPRI Mission

To increase and broaden the value that we provide to our members and other collaborative participants, EPRI's primary focus is on:

Simplifying and streamlining our structure and processes — reducing internal complexity so that we are easier to work with and are able to act more quickly.

Broadening participation in our collaborative activities — reaching out to grow our collective forum so that more traditional utility members can interact with a broader range of international, government, and supplier participants.

Strengthening our technical and scientific portfolio and staff — revitalize EPRI's technical and scientific programs to be even more relevant to the electricity sector's most critical issues, while developing and renewing technical staff talent.

Enhancing technology management and transfer processes — increasing the value members extract from EPRI programs.

For more information on the EPRI, visit our website at:

http://www.epri.com

Purpose of the PSMA

The PSMA is a not-for-profit organization incorporated in the state of California. As stated in the papers of incorporation, the purpose of the Association shall be to enhance the stature and reputation of its members and their products; improve their knowledge of technological and other developments related to power sources; and educate the electronics industry, academia, and government and industry agencies as to the importance of, and relevant applications for, all types of power sources and conversion devices.

PSMA Mission

The PSMA mission is to integrate the resources of the power sources industry to more effectively and profitably serve

For more information on the PSMA,
its members, its committees, recent activities, other publications,
or to join the PSMA, visit our website at:
http://www.psma.com

Table of Contents

Foreword

The PSMA and EPRI have a common interest and history in bringing awareness to the subjects of energy efficiency and power electronics. In 2001, PSMA and EPRI held a joint workshop on the role for switching-mode power supplies (SMPS) in creating a digital society. In 2007, the two organizations hosted a workshop on energy efficiency called "Follow the Power." That successful workshop featured ten experts presenting their knowledge on subjects that ranged from generation, transmission, distribution, and facility distribution to rack level and end-product consumption of electrical energy.

Since that workshop, the electrical grid has come under scrutiny for its resilience, capacity, and security. Many initiatives are attempting to address the "smart grid." What is the smart grid? How does it work? What does it include? How can power electronics play a part? The PSMA and the EPRI felt it was time for a follow-up workshop to explore the topic and co-sponsored the current workshop titled "Are You Smart Enough for the Smart Grid?"

This report summarizes that workshop and the effort of many volunteers and the 12 experts who presented their experience and knowledge on topics ranging from explanation and definition of the grid and the smart grid to drivers for change, management issues, communication, security concerns, technical challenges, and solutions.

The Steering Committee hopes you find valuable information in the presentations and narrative included in this report. The report is one of the benefits of your company's membership in PSMA. Please share this copy with others in your organization. If you or anyone in your organization is interested in additional copies of this report or any other PSMA publications, they can be purchased at psma.com or by calling the PSMA.

Acknowledgement

The “Are You Smart Enough for the Smart Grid?” workshop was the result of many hours of volunteer work by the authors who presented. The success of this workshop is due to their willingness to take time from busy schedules to consider the subject and objectives, create the content, and to present on a long Saturday away from family.

Our sincere thanks to both the EPRI and the PSMA for their support of the financial and technical aspects of the workshop. The coordination of the workshop would not have been possible without the many conference calls, coordination work performed, and services provided by the PSMA office of Joe, Judy, and Lisa Horzepa for workshop services, meeting rooms, and refreshments.

The workshop concept and development is the result of an environment created by the EPRI, the PSMA Energy Efficiency Committee, the Workshop Steering Committee of PSMA and EPRI volunteers.

The workshop Steering Committee was led by Ed Herbert, who not only formulated the concept and outline of the workshop, but led the process of creating the topics and organizing them into a cohesive flow. The members of the committee from EPRI and PSMA worked through many contacts to enlist the volunteer expert speakers and attended to the many details necessary for a successful workshop.

For the first time, the Steering Committee decided to use a paid facilitator to handle the coordination and organization of the workshop presentations and report. Doug McIlvoy, past PSMA chair, filled that role. Along with the facilitator role, the Steering Committee decided to include a stenographic accounting of the workshop and Lori Wessely ably performed the work. The report was formatted by Laurie House, who handled the editing and formatting of the 2011 and 2013 PSMA Technology Roadmap reports.

It is our sincere wish that this report provides you with the same in-depth knowledge the participants in the workshop enjoyed.

Section I
Workshop Overview

Workshop Overview

The Report

This report summarizes the PSMA-EPRI workshop in a new manner intended to provide the power electronics industry, PSMA members, and EPRI Members with the full content of the day-long workshop. The report format includes a transcription of the full day's proceedings interspersed with the presentations and panel discussions. Many reports contain just the slides from multiple presentations. This method of reporting allows the reader to follow along with the presenter, slide by slide, and understand the points being made. In many cases, side bars occur only in the transcription, but are key to understanding. The committee hopes this approach makes the report an easily understood and valued source of information.

The workshop addressed the topic "Are You Smart Enough for the SMART Grid?" and was organized into four sessions that addressed The Grid, The Evolution of the Grid, Grid Stability, and The Future. Four speakers in each session shared their knowledge, focusing on the topic for that session. Each session was followed by a brief panel discussion with the speakers for that session. The workshops ran from 8:00 am until well past 5:00 pm. Due to the intense nature and length of the presentations, the panel discussions could not delve into the many questions that remained. This report, therefore, includes a section of "Remaining Questions" the speakers and the workshop committee put together and provided with answers.

As the committee discovered in the workshop, there is much knowledge that is assumed and therefore missing in the transcriptions and presentations. Appendix A contains a section of Chairman's comments for the workshop, which endeavors to provide some missing information. We hope this section provides background for those readers who want more in-depth understanding. Appendix B is a listing of web links to provide more specific data on Smart Grid knowledge and current projects.

Dusty Becker of Emerson Energy Systems led the workshop, providing the introductions for the speakers and keeping the workshop on track. Dusty emphasized the purpose of the workshop which was:

To primarily educate power converter designers and users about how the "Smart Grid" may affect future power converter specifications and design. The audience, and readers of this report, will gain a thorough understanding of "The Grid," its equipment, its control, and its issues; with emphasis on dynamics, stability, and security.

An objective of the workshop was to create an environment in which grid operators and power converter designers understand each other's systems and designs well enough to discover cross-disciplinary solutions, trades, and optimizations.

Executive Summary

The information shared in the workshop left the participants with the knowledge that the grid as it exists today is aging... generators, transmission lines, substations, and transformers with average in-service time greater than the designed service life. The grid is very conservative, integrating new equipment and new technologies slowly due to risk aversion and liability. The increasing introduction of “smart” technologies opens the grid to cyber attacks that threaten the economy and well-being of all nations.

The power electronics industry plays a major and necessary role in the current grid. The power electronics industry and PSMA members can play an important role in providing equipment to support the “Smart Grid;” but utility operators and equipment providers must listen and respond to the needs of the other.

The Grid

Clark Gellings, fellow of EPRI, led the workshop with a very energetic, enlightening, and informative discussion of the grid. Clark gave the audience a look at the history, the current status and issues, as well as a comprehensive view of “the greatest machine ever made.”

Clark described the grid as “connecting the generation of electrical power to consumers through stuff.” The grid is complex, utilizing technologies that are tried and proven… and aging. Most of the equipment was designed for a 40-year service life…and most has 42 years of service aging.

Electricity is generated at about 13,000 volts, then stepped up through a transformer within a half mile to about 345,000 volts, and delivered through 250,000 miles of high-voltage lines to substations with distribution transformers that step it back down to 120 -480 volts for delivery to consumers. The purpose of the high voltage is to deliver more power at the same current, following Ohms Law. The higher voltages allow smaller tracks of land for transmission lines.

Loads on the grid are variable and utilization of the grid is low, with 20% of the assets used 10% of the time. Reliability is key, and maintaining the equipment is essential to reliability. Many simultaneous paths are needed to provide the almost-five points of redundancy, but that redundancy pays off... in 2001, nationwide outages were 1170 per 1000 customers; Con Edison had 10 outages per 1000 customers. Con Edison in New York serves 3.1 million customers (8.6 million people) with 35,000 miles of overhead line and 91,000 miles of underground line.

The national grid uses eight control regions for the US and Canada, coordinated by the North American Electrical Reliability Corporation (NERC). These regions are coordinated by entities, such as CAISO, that work with grid operators. Dr. Abdul-Rahman from CAISO describes balancing the generation of power to the load in the least costly manner. Complex computers and algorithms in control centers make power available between regions and operators. The control entities must be able to provide reserve capacity that can respond to load increases as well as reduce generation to maintain balance. Solar and wind generation varies dramatically (as much as 1520 megawatts in a four-hour period), forcing the control entities to anticipate and react very rapidly to keep the grid in balance. The control entity must provide for generation of real power as well as the reactive power (VARs) that are necessary to move the power to the end use.

There are many technologies that help balance the grid: Energy storage, HVDC, Flexible AC Transmission (FACTS). All of these are necessary and improve the reliability and efficiency of the grid. Efficiency of the grid is low: in delivering power for lighting and heating, the end user gets 3% of the power generated. More than 210 billion kWh per year are lost in delivery.

The grid today is a bulk system with power generated centrally. The grid of tomorrow is a distributed grid, combining micro-grids of solar and wind generators with storage of electrical vehicles, providing higher reliability through multipoint generation and storage.

Don Von Dollen of EPRI describes several drivers for change to the grid:

- Extreme weather events are on the increase.
- Environmental and energy policy is increasing demands for higher energy efficiency, more renewable energy, and lower greenhouse gas.
- The infrastructure is aging to a critical point.
- Growth in connectivity is dramatic, with each US home having an average of five networked devices. By the end of 2013, mobile connected devices will exceed people on earth.
- Customers expect more information faster than ever.

Electrical service needs to become more resilient, flexible, efficient, and customer focused and achieve these in a cost efficient manner.

Clark Gellings referred to the grid as the "dumb" grid; Dennis Symanski points out while we have elements of a "smart" grid, not all the pieces are in place and there is a long way to go. The "smart grid" needs to be bi-directional: with power and communications moving in both directions. The operators of the grid want to be able to adjust loads such as household appliances, heating, and lighting to reduce demand and balance the system. Consumers want more information to make better decisions on what to allow the utilities to control. Incorporating local storage and / or generation of electricity from electric vehicles can provide higher reliability of the local grid as well as lower consumer costs.

There are many projects in progress for the "smart grid." EPRI has 15 large-scale demonstrations going with 15 utilities. One of the most interesting is at the University of California at Irvine, where a full campus is being converted to a smart grid. The demonstration includes modification of homes and buildings to higher efficiency standards as well as incorporating storage, solar generation, and electric vehicles. This will be a great story to follow.

Many new technologies may or may not be proven. We must try new things, see what happens, determine what it effective, in function and cost, then keep what works and throw away what does not.

Evolving Loads

The grid, as we know it, is characterized by fully controllable central generation. It has mainly variable and unregulated loads with any transient balancing done by kinetic energy storage where no control is needed. Renewable energy is variable and distributed solar and wind generators. These act as "negative sources" in the traditional sense of the grid. The effect of these new sources is less kinetic energy storage and the need for more control.

The increasing number of loads and the speed of control for those loads presents a stability challenge to the grid. Traditional control is in the sub-Hertz range and the loads of the digital world are controlled in the kilo-Hertz range. This presents a stability issue. Dr. Jian Sun presents a unique solution: measure the small-signal impedance at a particular point and, through modeling, identify the grid impedance. He then presents a case study that uses the grid impedance to alter the control system of the load to correct for instability. His concept allows him to measure the small-signal impedance in real time and make the necessary adjustments. The concept allows a load to have autonomous response to changes in the grid. The concept is interesting and further emphasizes the need for communication between power electronics manufacturers and grid operators.

Harley Garrett points out that communication is becoming ever more important and is driven by policy and consumer demand. The world's many networking protocols provide a basis for diverse equipment to communicate. The speed of those communications and their complexity astound us. We can get 2.6 billion transistors in an integrated circuit and we can make systems on a chip with multiples of that. The communication speeds have increased from 45Mb/s for 1G in 1975 to over 1.05Pb/s for 5G in 2013. We have the internet, WANS, LANS, and HANS. We have wired, fiber-optic, and wireless. If you follow the protocols, you can make equipment that connects to and communicates through any of these standards. The Energy Policy Act of 2005 introduced the first two way meter as advanced metering (AMI). Under this policy; if you generate more power than you use, the meter runs backward and you can sell what you don't use back to the grid.

Smart meters that can communicate with the utility are being installed widely. Utilities know the usage in any location and we can tap into that and monitor it. Even with the smart meter revolution, there is still a long way to go with Home Area Networks (HANs) before control of private homes by the utility.

The world has changed. The world is digital and connected and this brings cyber threats. These threats are aimed more and more at Supervisory and Control Data Systems (SCADA) to disrupt our grid infrastructure and damage our economy. Jonathan Pollet provides insight to just how vulnerable our grid is to hacking. We can prevent some hacking; but, in most cases, the real prevention is more being vigilant and reacting rapidly when threats are identified. Large companies with IT specialists still get hacked. There is a network of hackers that make their living by hacking systems and selling information. Hackers can access financial systems, control systems, and sales information that are key to company survival. SCADA systems are less sophisticated than phone or business systems and send data in the clear ASSUMING no one is interested. Jonathan's example is that a dinner reservation made on your phone is more secure than a command to an RTU!

Jonathan recommends creating communication islands where control systems are isolated, using encrypted data, and requiring authentication. Even then, hacking is possible, so stay alert!

Improving Grid Stability

Improving grid stability is a major concern. The grid today is largely passive, but the addition of renewables and dynamic loads makes it susceptible to instabilities. Why do we care? Because instabilities can lead to cascading outages that already cost the economy about $80 billion per year. Preventing just 25% of them would allow large investments in improving grid stability.

Dr. Greg Smedley shows how a large solar generator can change its power delivery dramatically in seconds from 400 megawatts to a few hundred kilowatts. More solar coming online is creating a maintenance problem for switched capacitor banks and load tap changers. Dr. Smedley suggests the solution is to place widely distributed autonomous control devices, capable of providing a fast reflex correction, on the grid to restore stability. Providing large numbers of devices that can control the voltage locally prevents the grids substations switching large storage devices in and out and improves reliability, while lowering maintenance.

Dr. Satish Rajagopalan addresses storage, the types in use, the applications, costs, and research that is needed. There is bulk storage, usually pumped hydro, typically about 100 MWh. As you move to substations, the energy storage level drops to 2 to 5 MWh, usually electric storage or thermal. At the residential level, it is down to 10 kWh. At all levels, most storage is used for less than an hour in every 24-hr period.

Storage at the substation level can be very cost effective in providing for peak shaving of loads rather than upgrading the substation. Storage can also mitigate power variation in solar and wind generation, as well as minimizing cycles for tap-changing switches.

Pumped hydro storage represents 99% of the world wide storage capacity. This is followed by compressed air storage, then by sodium-sulfur batteries. EPRI is focused on electrical storage, where cost must be under $350 per kWh to be implemented.

Storage is a key element in the smart grid and it requires a significant amount of power electronics equipment for control and delivery. We need more research into cost reduction and power density.

Autonomous Response

The topic of Autonomous Response is covered many times, but Ed Herbert gives a unique insight. Ed defines autonomous response as the response of a load or power converter to a change in input voltage, phase angle, or rate of change. Autonomous response does not use a data link and achieves response by modifying the input characteristics of the load. A large benefit of autonomous response is the security offered by not being vulnerable to hacking.

Power electronics have included autonomous response in products for the past thirty years with power factor control circuits as interfaces between the power converter and the utility. Power factor control allows the input current to be shaped and the load can soft start and stop if enough energy storage is available. Ed maintains that autonomous response can be applied to sources as well as load.

The Conservation Voltage Reduction (CVR) technique is made ineffective by constant power loads because of their negative input resistance. Ed presents an approach that allows the load to alter its input impedance and act like a resistor under dynamic conditions. This can make CVR effective for converter loads. Ed suggests autonomous response is not limited to sensing voltage change, but can be implemented to sense other parameters; such as outside temperature, time of day, or grid phase angle.

The Future

Dr. Conor Quinn brings everything into perspective with a discussion of the smart grid and smart “super” grid. He points out that power electronics have come a long way and that efficiency has been improved by removing two thirds of the losses associated with conversion. Digital control of feedback loops and digital communication are everywhere. Interfacing the control loops to the outside world is common and can use many different communication protocols. Power electronics can sense, report on, and control many parameters. We can control HVAC, lighting, charging of EVs, and solar or wind generation in micro-grids. All of this exists today, but communication with the grid does not and much is still unclear. We need to have clarity in standards and protocols.

Dr. Keyue Smedley covered system optimization. She reminds us the grid is a collection of sources, loads, resistors, capacitors, and inductors that all obey the laws of physics. The grid is not able to stay ahead of demand growth and is, therefore, subject to lower operating margin and instability. Losses due to transmission and blackouts cost the US $126 billion per year. Congestion cost, the cost of NOT being able to move power as needed, is estimated at $4.8B per year. The opportunities and challenges are huge.

Communications is growing dramatically with huge amounts of data being transmitted. Dr. Smedley asserts widely distributed, local, autonomous actuators with global set points could be used to create a “power electronics super highway.” She suggests using four-quadrant converters as the actuators. Power electronics can handle that task due to power handling capability and fast response.

Summary

After a full day, Dr. Gellings energetically summarized the workshop. The complexity of the smart grid is of concern, but there are many opportunities for power electronics in tomorrow’s power system... phenomenal opportunities, driven by technology. There are opportunities in campus-level local energy networks, building-level local energy networks, and bulk power systems.

The implications for power electronics are that products need to energy efficient, cyber secure, and communications enabled and interoperable.

Dr. Gellings suggests it is time to establish a PSMA-EPRI joint task force to address smart products for the smart grid. The task force could outline a roadmap, draft guidelines, and discuss what attributes those products should have.

We need some leaders....volunteers can contact the PSMA office.

Welcome and Introduction

DUSTY BECKER

I'd like to welcome everyone to the "Are You Smart Enough for the Smart Grid?" workshop, which will have a series of presentations. We'll have three presentations, no questions; then move to a group of presentations with, if not a common theme, a common thread. When those three presentations are done, we'll have an interactive group discussion where we make comments and ask questions, that sort of thing.

We're scheduled from 8:00 to 5:00. My experience is that we'll probably use most of the time.

I'd like to thank everyone for being here. This is an outgrowth of the "Follow the Power Workshop" we did in 2007. We got together and decided we needed to follow up, but where to go next? After we threw a bunch of ideas around, we came up with the idea of an interactive workshop on the smart grid and thought it might be quite useful.

The agenda is embedded in a couple of places. The first papers, as you can see, are a series of four papers on the grid itself. The next grouping is evolving the grid. Then we have grid stability and then the future, with an adjournment.

Now, I sit on a bunch of different grid committees, but I always feel like a novice because it is a huge, huge subject with a lot of ins and outs. My personal interest tends to be more in microgrid discussions and how they interact with the grid. So, that's where I come from. Plus, as you know; at Emerson, we do a lot of data center work and stuff like that in the power industry. One might consider us to be -- amongst many of the users -- victims of the grid. The grid does stuff to us.

So thinking about it, what the experts say about the smart grid, I've got a bunch of different things here. They say it's self-healing. A user might think, "Will I see fewer outages? Will the outages be shorter? Will I still need a UPS?" What is my response to that? What's a "self-healing" grid mean to me? It allows active customer participation. I've heard that a lot… it's a great catch phrase. The question is: What does that mean? What exactly does customer participation in grid decisions mean?

- Resilience to the human issues, natural disasters, and cyber-attacks. The ironic answer is that a lot of guys are still recovering financially in the "East from the 2003" event.
- It enhances quality and reliability for twenty-first century loads. The answer from the user is, "It's about time! We've been doing advance loads for quite a while now!"
- It'll support all generation and storage options. Generation? Storage? "Help us to find what that means." I sit on the NEMA Storage Committee and we're still trying to figure out what it all means.
- It can enable new product services and markets. When somebody tells me that, I figure I'm going to have a vendor talking to me really quick.
- Asset utilization, operational efficiency... How do I better use my assets? How do I spend less money on my operations? How can I be more efficient?

Often times, we get caught up in the smart grid ideas, as engineers are prone to do, as "It's pretty cool! We can do a lot of cool things!" and "Oh, boy! Let's have some fun with this!" Users are all about: "How can I utilize my assets more efficiently? How can I save money? How can I reduce the risk of serious and financial impacting outages?" That's what they want. So let's not lose sight of exactly what it is; it's the reason the smart grid is of interest to the industry.

Section II
Session 1 — The Grid

What is the Grid?

DUSTY BECKER

The first speaker we have today is Clark Gellings. Clark’s got a tremendous curriculum vitae. He’s been around for a while. He’s a fellow at the Electric Power Research Institute and he’s responsible for technology innovation, scouting, and scenario planning.

Since joining EPRI in 1982, Clark has had a lot of different technical management executive positions. He’s been a VP multiple times. He was Chief Executive Officer of several EPRI subsidiaries and, prior to joining EPRI, he spent 14 years with the Public Service Electric Gas Company.

Clark is also a recipient of *EnergyBiz* magazine’s Knowledge, Innovation, Technology, and Excellence Lifetime Achievement Award. He is a member of the board of the University of Minnesota’s Technology Leadership Institute.

Clark has both an Electrical Engineering degree and a Mechanical Engineering degree, as well as a Masters in Management Science. He’s a registered Professional Engineer, a Lifetime Fellow of the Institute of Electrical and Electronics Engineers, a fellow emeritus at the Illuminating Engineering Society, an Honorary Member of CIGRE, and immediate past-president of the US National Committee.

So please welcome Clark as our first speaker.

CLARK GELLINGS

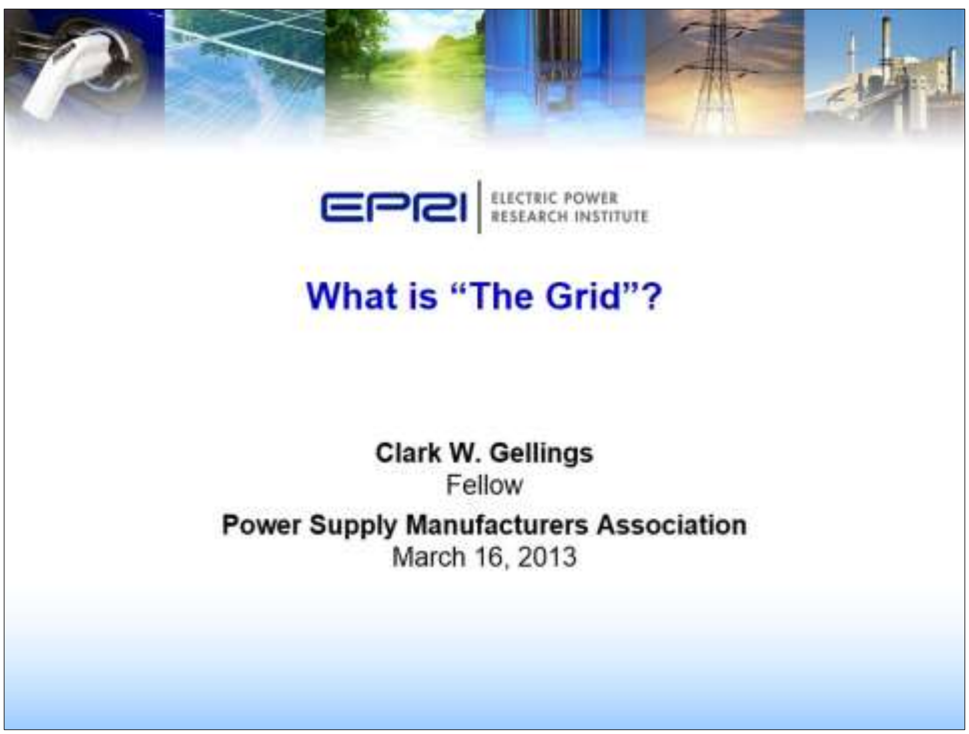

Thank you. I get the easiest -- or maybe the hardest presentation -- to tell you what the grid is. But I’m not going to tell you what the smart grid is; I’m going to tell you what the dumb grid is... pretty much as it is today! There are exceptions to this.

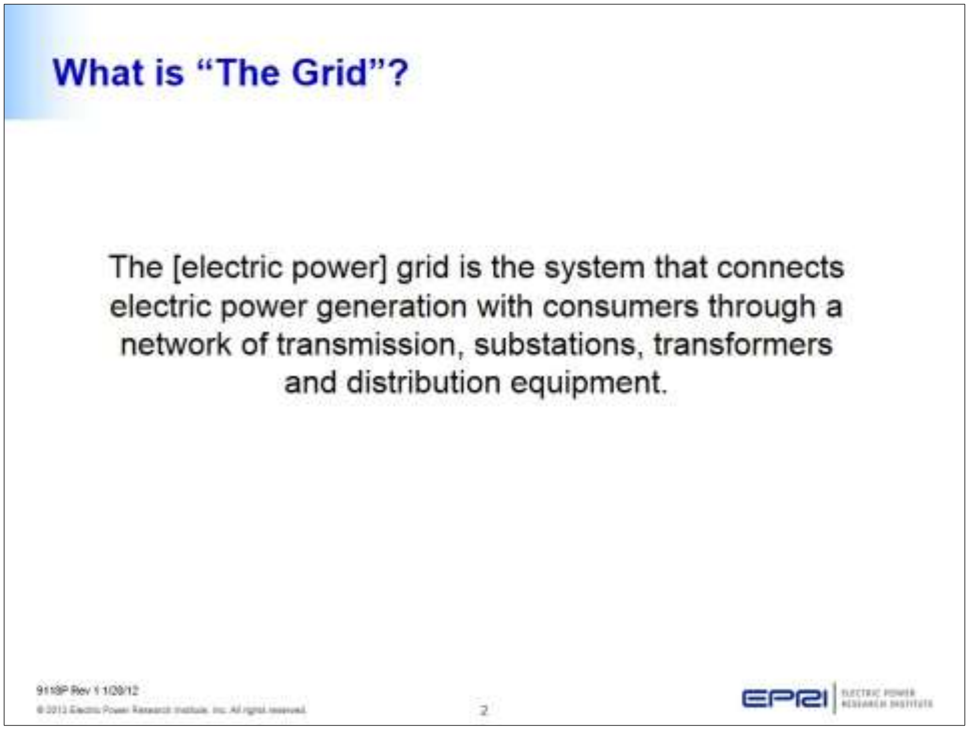

Basically I'll go through this real quick. Some of you may be insulted by how rudimentarily I do this, but I have a feeling that some of you don't really have a good feel for what a power system looks like. By the time I'm done, and my colleagues who follow me, perhaps you'll have a better understanding.

Let me try first a simple definition of the grid, the electric power grid, if you will. There are other power grids in the world. The grid is a system that connects electric power generation with consumers through a network of stuff: stations, substations, transformers, distribution equipment. As this morning continues, you'll hear about some of the advances in the interface between the power system and customers that will be of particular interest to folks.

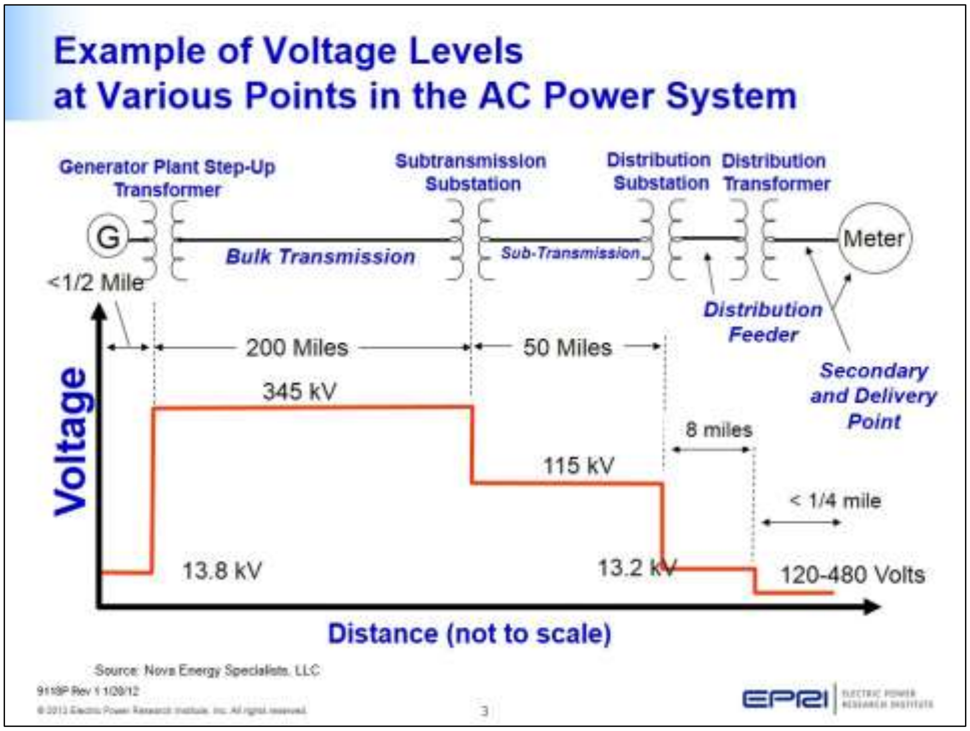

So let me give a quick example of what a power system looks like; in case you haven't thought about the voltages and the like. Typically, what we're trying to do is connect power generation – that's also evolving, as you'll hear -- connect power generation with consumers through a series of wires and cables.

We start by generating, usually at about 13,000 volts. Why so low? Because it's really difficult to create rotating machinery and provide adequate insulation and the like. Otherwise, we'd be happy to generate at higher voltages. Because almost immediately, certainly less than a half a mile, usually less than that, we begin to bump it up to higher voltages; maybe 345,000 – that's popular in the US.

We have with 250,000 miles of high-voltage line that's over, say, 230,000 volts. That's the bulk of the power system, but you'll hear a lot of exceptions. You'll hear 500,000. In fact, I'll show you a 765,000-volt example.

Then we take that into some kind of substation, usually a sub-transmission sub-station. There are a couple of different configurations here. We bring it down to lower voltages, maybe 115,000 volts. There are a couple of other variations, no question, depending on where we are.

We're getting close to the customer now, so we go to distribution equipment. Distribution sub-stations are usually down to 13,000-13,200 volts. There are still a lot of 4,000-volt lines; you'll hear that. Even 2,000. More or less, we're going with 13,200. Sometimes a little higher, maybe as high as 35,000 or so; but just to give you a general feel. Now we're really close to customers and eventually we'll go to a distribution transformer, and we'll bring it down to 120 to 480 volts.

You see I don't say 110; there's no such thing. Where the hell do people get 110? I hear professionals in our industry talk about 110. What is that?

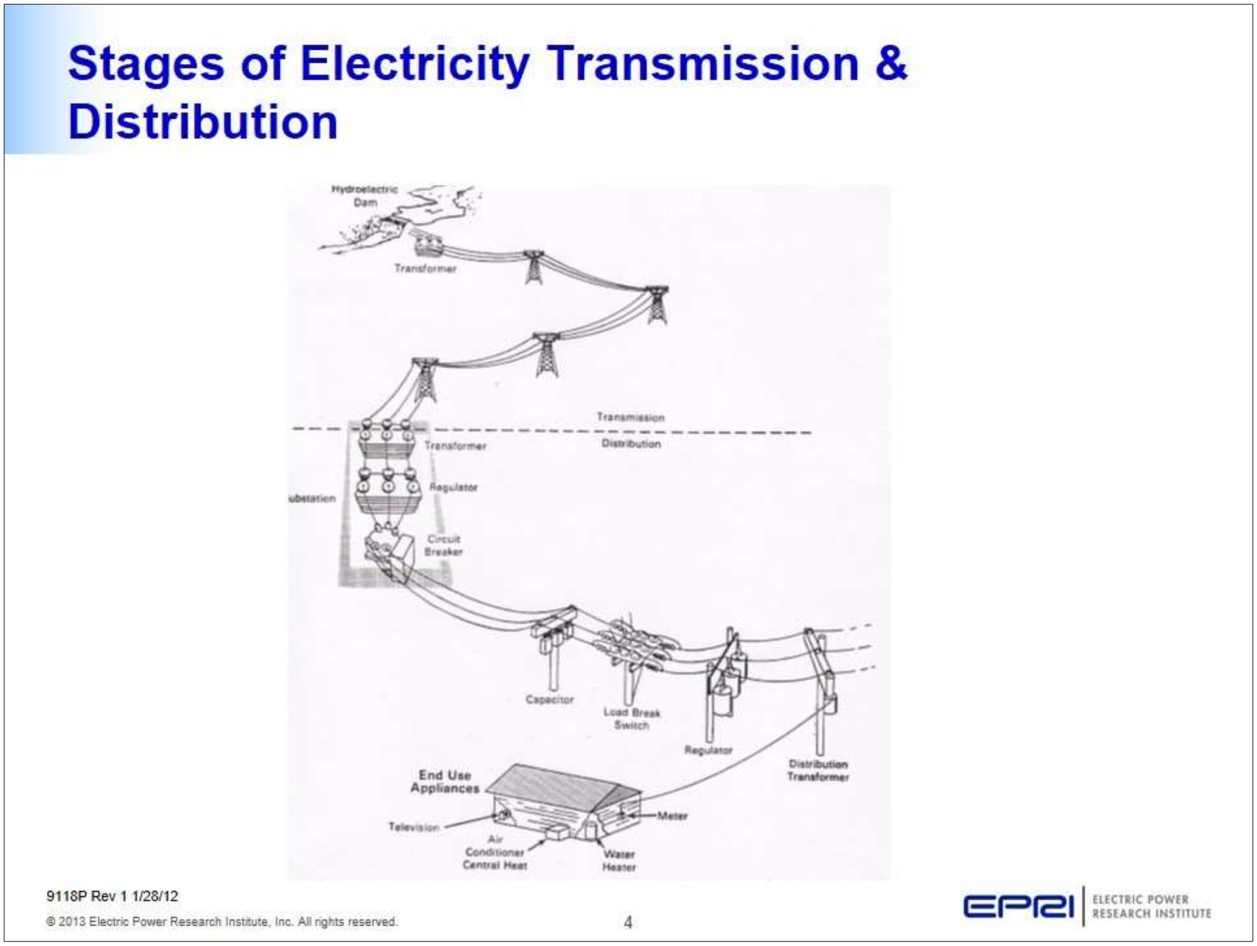

All right. I could have of done it this way… I know this isn't real clear. Basically, I've got a power plant up here, a hydro-electric dam. The transformer is my high-voltage station. Then I go to a sub-station and distribution starts. I've got various equipment, lots of transformers. I'll show you some pictures. I might have other devices. There are certainly circuit breakers along the way, capacitor banks, load switches, regulators, distribution transformers. Finally, I'm connecting to the customer.

Transmission & Distribution System Voltages

(1 kilovolt or KV = 1,000 volts or V)

System Element	Voltage
• Secondary	120 – 480 V (low voltage)
• Distribution	4 KV – 35 KV (medium voltage)
• Sub-transmission	35 KV – 69 KV
• Extra-High Voltage (EHV) Transmission	138 KV – 800 KV (most common levels are 230, 345, 500, 765 KV)
• Ultra-High Voltage (UHV) Transmission	Above 800 KV

9118P Rev 1 1/28/12

5

Some terminology we use, which is probably a little different... Secondary voltages are 120 through 480. We refer to that as “low voltage” — I thought you’d get a kick out of that, having a rough idea of where your business is! Distribution, as I said, is 4 kV to 35 kV. We call that medium voltage. Sub-station is 35 kV to 69 kV and there are variations on that. Extra high voltage is 138 kV to 800 kV. The most common levels are 230, 345, 500, and 765 kV. Go to China and see a million! And you also see a lot more dc. All of this so far is ac.

Ultra-high voltage would be over 800 kV, which we don’t have in the US. Actually, at EPRI, we have a laboratory capable of simulating 1.2, over 1,200,000. We thought someday we might get there, but we’re not apparently headed that way just yet.

Higher Voltage Means More Power

Phase to Phase Voltage (AC Volts RMS)	Current (Amps)	Power Per phase (watts)	Three Phase Power (watts)
208	400	48,036	144,107
480	400	110,851	332,554
4800	400	1,108,513	3,325,538
13200	400	3,048,409	9,145,228
34500	400	7,967,434	23,902,301
46000	400	10,623,245	31,869,735
69000	400	15,934,867	47,804,602
115000	400	26,558,112	79,674,337
230000	400	53,116,225	159,348,674
345000	400	79,674,337	239,023,011
500000	400	115,470,054	346,410,162

Source: Nova Energy Specialists, LLC

9118P Rev 1 1/28/12

6

It’s a simple point: higher voltages mean more power. P = I x V. We know this, right? Naturally, I want to get the voltage up as much as I can because I want to move a lot of power around.

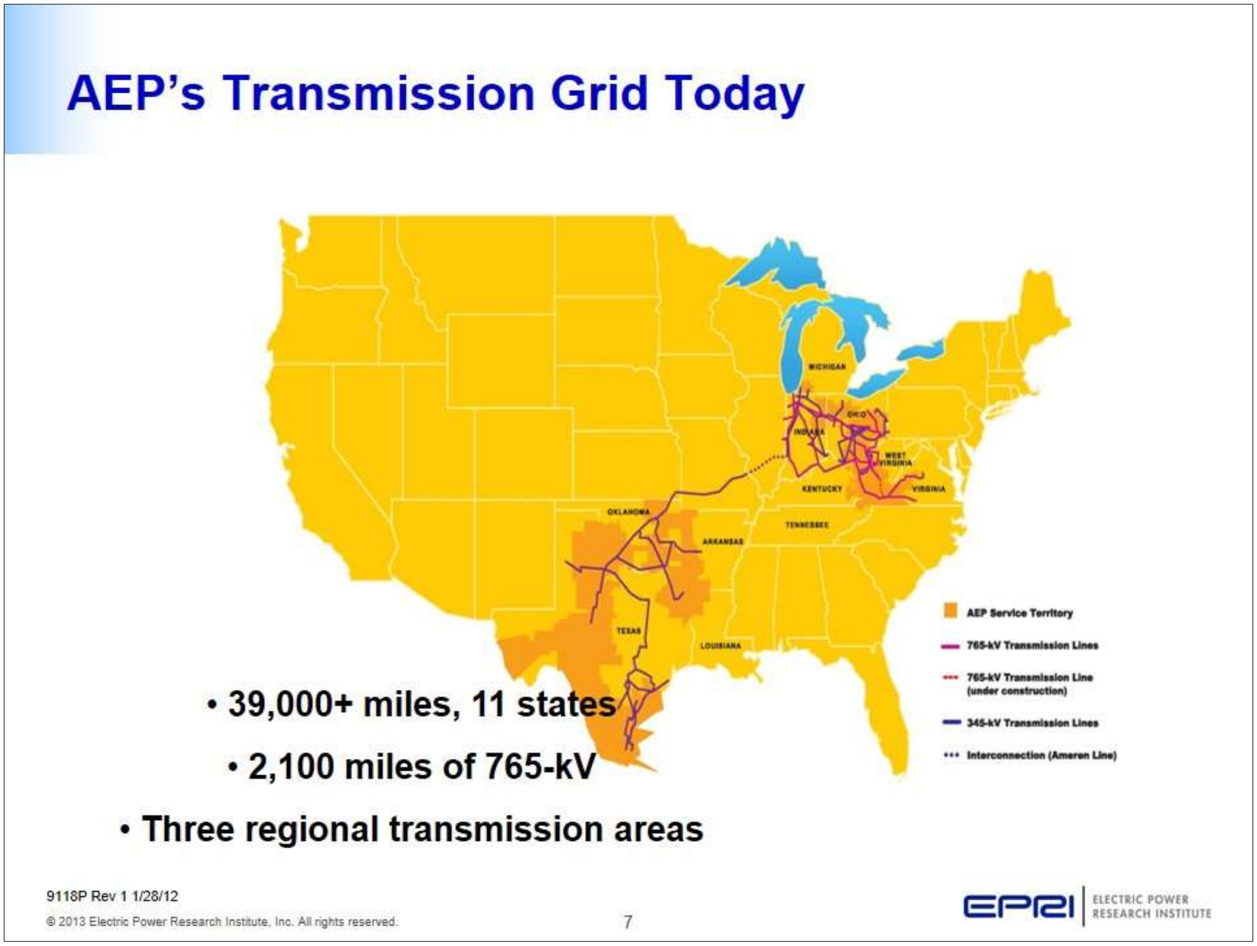

What does the system look like? You'll see better examples in a few minutes. This is the American Electric Power in 11 states. Over 39,000 miles of high voltage; 2,100 miles of that is 765,000 volts. This bulk system is really what supports the overall power needs of the consumers in those communities. They use a lot of this 765 because they think it has the highest capacity and it's the most efficient.

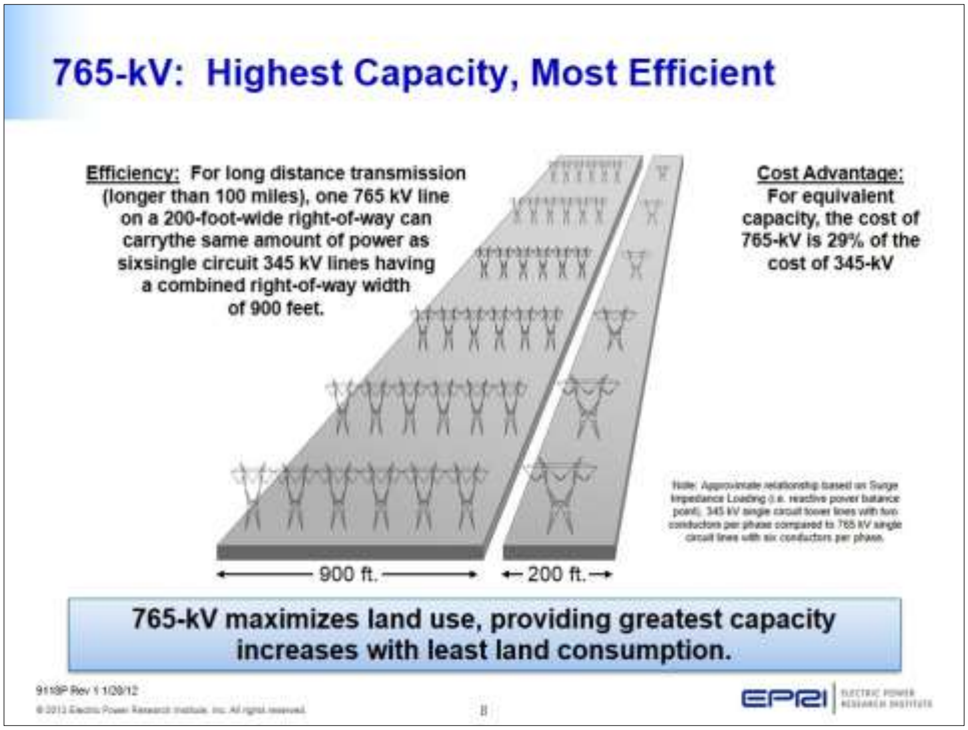

Here's an example of the right of way that's necessary for one 765-kV line and the equivalent that would be necessary with 345. It would take six single-circuit 345s to do what one 765 does. The towers are bigger. Then you get into problems with approvals of communities, reactions, the government, electric and magnetic fields, and so on.

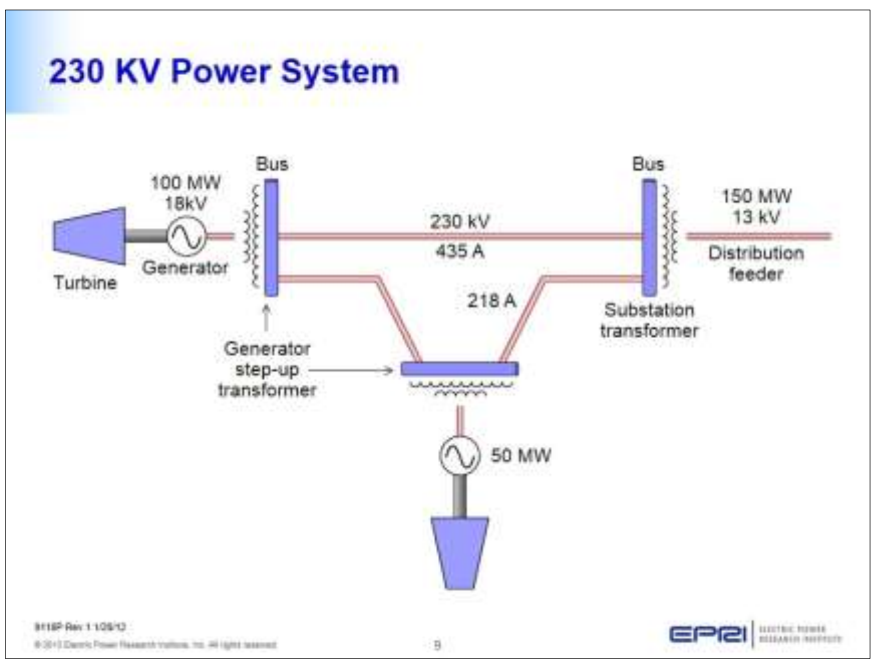

So these systems, typically, if I dive down just a little bit deeper, look like this: transformers, generators — in fact, more than one generator — a very complex system! The National Academy of Engineering referred to this system as the most complex machine man has ever built. They said "man" — pardon me, that *we* ever built.

How did this start? Well, we hear a lot of about our friend Thomas Edison. Yeah, he invented the light bulb and so on. But what he really did with the Pearl Street Operation in New York City was invent a working power system. This wasn't the first power system; there were others, small systems and the like. But he was the first to organize a business around it and to put all the little pieces together that were necessary, including the underground cables, use of transmitted electricity. That was a big deal in the evolution of the power system because there is something referred to as the "overhead" of the system.

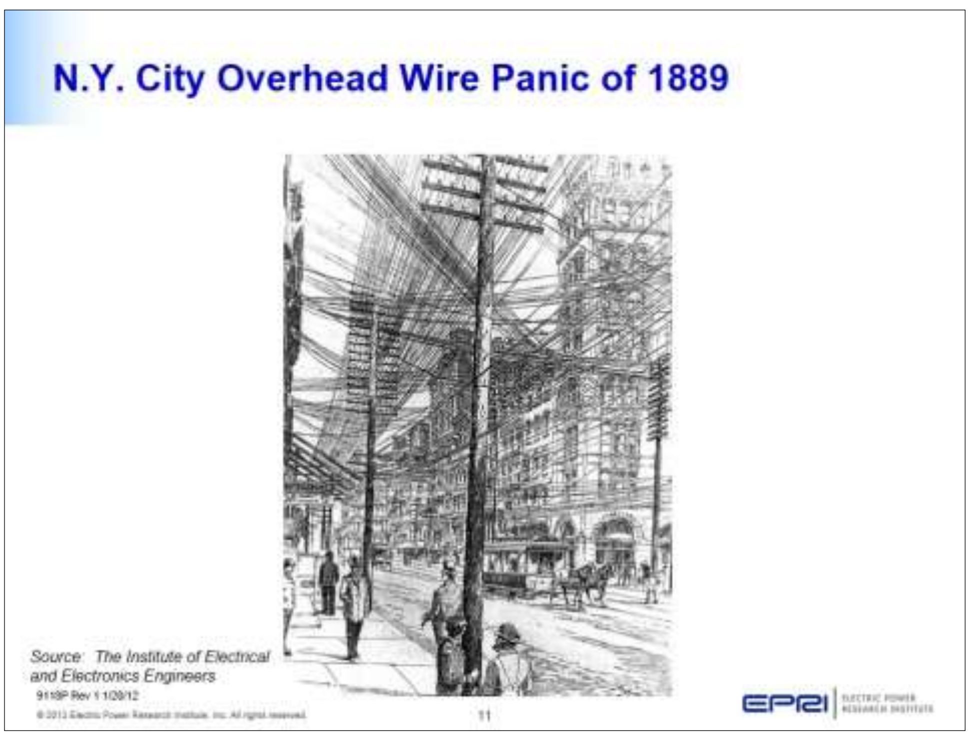

The overhead wire panic in New York City of 1889... Apparently several times within a few-week period, someone got electrocuted at work. A workman got electrocuted and his body hung up in the wires; people didn't like that. They objected to it. So there was a resolution passed by the city of New York that said that eventually this stuff all had to go underground. Well, Edison had a jump on it already.

Transformers are everywhere, all different kinds. You see them around. You probably don’t pay a lot of attention to them. They’re the backbone, right? Most of you know what the Lorentz force is? Here’s the Lorentz force at work. It is because of the Lorentz force that we are able to take voltages as we boost them and make use of higher voltage transmission with smaller conductors.

There are a lot of other components. And these transformers get pretty damn big, expensive, and vulnerable. It’s pretty easy to disrupt something like that with a high-powered rifle. Not that I'm going to give you a cookbook on how to do it... but it’s out there.

They have a couple different configurations. This is a single-phase one. The previous one was three phase. This single-phase generates a step-up transformer. This is a highly vulnerable piece of equipment for a lot of reasons. But, in a power plant, I can go right off the generation bus bar into something like this, so I can get into higher voltages.

Through the generation switch guard, it's called, I've got a whole set of terminology I can the throw at you. I'll try not to. I also have circuit breakers and switches here, right? But this is a gas-insulated circuit breaker. These guys are small because they're using a gas, sulfur hexafluoride, SF_6, as it's called. It's a deadly greenhouse gas. Twenty-thousand times worse than methane. But it's a really damn good insulator. If I didn't use that gas, these things would be four or five times that size... getting into really big equipment.

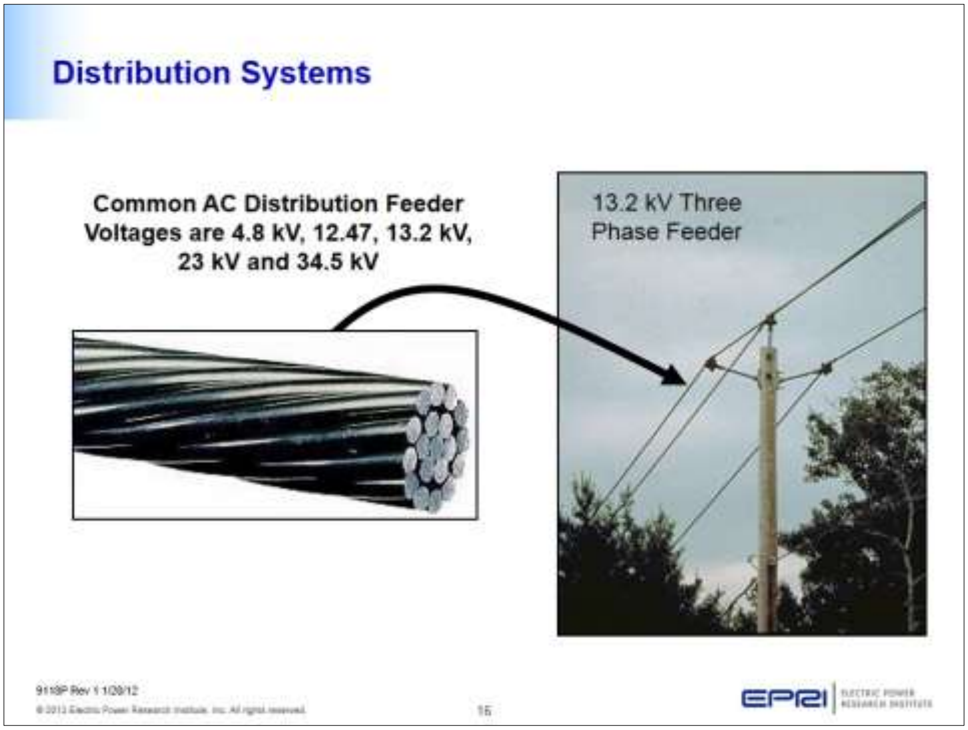

Distribution systems, often overhead, such as a 13,000 kV system, look something like that. We see them all the time. You probably don't pay a lot of attention to them. That is, in fact, the backbone of what delivers power to most customers. I already mentioned what some of the common voltages are.

Of course, we can do this stuff underground too. We do it in duct banks, manholes, and intro points. We direct bury some of it. Each of these methods of distribution system has its own foibles. There are problems with each… especially through things like Hurricane Sandy.

I get asked "Why don't we put it all underground?" Sure. Go ahead. We can do it… although I don't think that solves the problem. As I've often said, "Don't ever ask me to design a power system that works under water. I don't want to do it." Most of the problems we saw in New York and New Jersey weren't just from wind; they were also from water. Water plays a serious role in power systems. So go ahead and direct bury it if you like, but it's ten times more expensive than it is to put it overhead. But we don't care... happy to do it! It's just an engineering problem... well, no problem at all.

Just to remind you that there is always nagging about utilization of the asset. In fact, the asset is terribly under-utilized. This is what we might call a "low-duration curve" here on the bottom. Percentage of time that you might use the peak of your capacity. That's an illustration, but it's pretty representative. I would say there's 20% of the system that you only use 10% of the time. That's a tremendous expense.

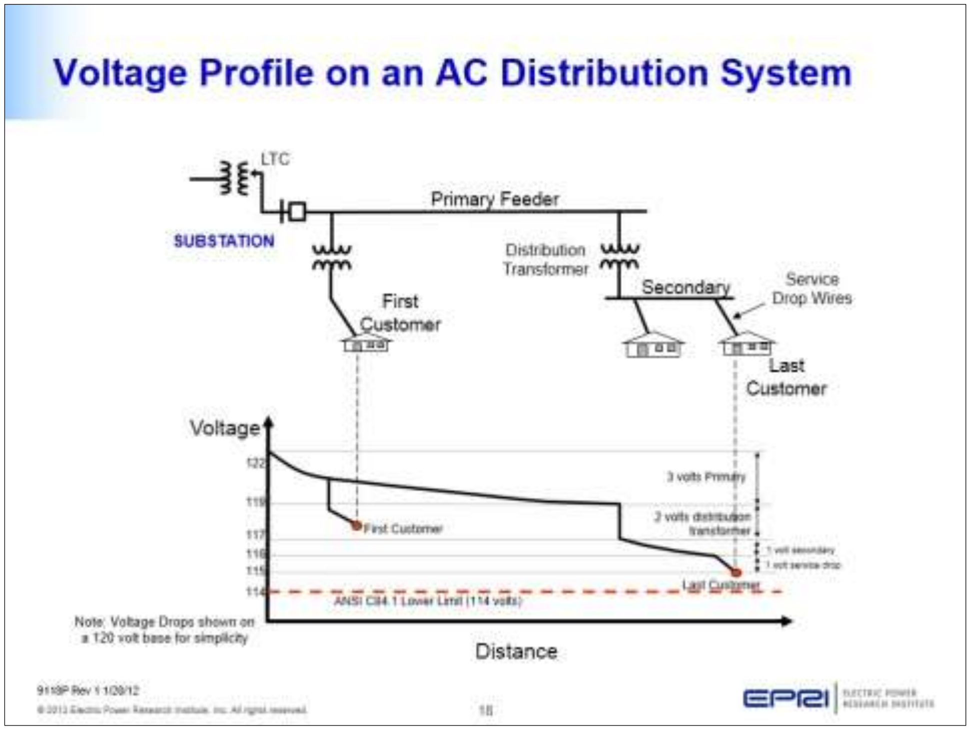

There are things we can do, and you'll hear some of those from my colleagues, to make better use of how the load is distributed. It's very important.

On top is a typical residential daily demand profile. I think this is actually Dennis Symanski's house. I don't remember, but it's close. Why is there such a bump here? You go home for lunch late in the day?

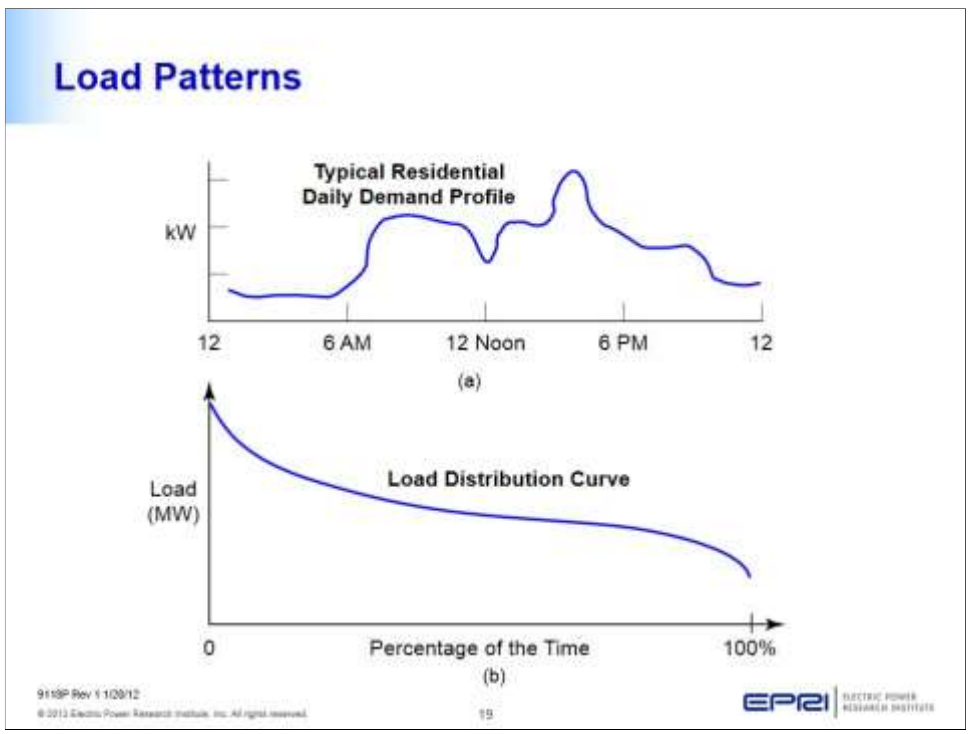

Actually, you want to reach into there and you want to say to yourself, "Let's smooth that out." If you smooth that out, you're going to affect this, right? So techniques for smoothing that out will be touched on by some of my colleagues.

We have a variety of systems. Some are quite remarkable. This is Con Ed, Consolidated Edison Company out of New York. It's somewhat unique; although, also from an electrical perspective; it looks a little like Singapore and a few other places in the world. It's got an awful lot of people; 1.8 million or so, maybe a few more, now; with over 3 million electric customers.

They still do, by the way, have 35,000-plus miles of overhead line. Not in Manhattan, Stanton Island, Westchester. They have 91,000 miles of underground line. What they do is put in a very elaborate system to ensure almost perfect reliability. They achieve this by having an overlay system with almost five points of redundancy. Very expensive, but wise given the criticality of supplying the electricity.

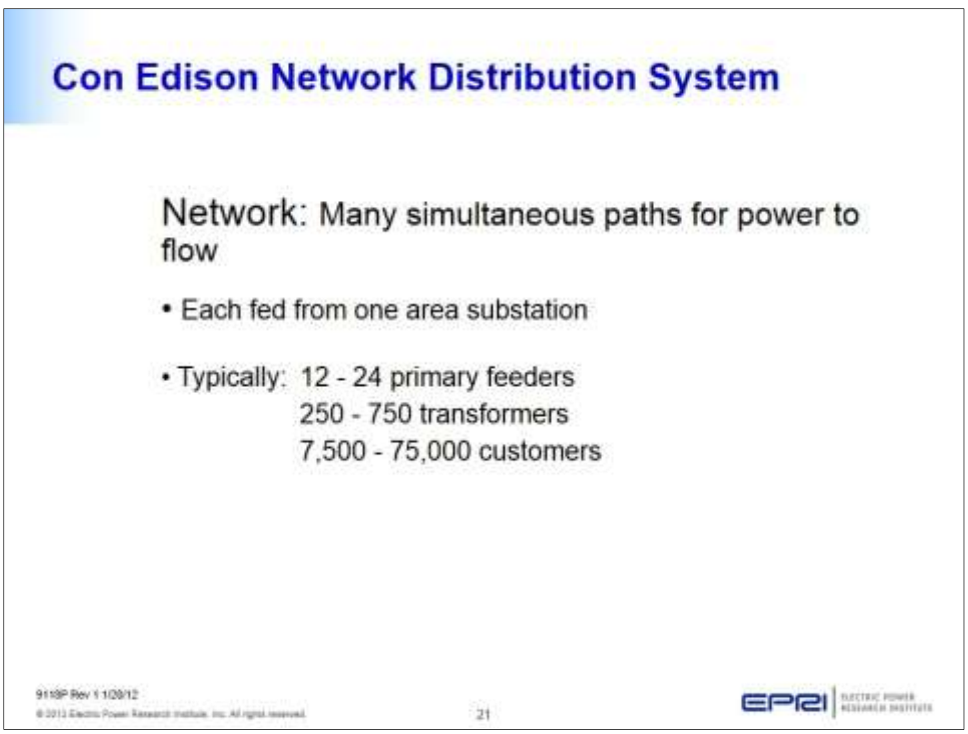

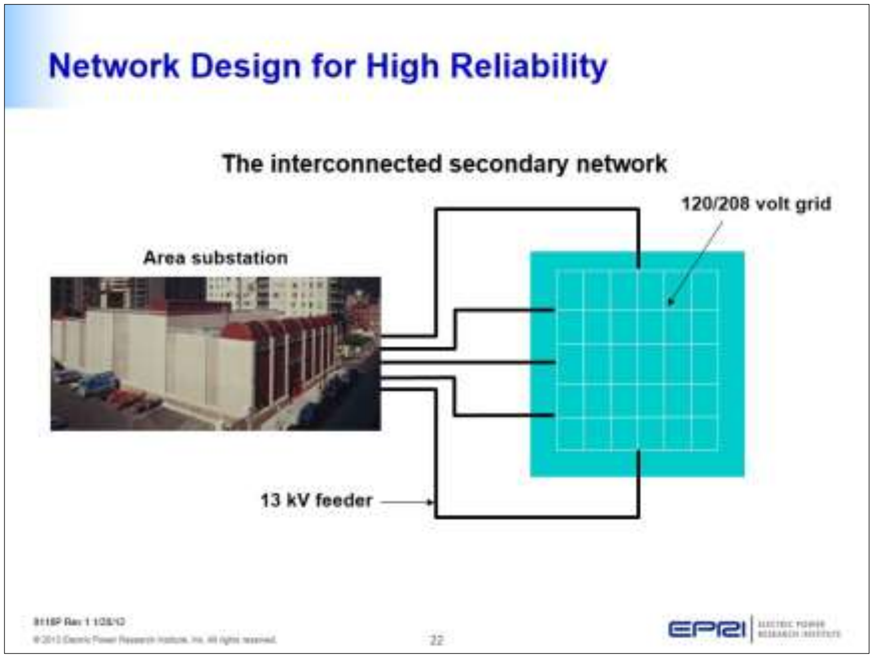

Think about it for just a second: If you live in a 20-story building and maybe you're old like me, you don't really want to be climbing up and down those stairs with jugs of water or milk or whatever you want to buy. Maybe some beer. Because, obviously, that's going to be a hell of a strain on those people. So power supply to the citizens of New York becomes critical. They can afford to and they're encouraged to put in a much more reliable power system.

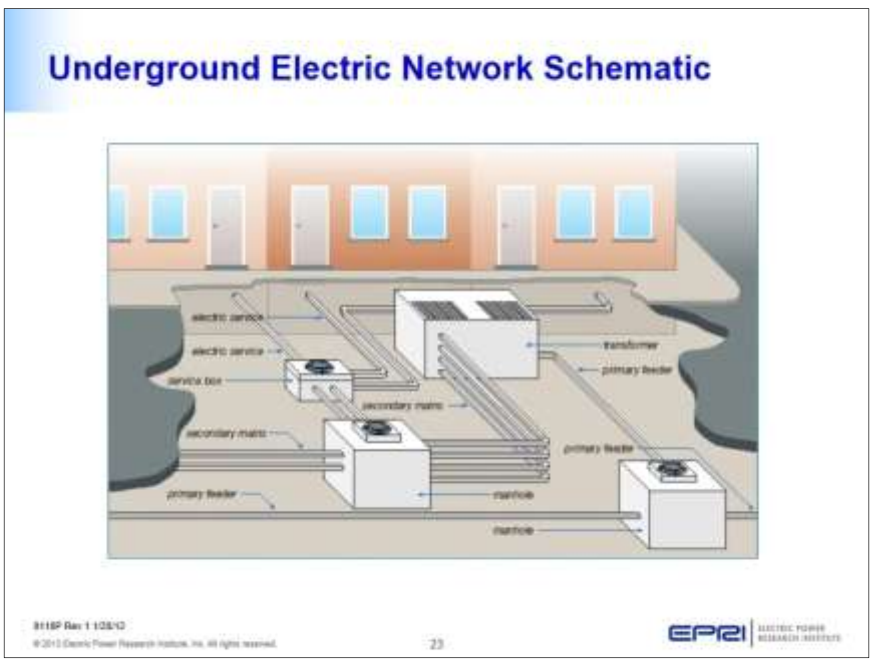

So they have this really interconnected network system. They actually build off a 120/208 grid with 13 kV feeders in area sub-stations that look like buildings you visit in New York City. You drive by and won't even know it's a substation. You wouldn't have any idea what's in there.

If I were to break that down a little bit more, you'd see a series of manholes and hand holes. They call them surface boxes. These are for splicing the secondary. These are a sensitive point in the Con Ed system, or any system that uses this kind of configuration because those splices are made by hand and often immersed in salt water. They put salt on the streets and salt water corrodes those connections, weakens them. You can end up with a voltage potential difference between that, let's say the cover of that surface box, and something else, say the metal edge of a curb. In fact, there was a young woman electrocuted in that way two years ago, which caused quite a stir.

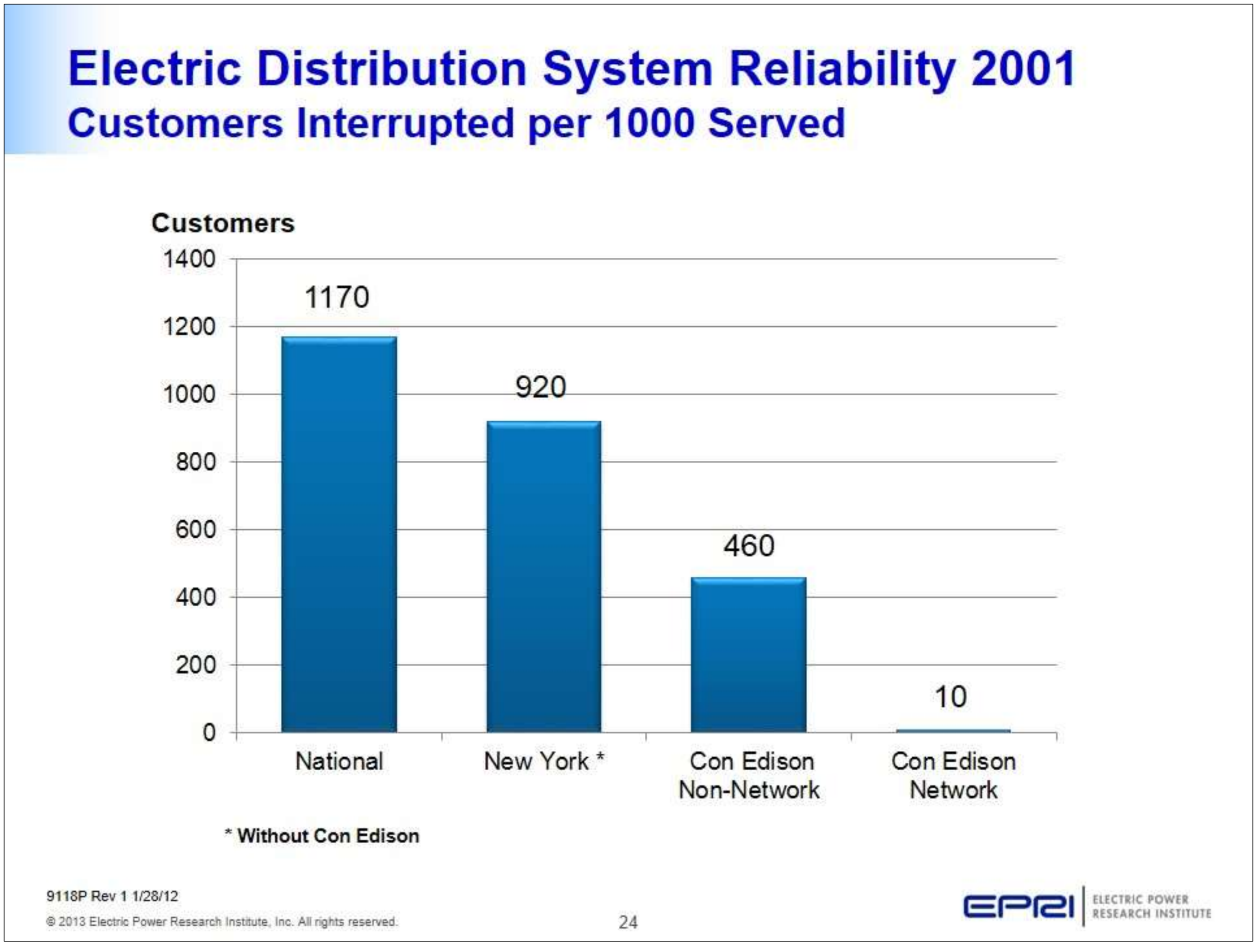

So there are some foibles with all of these. Reliability, however, as a result of a system like that, is phenomenal. So these are customers interrupted per thousand served. This is a couple of years old. The average consumer in the United States is without power for about 110 minutes, if you set aside hurricanes and stuff like that. 110 minutes! So any given day, there are potentially half a million customers in the United States without power for at least 20 minutes.

We do hear about reliability occasionally. My colleague mentioned it in his introduction. But, for the most part, you can make a system very, very reliable — like Con Ed's network.

More likely you're going to live with lesser levels of reliability. It's not because the industry consists of bad people. This is mostly because of distribution outages: 85% of all outages are from the distribution system. By the way, 85% of those are due to squirrels and other animals. The squirrel is public enemy number one for power distribution.

This makes things extremely complicated. This is an old photo, but I'll show you in a second that this is what you'd see if you opened up a street in New York. It is phenomenally complex because you have all this other stuff going on, like cable TV and the telephone system. It all wants to occupy these spaces. Here are Verizon's duct banks near the World Trade Center. That picture was taken during the reconstruction of the area around the World Trade Center.

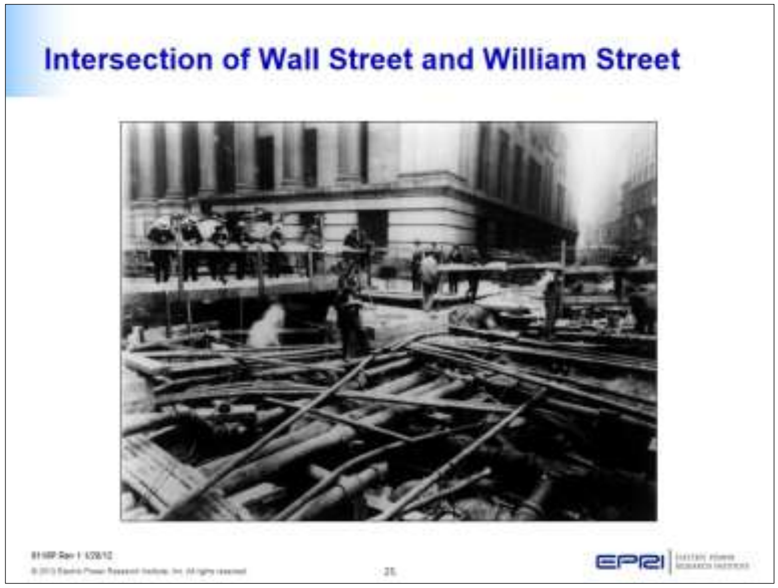

It's a phenomenal problem that we have with where to put this stuff underground. We've gotten so elaborate now, we're actually boring tunnels in some place in the world. We've done this in Berlin. But here in New York, the same kind of tunneling technology is used for automobiles. In the case of Con Ed, this first tunnel is way below any of the other infrastructure so they can move gas, steam, power, etc. under the streets without the problem of interfering with all these other infrastructures.

This comes with another problem. Do you remember, after Sandy; there was an explosion and a high-voltage transformer blew? They couldn't get to that transformer fast enough to de-energize it and take it offline. And, as I've already intimated, water and electricity don't mix real well, so these tunnels fill up with water and they have to be emptied somehow.

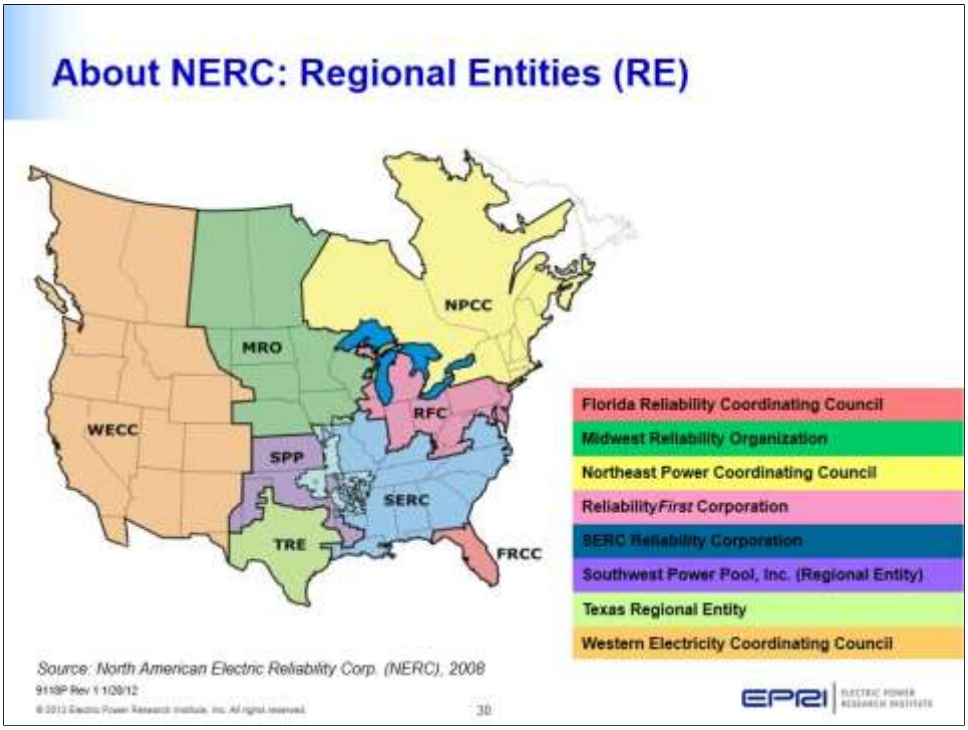

You're going to hear a little bit about how a power system operates in a few minutes, but the power system basically breaks up into pieces we call control areas. These control areas are a logical combination of generations of sources and some of that high- and medium-voltage stuff I mentioned. All those wires! There is an operator, a series of operators, computer systems, and the like, that sit together and work through that control system. These control areas are coordinated by these regional entities. These are reliability organizations that exist regionally in the United States so somebody is watching the system. There are, as you'll hear about the California ISO in a few minutes, also state-level organizations that coordinate how these individual control operators work.

This is an old image of control areas, but it gives an idea that there are quite a few of them and coordinating them is quite a job. Don't jump to the conclusion that we should do away with them, combine them, or whatever… because there are some complexities with doing that as well.

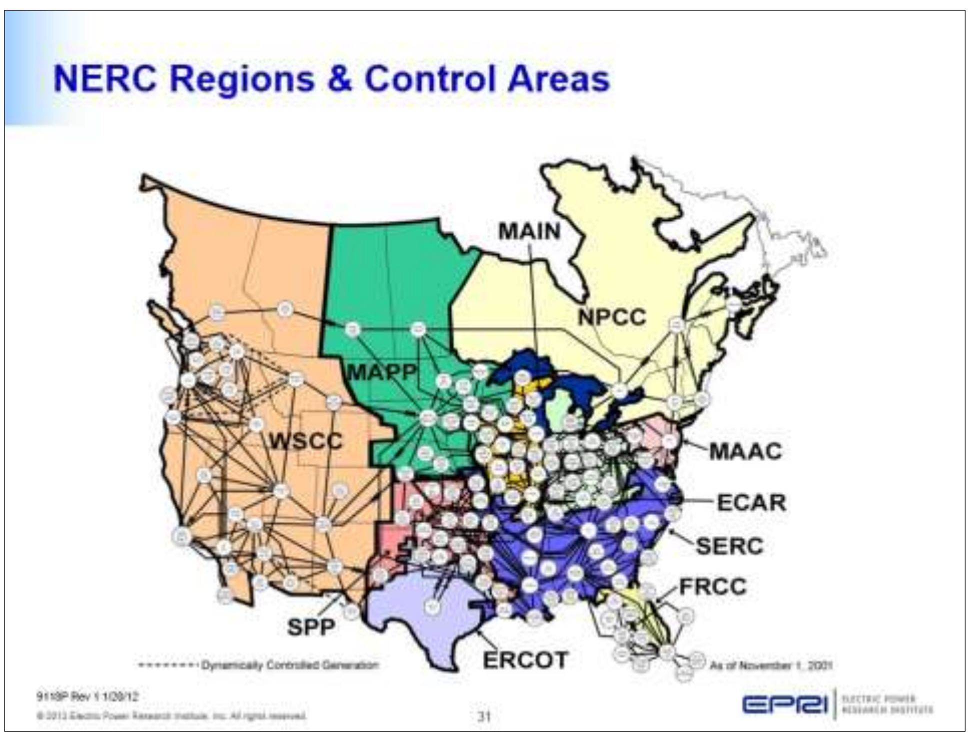

You'll hear a little bit more about how this all comes together. Obviously, there is quite a complex set of operations that take place, driven by some master computers that are always under development, very expensive, and very complicated. They're part of the power system and they break down into various aspects of the power system: generation, transmission, and interfaces with customers. Their job in these control centers is to balance the load.

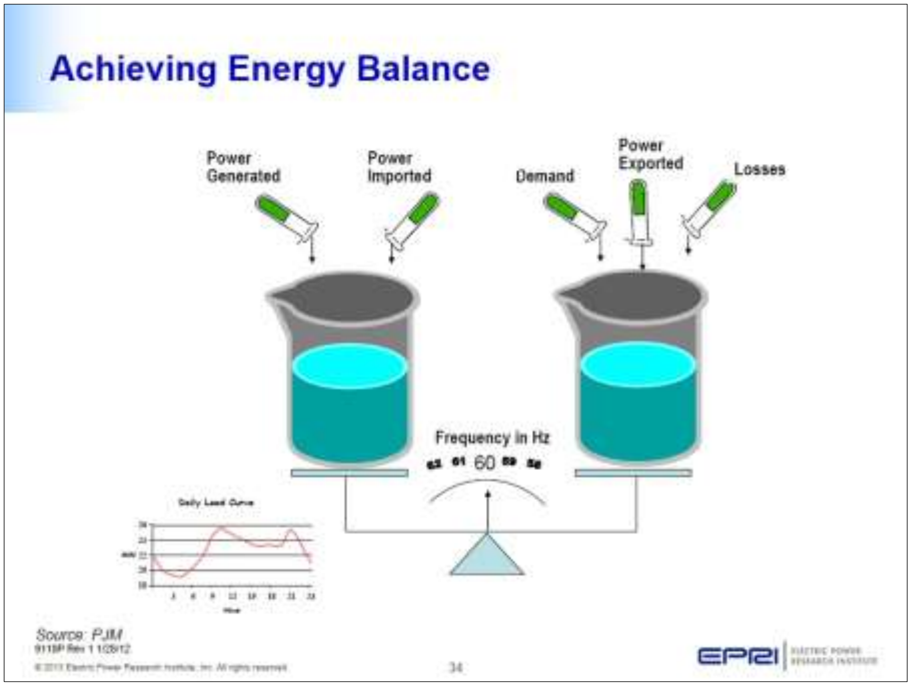

Storage is somewhat minimal. It exists, we have some storage, but we'd love more cost-effective storage. We really have to balance supply and demand. Balancing supply and demand allows us to keep the frequency steady and the system stable. Start swinging because of malfunction, a short, a disturbance, a loss of a major generation piece… and you throw this frequency out of whack and out of balance and the system begins to collapse. You collapse a system like any of ours and it's a hell of a job putting it back together and getting it all running again.

So the idea, always, is to worry about demand and supply; both what I've got in my area as well as what I'm importing from other areas.

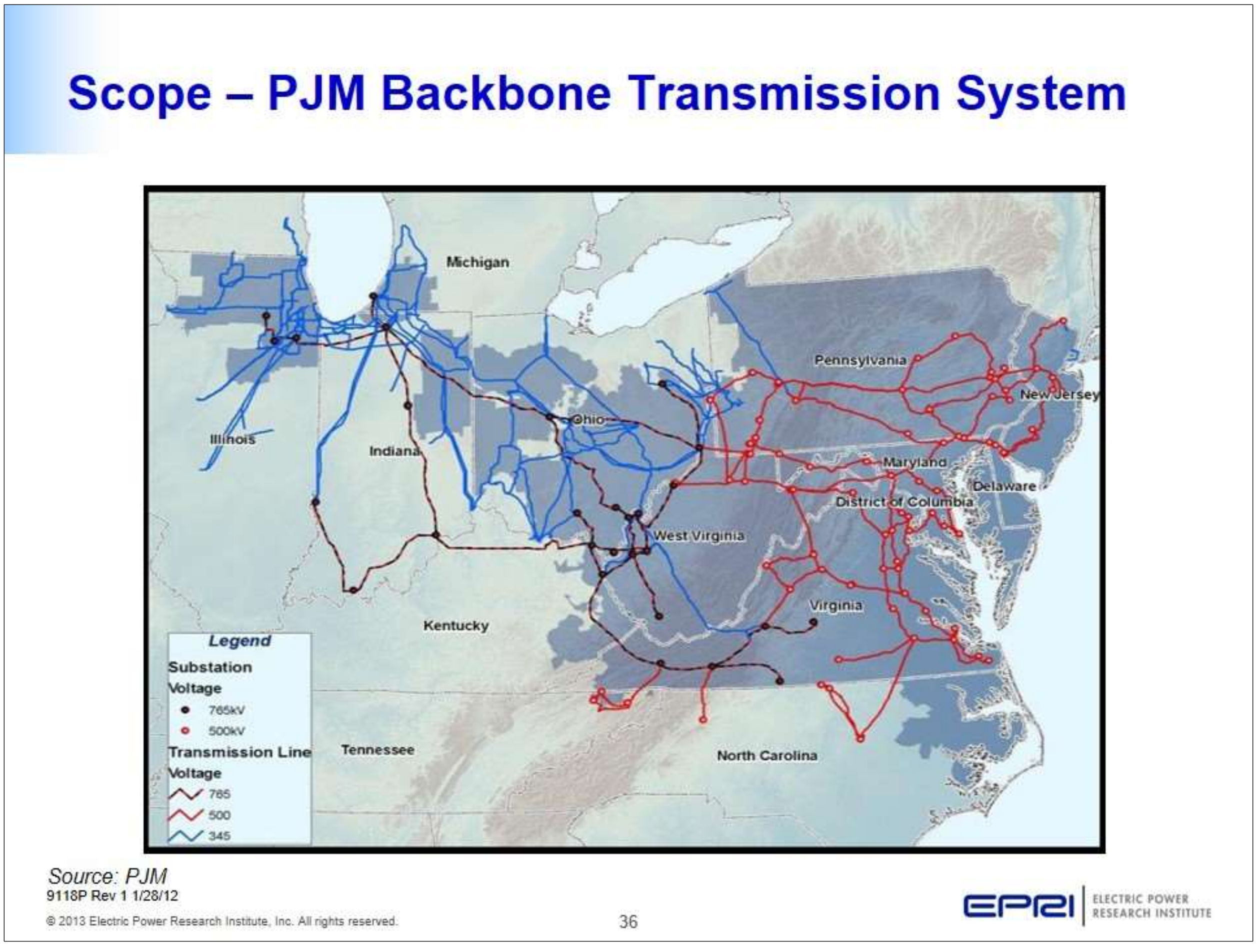

I think that you have a taste for some of the transmission backbone now.

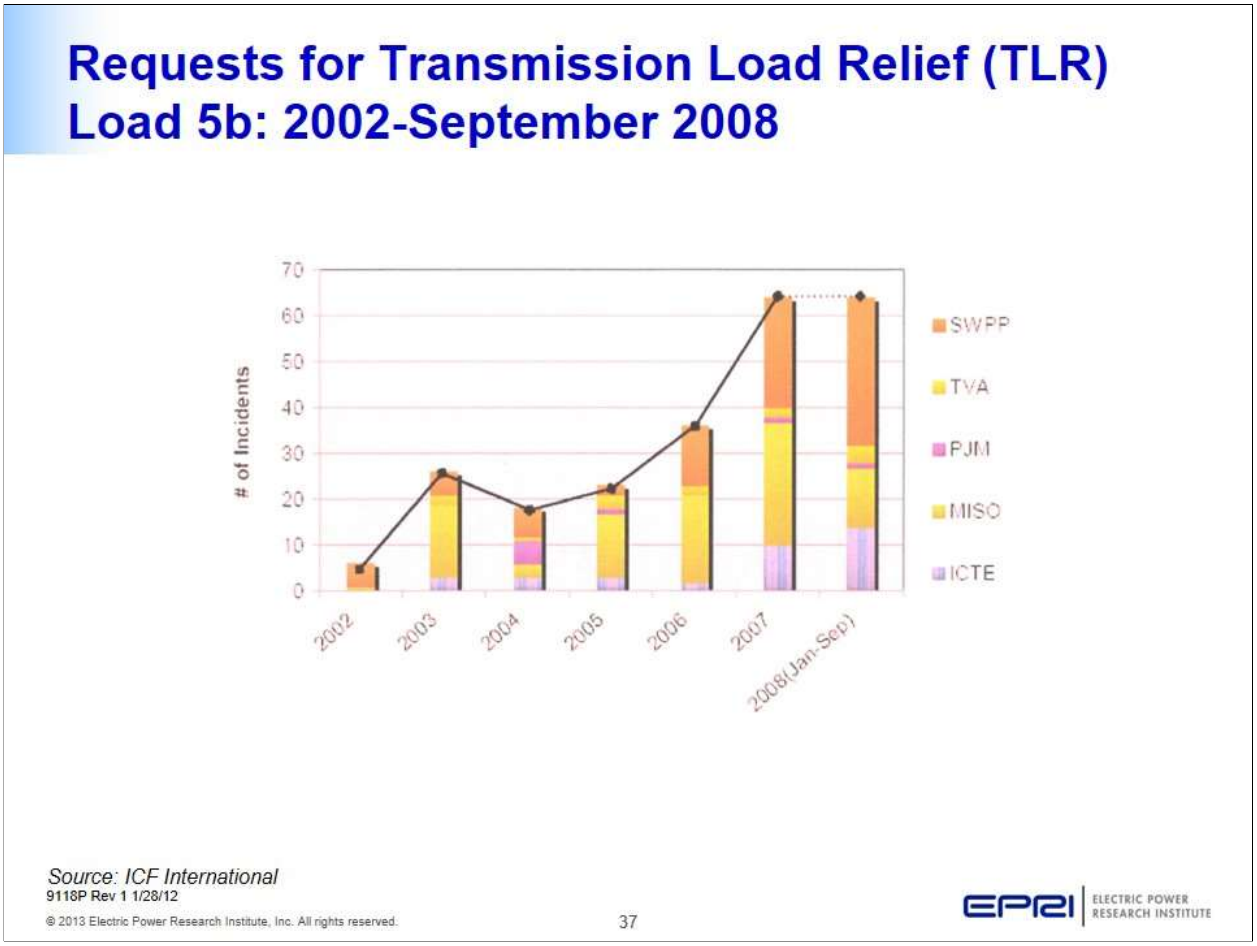

One of the constraints in a power system is that people like to move power for economic reasons. Unfortunately, power moves according to physical laws. Remember Kirchhoff? Some of you have studied Kirchhoff. There's the old joke about the politician who asks why power can't move from here to there instead. The answer is, "Our flow follows Kirchhoff's law." The politician replies, "Well, if it's a law, we'll just repeal it."

We try to operate the power system so that everybody can play in the market, so that any generator can bid in. Supply, I want supply; I'm willing to pay. I'm willing to do it for so much.

However, it's constrained by the system. You can't afford to build an infinite bus. You can't afford to have a copper plate cover the U.S. so that we can use anything anywhere anyone wants to. It just can't be done. So we set up a system that allows people to bid in. When there's a constraint, the operator can say, "Nope. I can't do it for this good reason." That does inhibit the market, but we've got a system where there are set rules that can be governed.

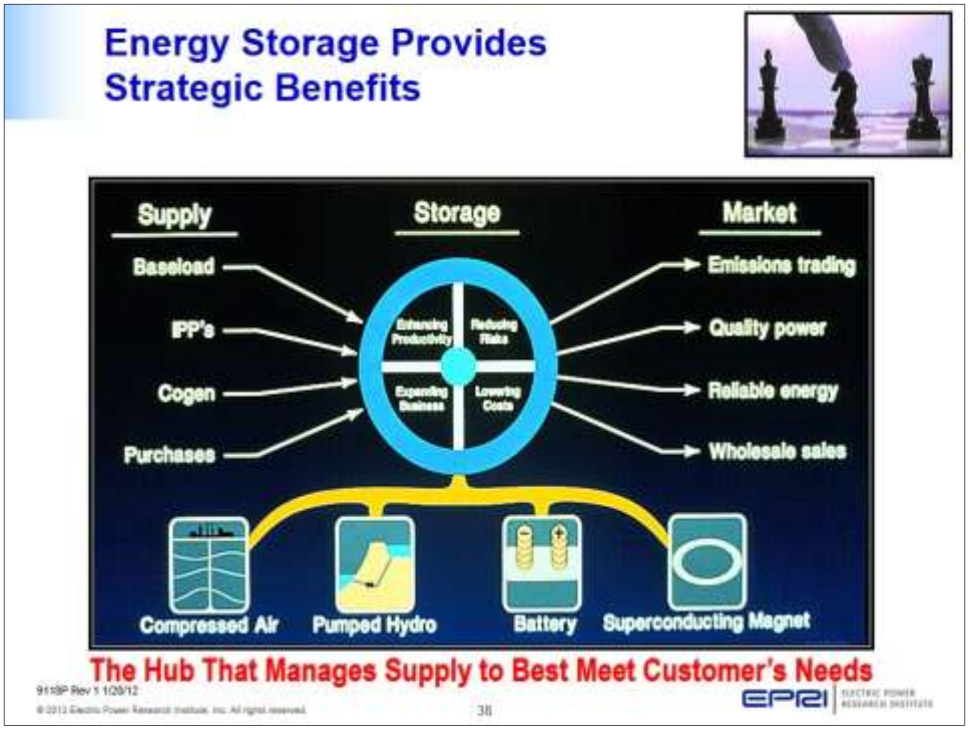

Storage would be great. It would provide us energy balance.

If we could do storage effectively -- it's the Holy Grail, there's no question -- we could really improve things a whole lot.

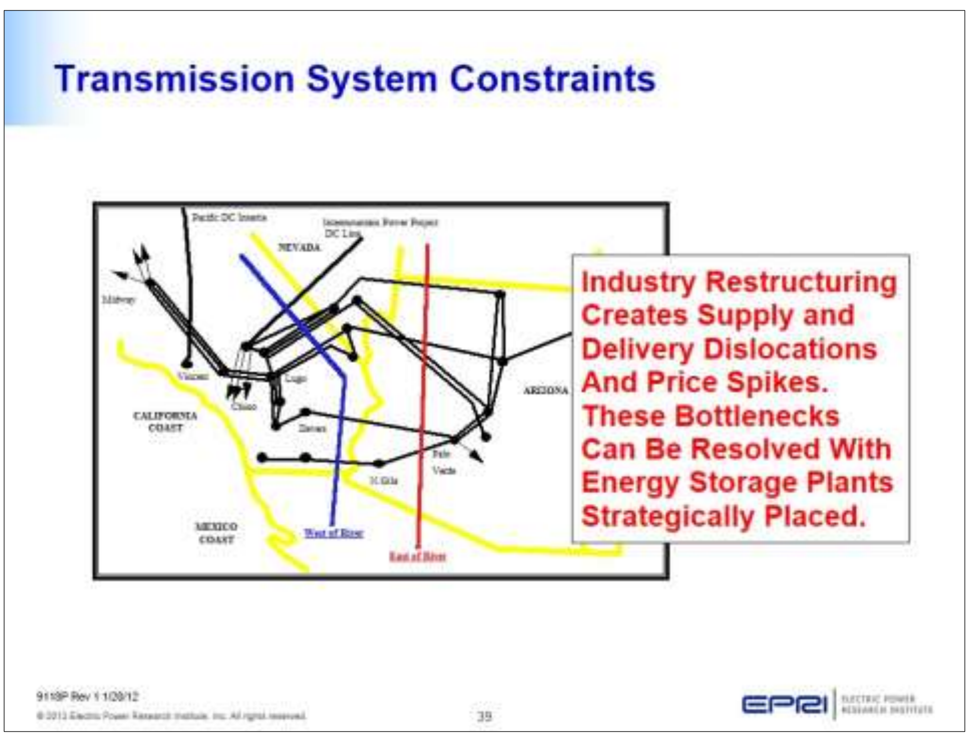

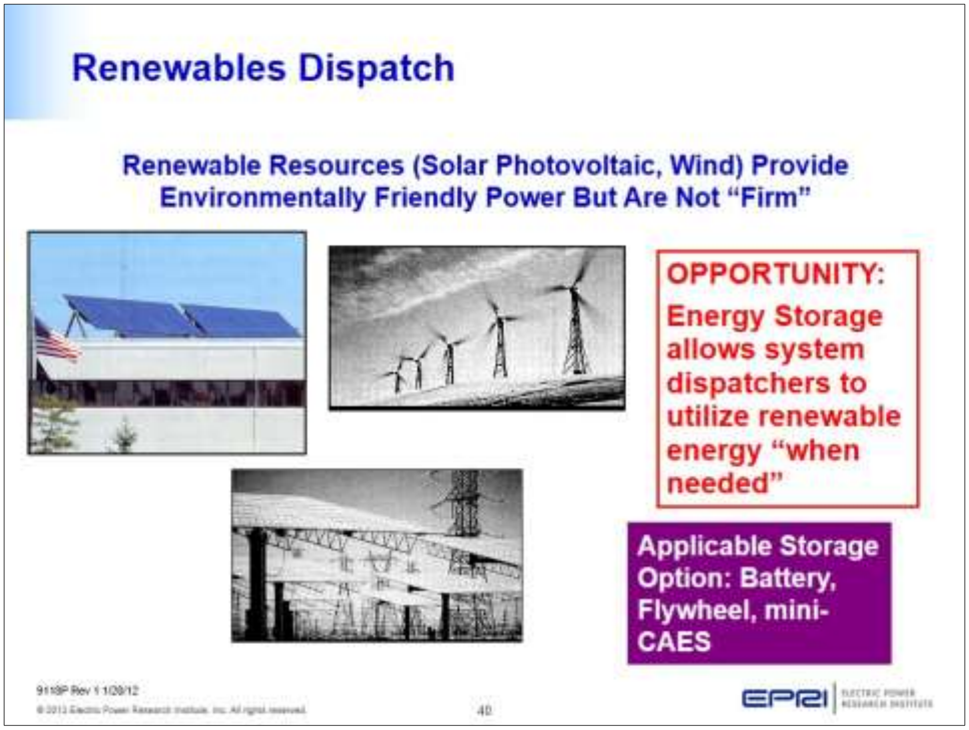

You're going to hear a little bit about the smart grid and that renewables are inherently variable. Some have said intermittent. If I say intermittent, that almost implies a bias. I don't want to do that; but the sun doesn't shine all the time, right? It doesn't shine when clouds go over, so photovoltaics are sensitive to this. The wind is surprisingly variable. If you look at data on a wind power output, you'll be amazed at how much it varies. Even if you stood on a wind farm and thought, "Wow, that's pretty consistent." …actually, no. The pressures really vary.

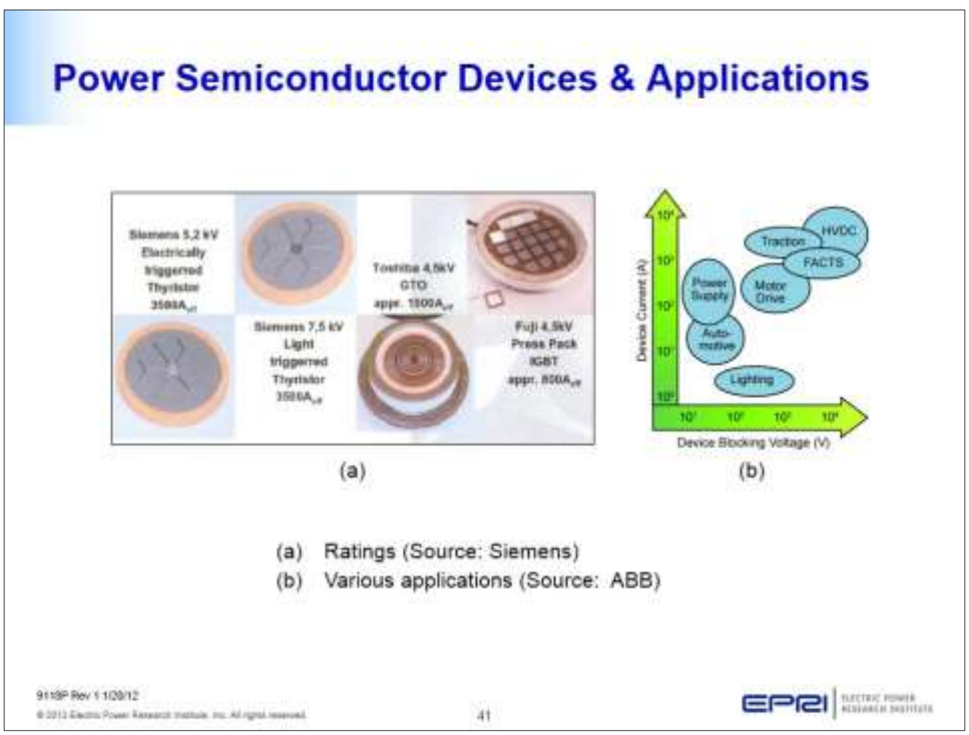

Another piece of technology you'll hear about is the use of power semiconductor devices. You're already familiar with them on a different scale. If we could use those really effectively in the bulk power system, there's a lot we could do, like controlling the power system better, and a lot to make more effective use of other devices in the power system.

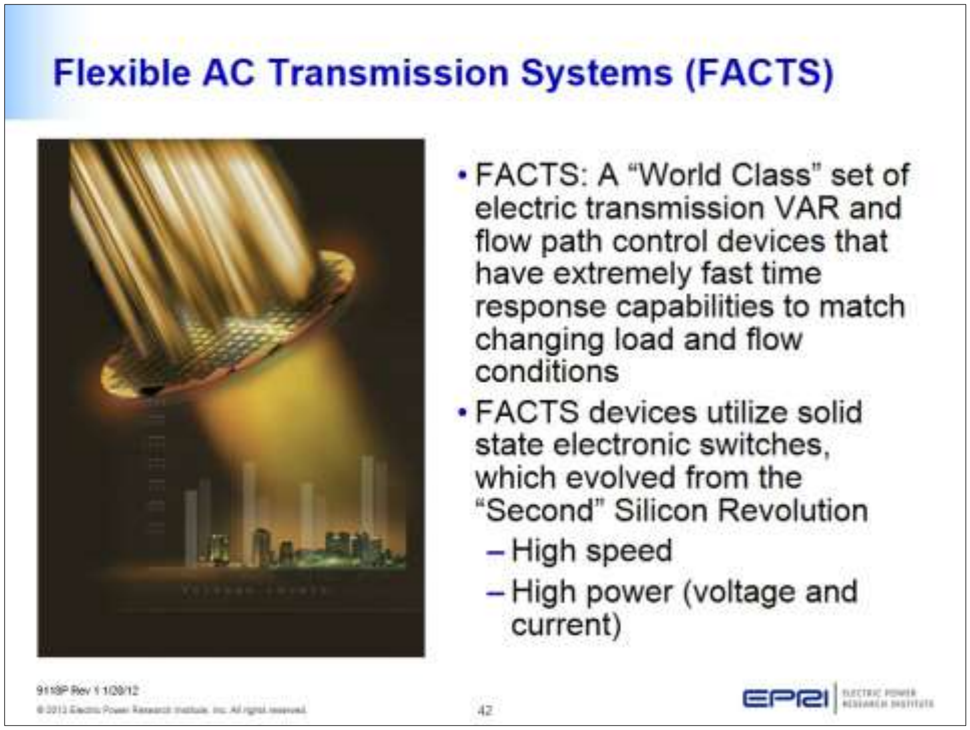

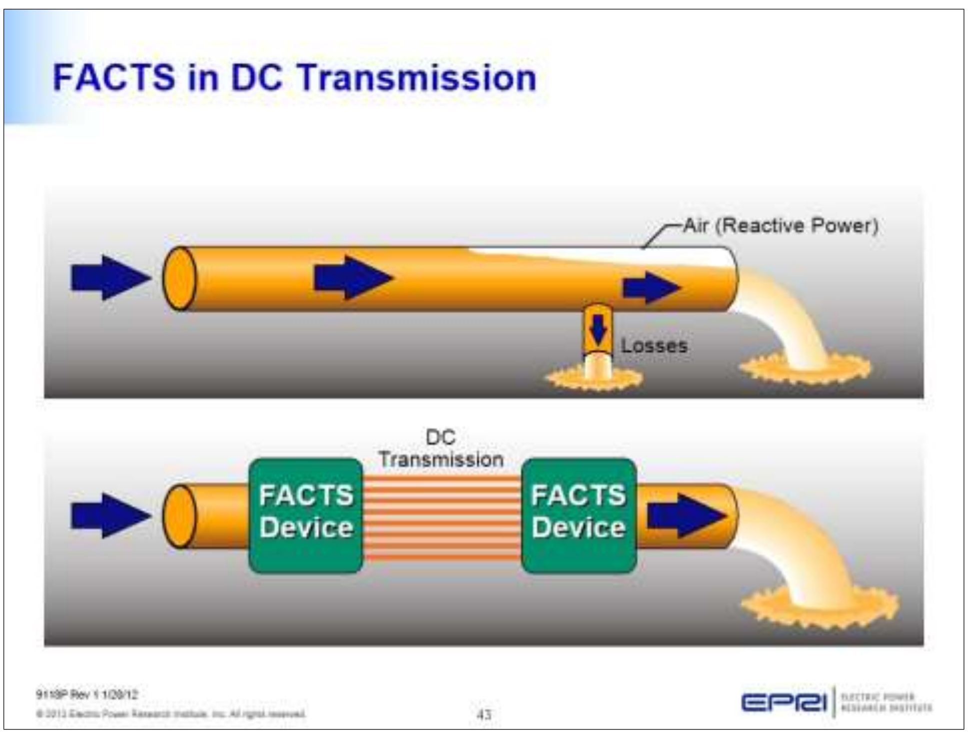

There have been inventions. These are EPRI, actually, flexible ac transmission systems (FACTS), a class of devices that can actually take… Basically, you can take ac power from one place and convert it to dc, could chop it up, and put it back on the same power system, adjusting phase angles and so on. Or we could inject it into another line going nearby. If we use this, we can always use it in dc transmission.

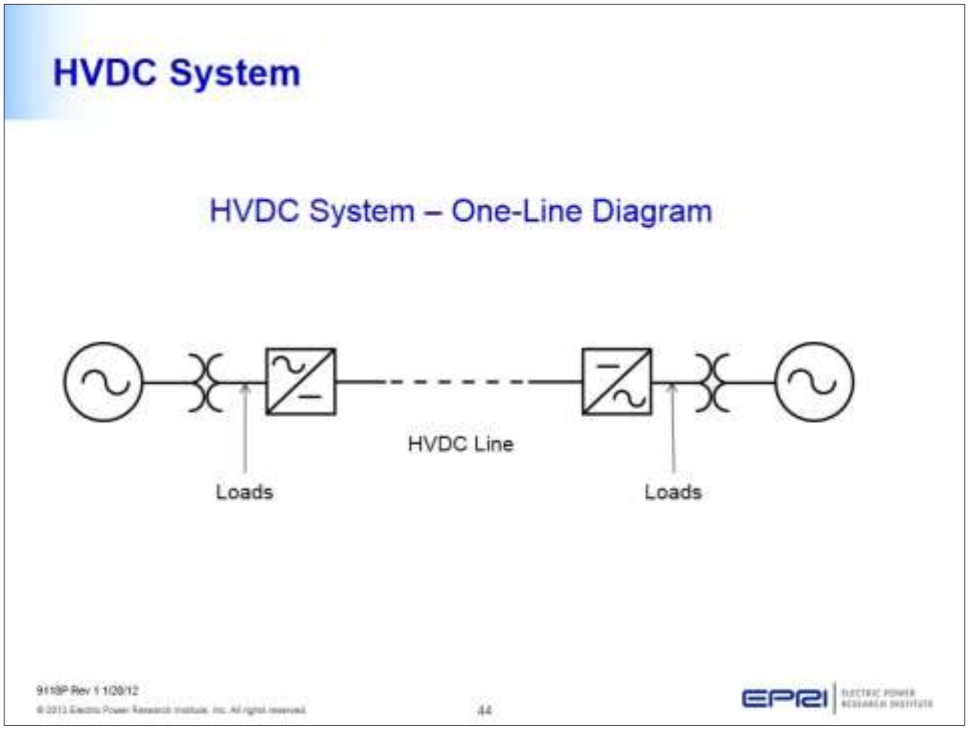

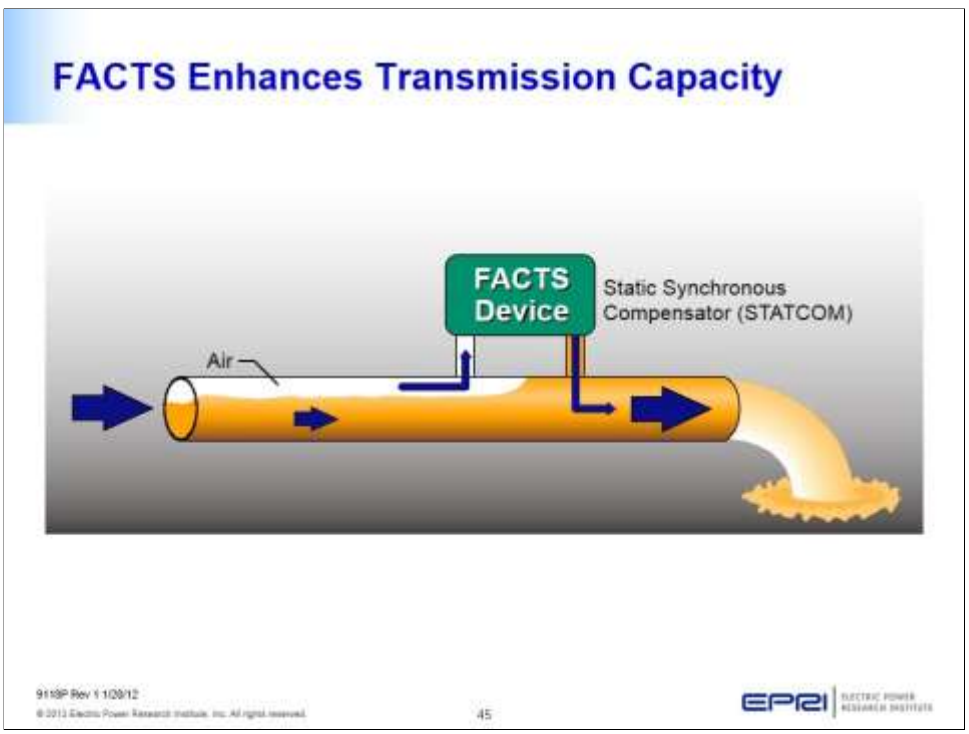

I use the analogy of air reactive power. You got the real and reactive stuff and I can employ a FACTS device to use dc transmission. When I do this, I can move a lot of power from one point to another with a very strong system. I wouldn't have the vagaries I would have necessarily with an ac system. I could use FACTS devices to enhance transmission capacity even on one line, by taking out power and re-injecting it at a different phase angle. I can control the amount of power that I have on that line. I can redirect flow from one device to another.

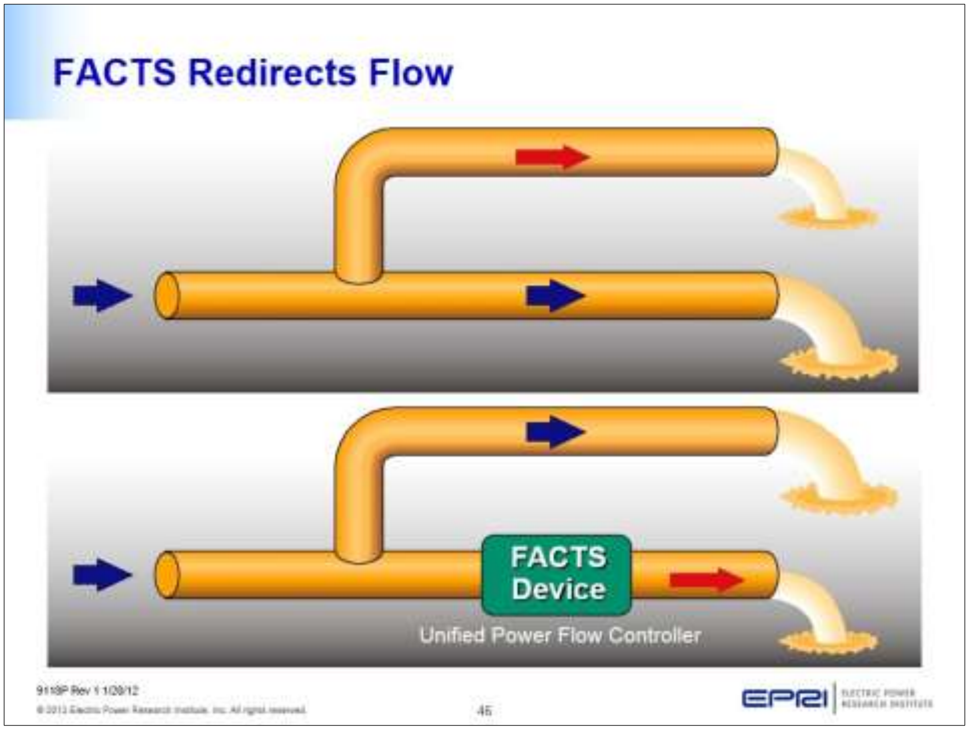

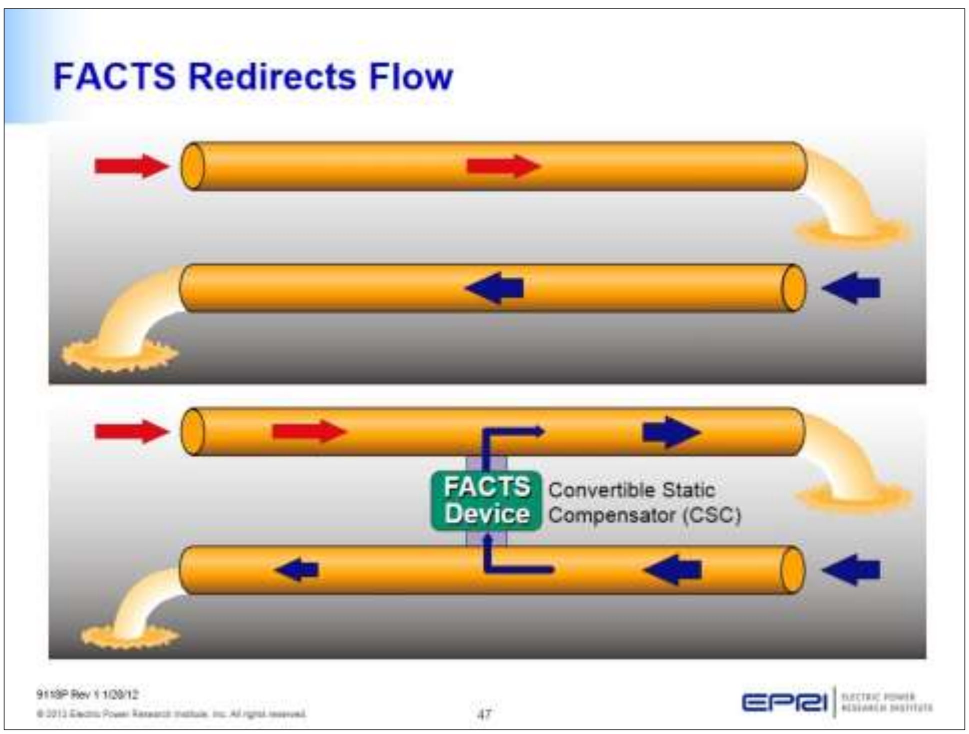

But these devices are expensive and they're expensive because they contain power electronic devices. So one thing in the future is the whole system of power electronics, the valves and all of the power supplies and the like, that will go into these devices.

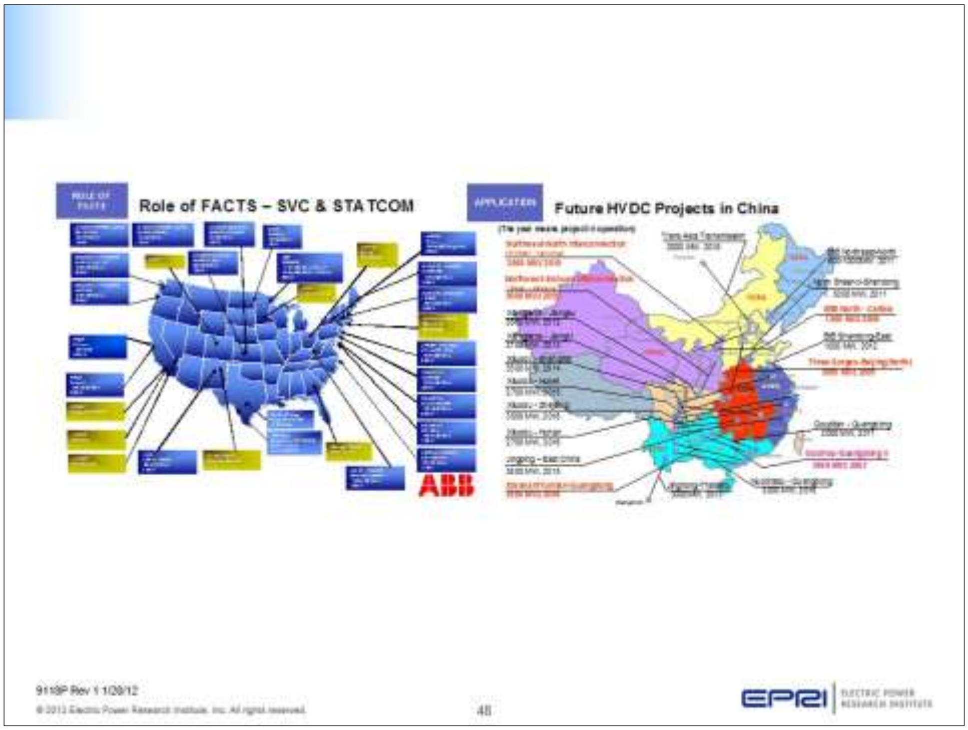

We see increasing numbers of these being deployed. Not all exactly FACTS, but something called a "static core compensator," which basically adjusts the power factor. It will have a huge impact because of what's going on in China; the technology development, at this point, is going forward in China.

We have a few projects that are high-voltage dc that use these devices in North America.

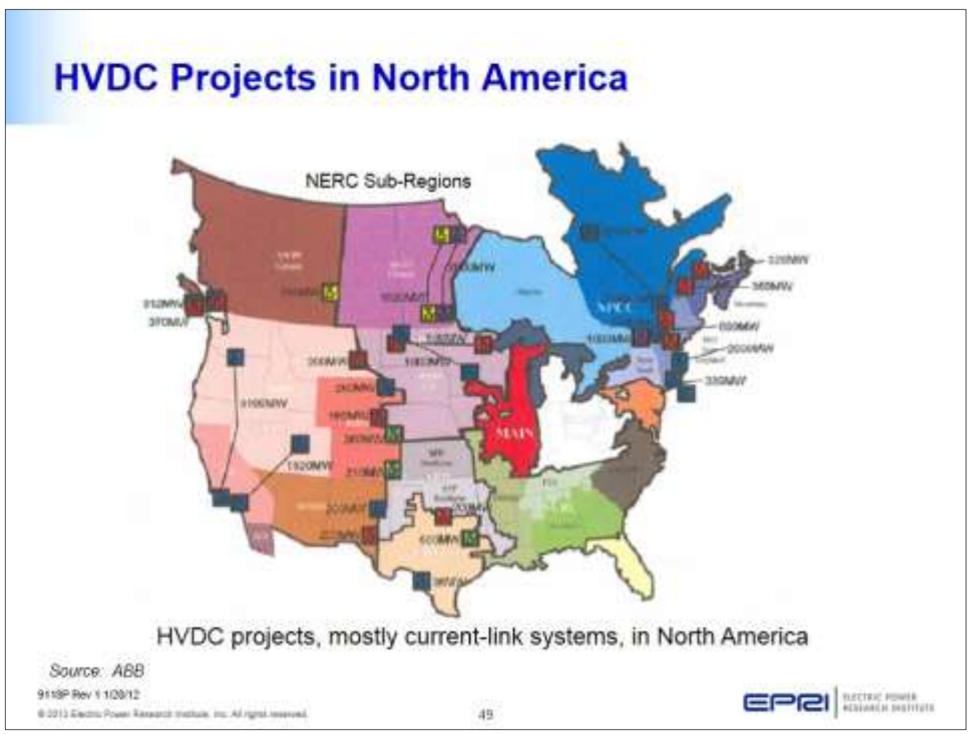

One of the things that I think we'll touch on during the day that's really of concern to us… is the losses in this value chain. We lose a lot in generation.

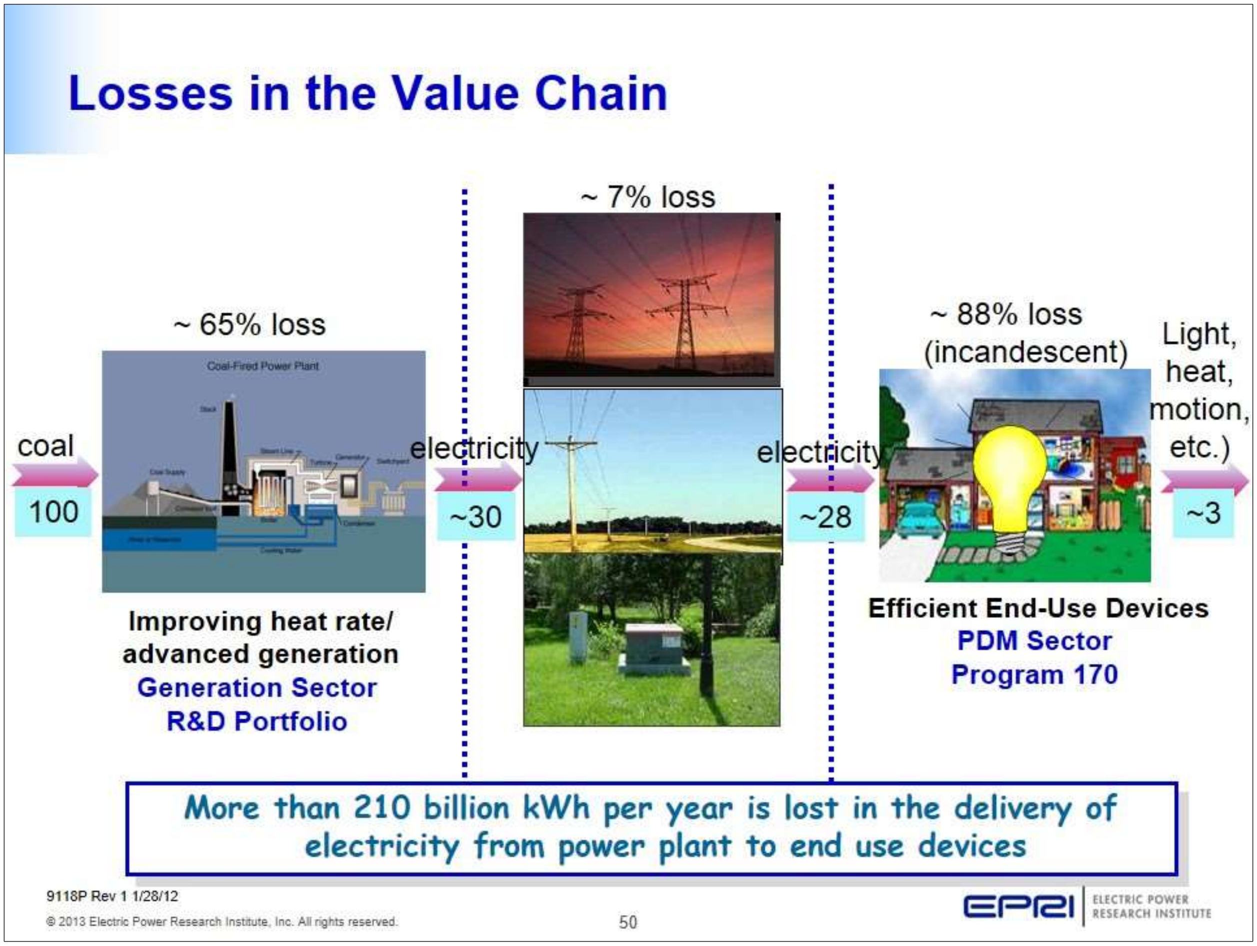

Being very unkind to us here, I took an old coal-fired plant. I lose 65% of my energy from input to output in generating power at an old coal plant. Those numbers are a lot better from newer, combined-cycle, gas combustion turbines. I then lose 7% of my electric energy through transmission and distribution (national average varies). I end up, so far, with 28 of my 100 units in. If I'm using an incandescent light -- yes, I know, we're moving away from them — but still, I end up with effectively only using 3 units of energy!

One of the things you hear people from EPRI talk about is that it can't go on like this. We need technologies to boost efficiency. I point to one here that we've been improving, but we need technologies that are going to dramatically change this equation.

By the way, we use 11% percent in total. We lose another 4% of our electricity inside pumps, fans, and stuff like that. And that, we think, can be changed by as much as 20%.

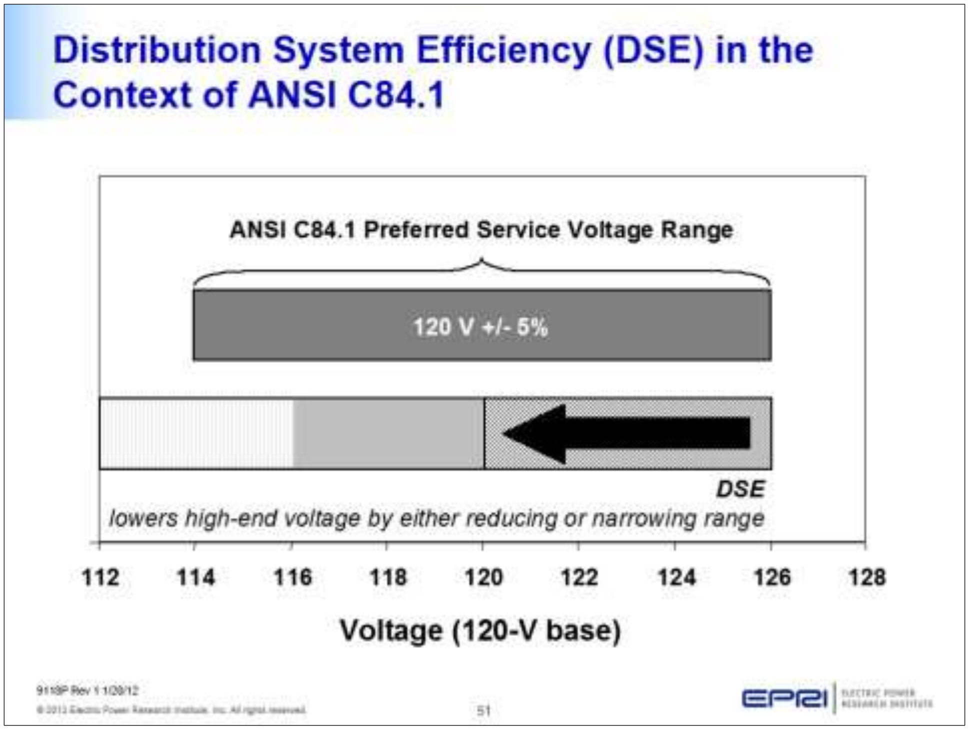

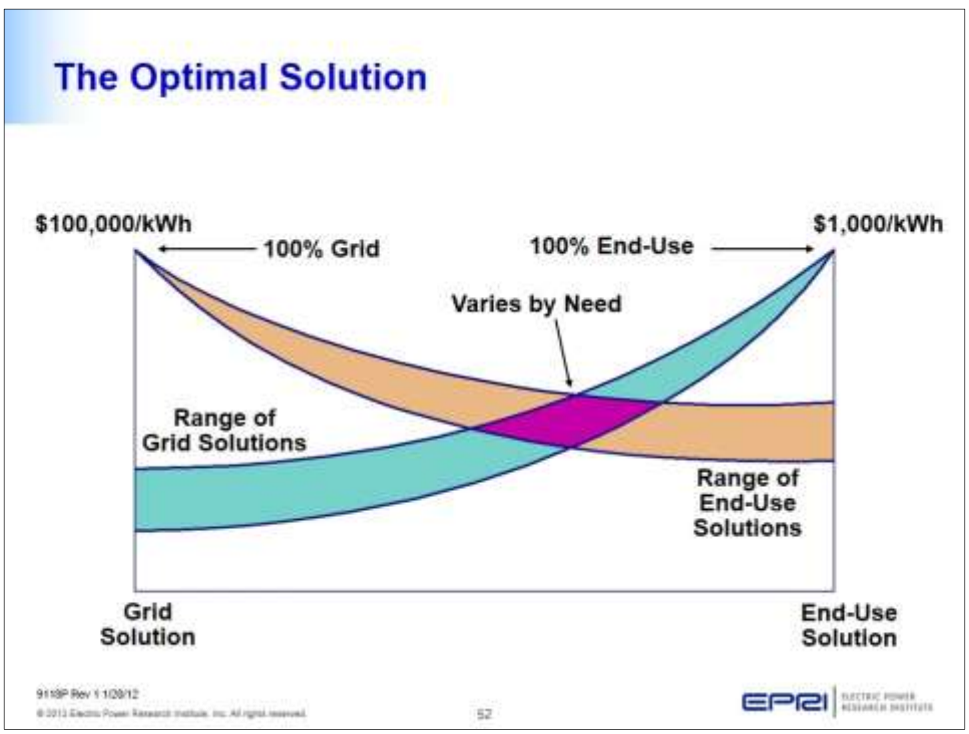

The optimal solution to us for the power system of the future is probably one that doesn't just build grid, but also recognizes end use. The optimal solution for us is somewhere in between.

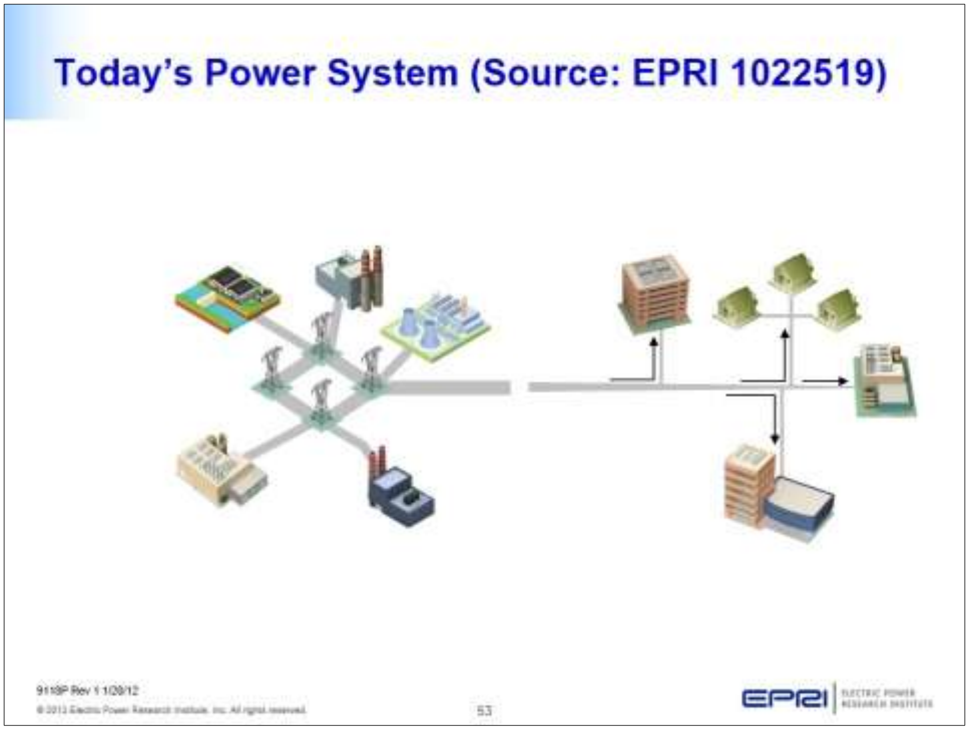

Today's power system, as I've described it, is bulk power generation. We're using a system; we often call it the grid. We're using that system to bring that power to consumers of all kinds.

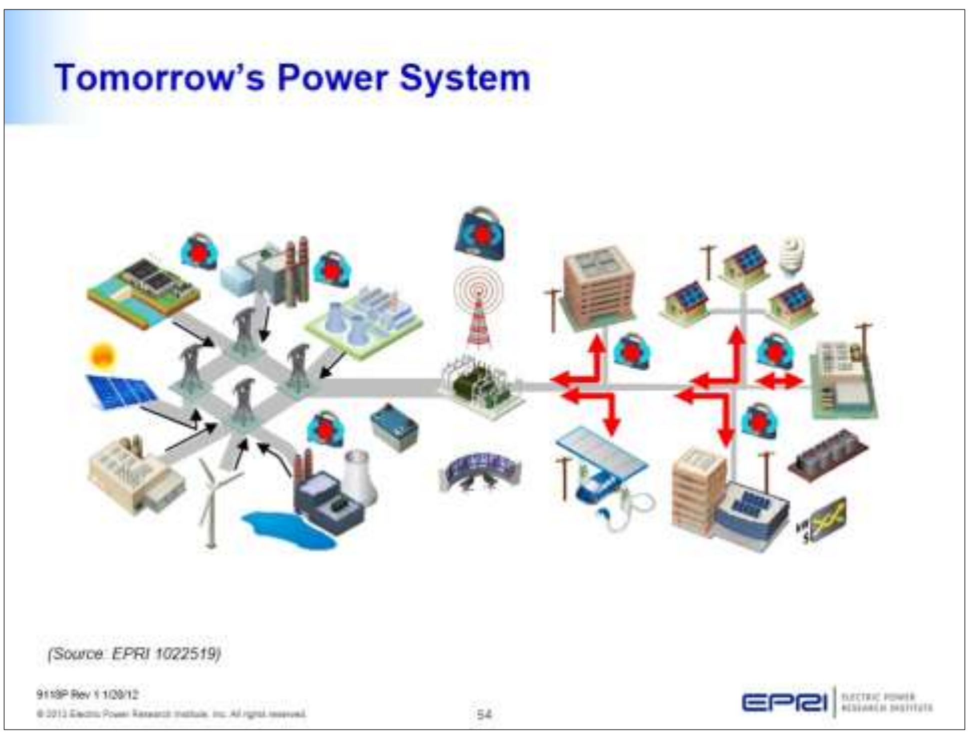

Tomorrow's power system is going to be an awful lot more complex. We already have 380,000 buildings in the United States that have photovoltaic. We are seeing increased adoption of plug-in electric vehicles. We're seeing more interest in things like microgrids and the complexity these bring is phenomenal.


Technologies Involved in Distributed Energy Resources (DER)
Renewables
PV
Building Master Controller
Advanced Metering
Distributed Generation and Storage
Plug-in Hybrids
Smart End-Use Devices
Thermal Storage
Ground-Coupled Heat Pump
9118P Rev 1 1/28/12
© 2013 Electric Power Research Institute, Inc. All rights reserved.
55
EPRI ELECTRIC POWER RESEARCH INSTITUTE

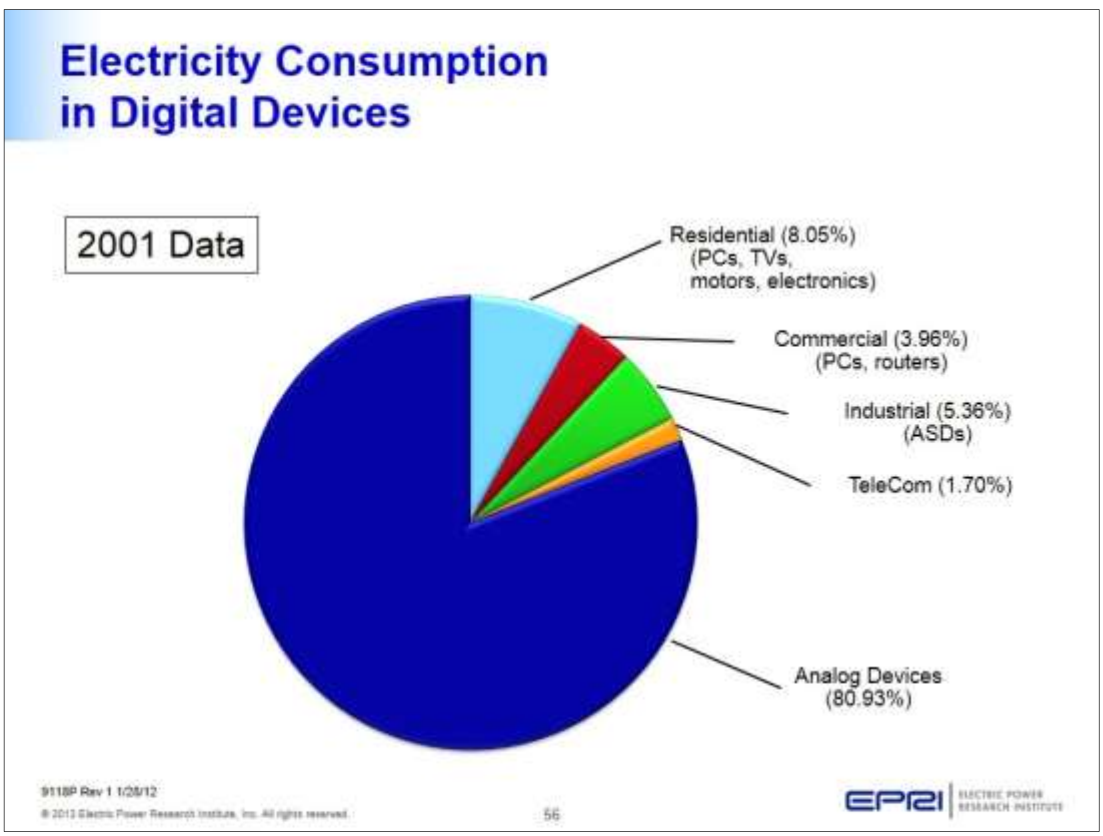
Electricity Consumption in Digital Devices
2001 Data
Residential (8.05%) (PCs, TVs, motors, electronics)
Commercial (3.96%) (PCs, routers)
Industrial (5.36%) (ASDs)
TeleCom (1.70%)
Analog Devices (80.93%)
9118P Rev 1 1/28/12
© 2013 Electric Power Research Institute, Inc. All rights reserved.
56
EPRI ELECTRIC POWER RESEARCH INSTITUTE

Rise of the Digital Economy

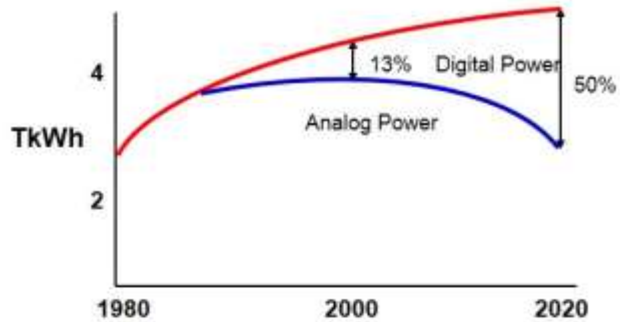

- Demand for "digital quality" power is growing rapidly
- Requires much higher reliability and quality
- New devices have different characteristics

A Toll Felt Throughout the U.S. Economy

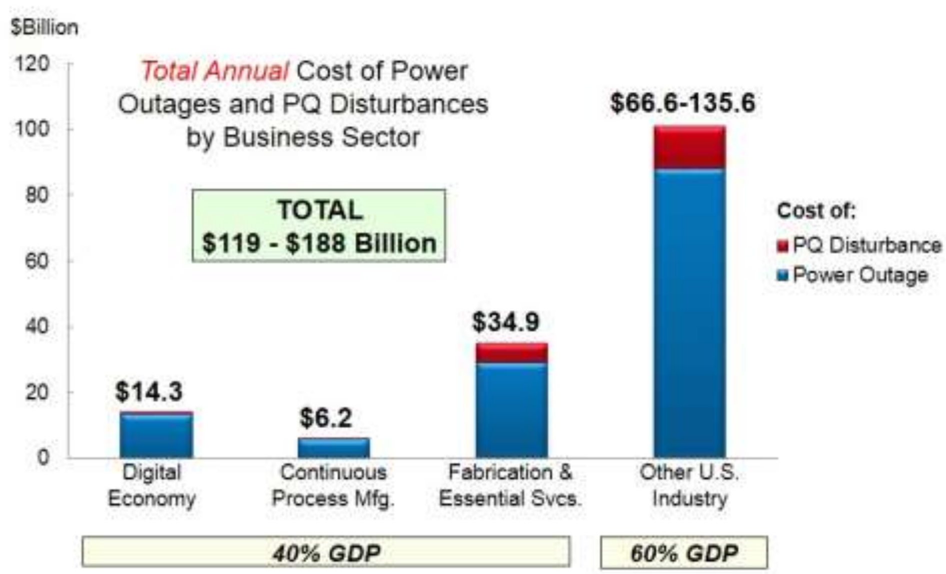

Source: Primen Study: The Cost of Power Disturbances to Industrial & Digital Economy Companies

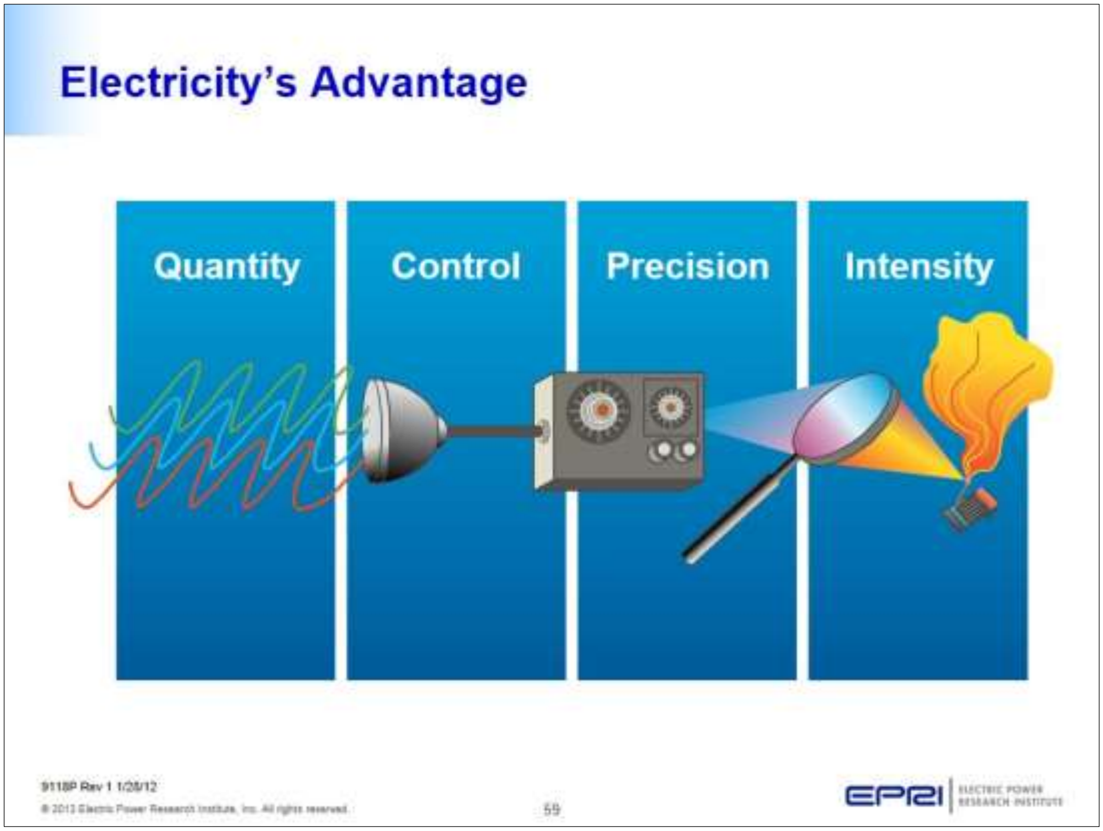
Electricity's Advantage
Quantity
Control
Precision
Intensity
9118P Rev 1 1/28/12
© 2013 Electric Power Research Institute, Inc. All rights reserved.
59
EPRI ELECTRIC POWER RESEARCH INSTITUTE

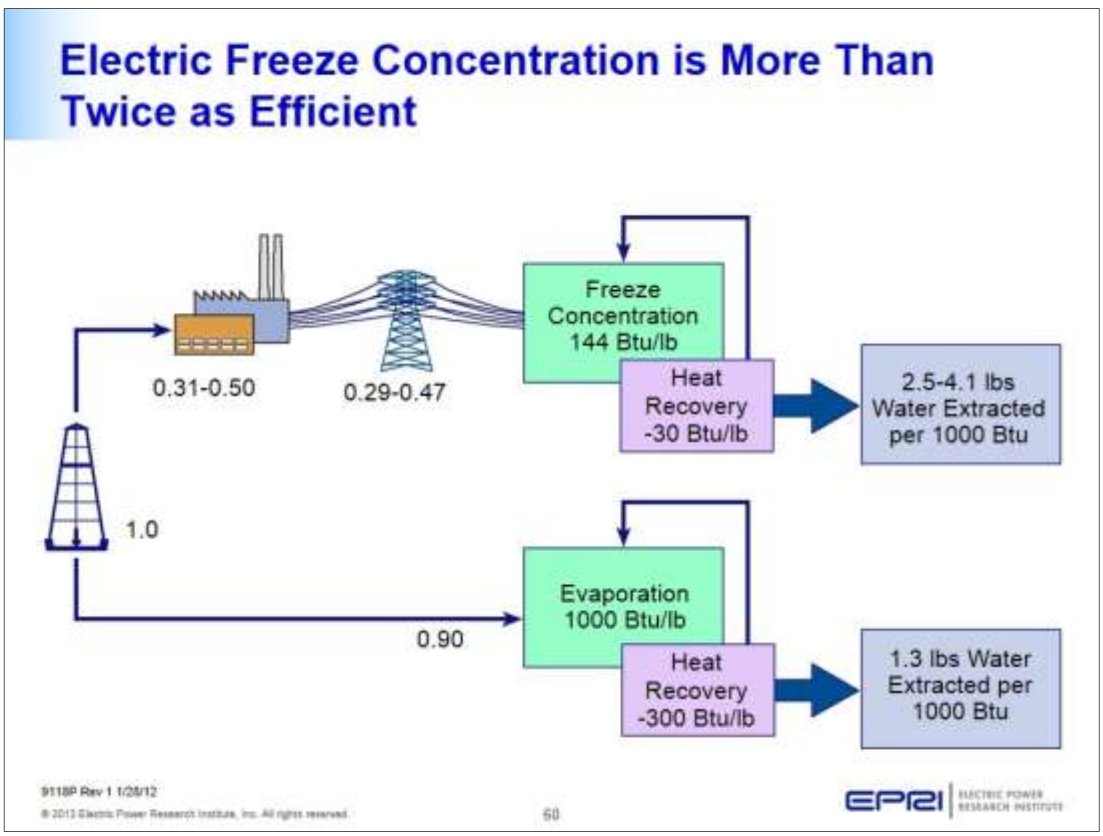
Electric Freeze Concentration is More Than Twice as Efficient
0.31-0.50
0.29-0.47
Freeze Concentration 144 Btu/lb
Heat Recovery -30 Btu/lb
2.5-4.1 lbs Water Extracted per 1000 Btu
1.0
0.90
Evaporation 1000 Btu/lb
Heat Recovery -300 Btu/lb
1.3 lbs Water Extracted per 1000 Btu
9118P Rev 1 1/28/12
© 2013 Electric Power Research Institute, Inc. All rights reserved.
60
EPRI ELECTRIC POWER RESEARCH INSTITUTE

That's why the smart grid! I hope this is useful as foundation to go on and talk about what the smart grid looks like. I think you know the equation that describes where end uses are going… increasingly more digital, less analog.

Thank you. I hope you found that of some value.

Keeping the Grid Working

DUSTY BECKER

Our next speaker is Dr. Khaled Abdul-Rahman. He received his BS and MS degrees in Electrical Engineering from Kuwait University. He received his Ph.D. in Electrical Engineering from Illinois Institute of Technology, Chicago, IL.

Dr. Abdul-Rahman is the Executive Director of Power Systems & Smart Grid Technology Development at California Independent System Operator. His current job responsibilities include overseeing software implementations of forward and real-time market applications; advanced grid applications, including online voltage and dynamic stability analysis; market and grid operator training simulation; integrated optimal generation; and transmission outage coordination. Dr. Abdul-Rahman's responsibilities include ISO smart grid projects and pilots, external technology collaboration, and external research support.

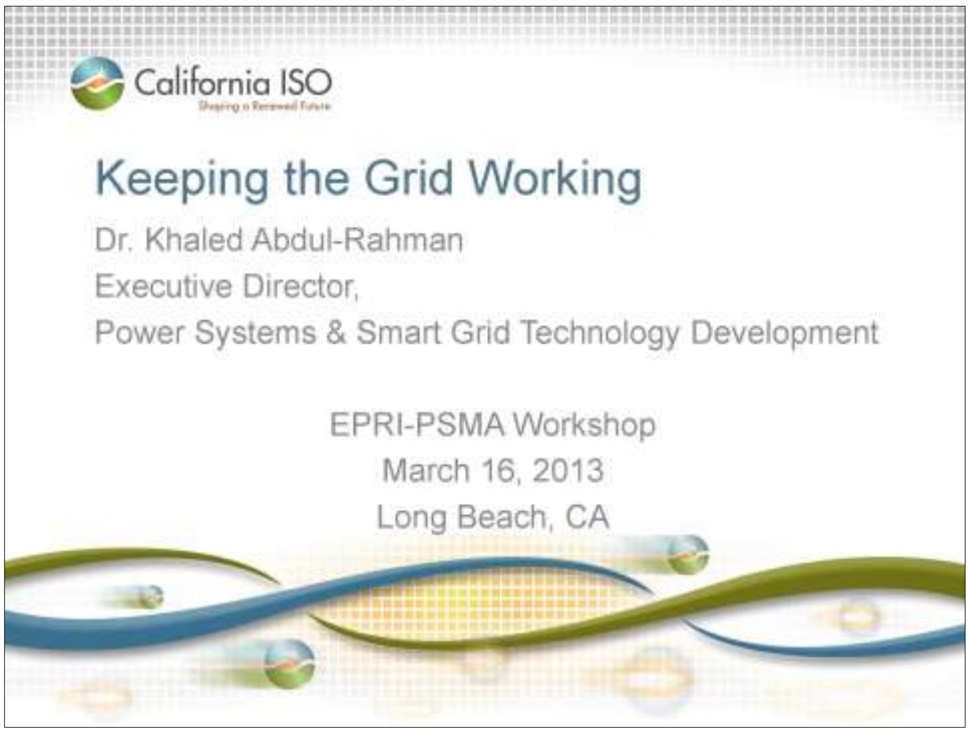

Dr. Abdul-Rahman has many publications in *IEEE Transactions on Power Systems* and has spoken on many technical panels and conferences related to power systems.

Please welcome,
Dr. Abdul-Rahman.

DR. KHALED ABDUL-RAHMAN

Good morning, everyone.

> *ISO's responsibility is to manage and move power from where it is produced to where it is consumed, in the least costly manner, while maintaining the reliability of the grid.*

Now that we have seen the first presentation and gotten an idea about what is the "grid," I would like to build on that and introduce some of the challenges that the Independent System Operator (ISO) faces while managing the power flow on the grid. The ISO's responsibility is more on the transmission side of the grid: How to manage and move the power from where it is produced to where it is consumed in the least costly manner, while maintaining the reliability of the grid.

I'll be talking more from an ISO perspective while touching base on some fundamentals and basics of electricity flow to give you an idea about some challenges and opportunities we see in the future grid. My presentation will give you the necessary basics that other presenters will build upon to explain issues related to maintaining voltage profiles and reactive power supply.

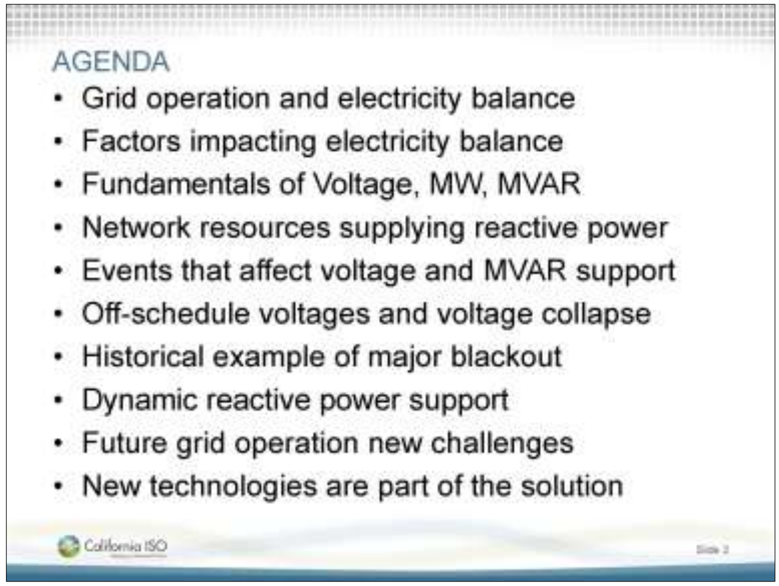

I will quickly cover the fundamental concepts: How the grid operates and how the ISO balances the electricity — the factors that are impacting the electricity balance.

Also, I will expand a little bit and zoom in on the voltage or reactive power concepts. The reactive power is a confusing concept and I would like to spend some time to clarify this concept and explain its importance for the grid. I will also list the major network resources supplying the reactive power; events that can impact the voltage, and the MVARs; The things that could happen if the voltage goes up. What happens when we have over voltages, under voltages, and what happens when the voltages get out of control and a voltage collapse occurs and the system starts breaking apart?

I'll give an example of one of the major blackouts and why we need dynamic reactive power support to help us resolve issues.

Then we'll talk about some of the future operations and new challenges and how the new technologies need to integrated and function as part of the solution rather than being part of the problem.

The grid operator must operate the grid so that no single contingency causes a disruption of service.

The grid operator is responsible for maintaining the reliability and the continuity of the service on the electricity grid at all times. In other words, we have to make sure that we operate the grid reliably and to do this, the grid operator must operate the grid so that no single contingency causes disruption. If we lose equipment in the transmission network; one piece of equipment, whether a transformer or a 500-kV line or 345-kV line... If we lose that equipment, the system must sustain that loss and will continue to operate normally. We will not have to lose or shed load.

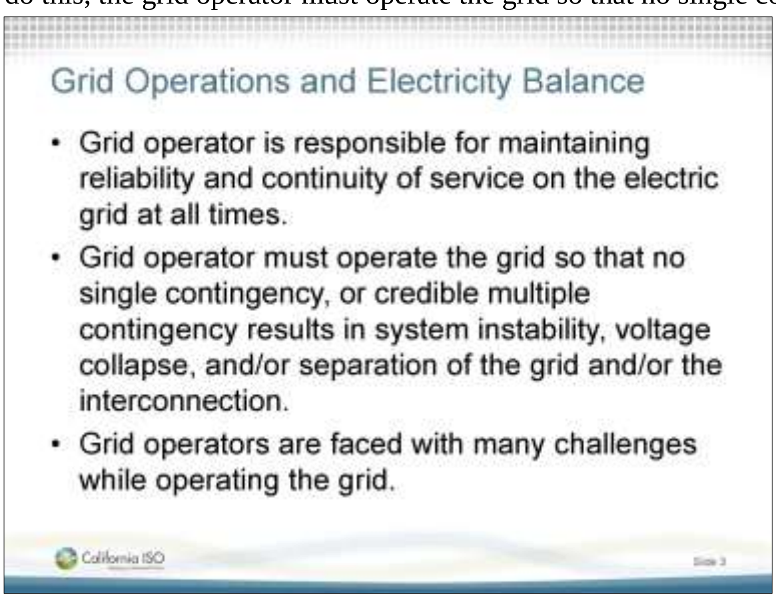

"Contingency" is any loss of equipment, lines, or generators, so we are trying to protect against it. Sometimes there are critical, multiple contingencies. It is very expensive to design a system to withstand double or triple contingencies all at the same time. Usually people design for one major contingency, and sometimes to withstand credible and critical multiple contingencies.

Contingencies are defined and grouped as a contingency list and we make sure that, when these occur, the system is still secure and able to provide the power to the load.

When we do engineering studies, we consider if there's branch overload or a bus bar voltage drop or increase beyond normal limits and these get reported. We try to adjust our base case (the normal case without any contingency), such that, when a contingency occurs, the voltage drops or overloads are prevented. While we are doing that, we face many challenges operating the grid.

A major function is to balance the supply and the demand.

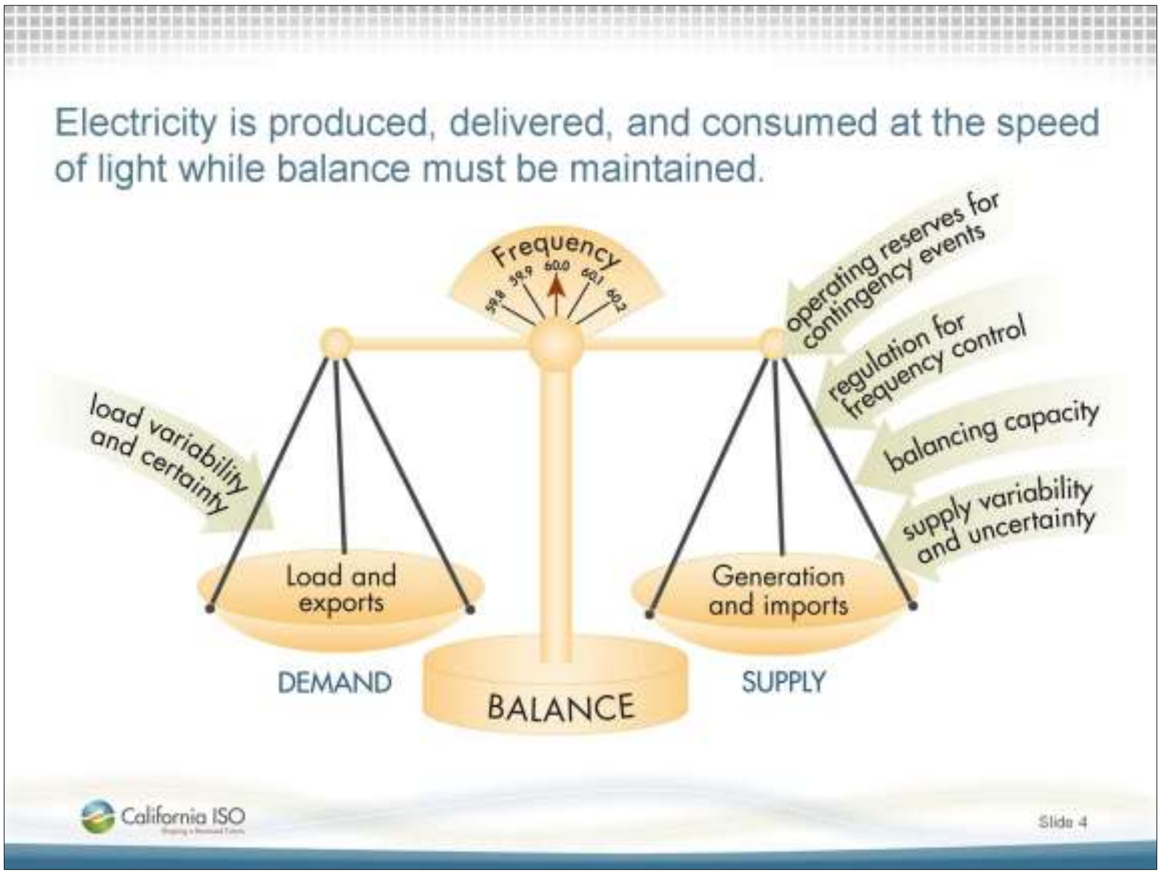

This diagram shows that the major function of the ISO is to balance the supply and demand while keeping the frequency of the system at 60 Hertz.

Supply means generation, plus whatever is imported from neighboring control areas.

The load is our demand, as well as the exports to the neighboring controls.

The load is changing as the customers are turning on and off equipment. The load is changing continuously. That is uncertainty we deal with on a continuous basis.

Generation is the other variable.

There are multiple challenges that the operator has to manage. We have to make sure that we not only have enough supply to meet the current load, but enough supply to meet the forecasted load. Some of the resources, the generators, for example, need time to start up before they begin generating. We have to take that into consideration so that when the forecasted load becomes a reality, there are enough online units to provide that supply with the power demand.

We have to have some reserve capacity, unloaded, waiting in case of a contingency.

I also talked about generation loss contingencies. We have to have some reserve capacity, unloaded, waiting in case of a contingency, like the major generator going offline. Then we can supply replacement capacity from the reserves. We call it operating reserve. There is spinning and non-spinning reserves, but let's just call it operating reserve.

The load is changing continuously and instantaneously. Because of that, we have a very complicated system of automatic generation control. If you remember from the previous presentation, there was a picture of the control center. The operators are watching those changes all the time. To maintain regulation, they keep track of these changes and move some of the generators that are fast responding. Those are generators that are controlled by AGC or regulation signals. We are controlling them from the control center directly. We send them a regulation signal to increase or decrease to follow all of these instantaneous changes of the load.

We have to make sure that we have enough capacity overall to meet our largest demand. We use the largest demand of the year in the planning studies. On top of that, we need some margin for the reserve. We need to make sure that the system has enough capacity to cover all that. That's part of the planning and resource allocation.

We have another challenge. We also have to forecast the supply. In the past, we were controlling the generator, so we knew that this generator was producing 100 megawatts or 200 megawatts. Now the supply itself can include variable energy resources, with more and more intermittent power sources connecting to the grid.

Not only do we need to forecast demand more accurately, but also we have to forecast solar and wind generation more accurately.

Not only we need to forecast demand more accurately, but also we have to forecast solar and wind generation more accurately. While we're doing all of this, we have to make sure that the supply and demand are balanced so that our frequency stays around 60 Hertz. If the generation exceeds the demand, the frequency will increase. If the demand increases, the frequency will decrease.

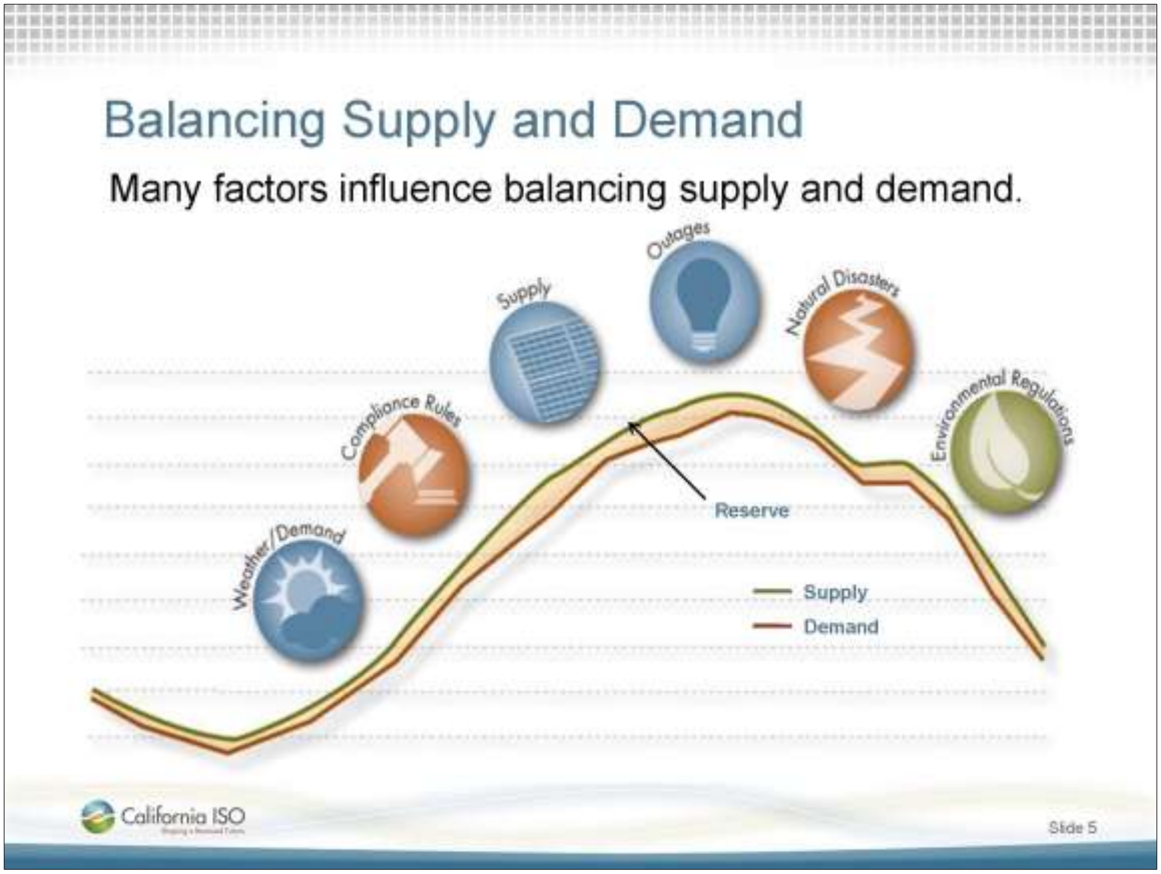

In this slide, the red line is the supply that we need to meet the varying demand. The green line is the supply that we need to have ready; the demand plus some margin is called reserve capacity.

There are many factors that influence the supply-demand balance. The weather is one, because the demand is typically conforming loads impacted by weather conditions. The weather also impacts variable generation, e.g. clouds cause rapid changes in the generation of solar panels.

The ISO operates the transmission grid, and must comply with applicable reliability and regulatory rules.

The ISO operates the transmission grid and must comply with applicable reliability and regulatory rules. We are regulated by the Federal Electricity Regulatory Commission (FERC). We are also regulated by the state of California. Our board of governors is appointed by the state. We also have a lot of regulations related to the environment, clean power act and others.

The availability of supply and transmission outages that occur in the networks both impact the supply-demand balance.

Natural disaster is another factor impacting the balance. If we have fires during the summer, for example, a fire could get close to major transmission lines. Sometimes we open a transmission line in case a fire cannot be controlled. If it gets very close to major equipment, we shut down that equipment and we operate without it.

The AC terminal voltage is really the driving force that causes power to flow in the electric grid.

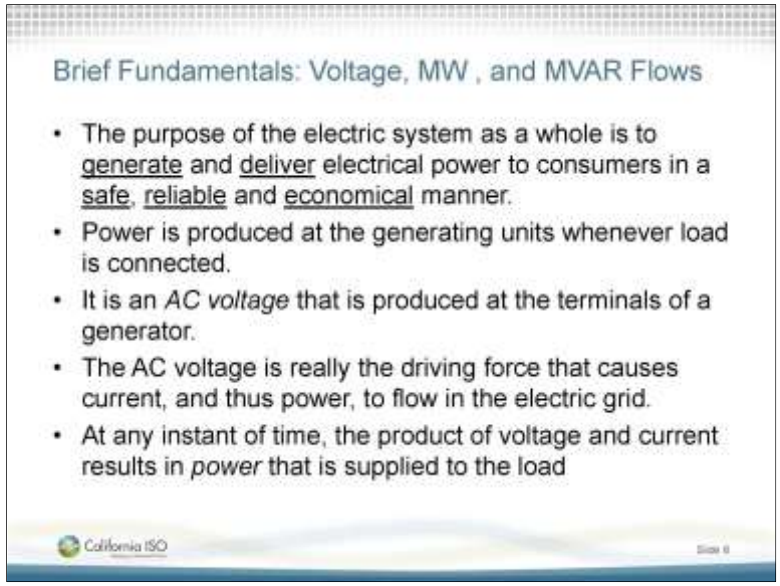

There are many topics that we can discuss for how the grid operates, but for this workshop, I will focus more on voltage, megawatts, and megaVARs; and how, from a transmission point of view, these are important to recognize. We can reflect on how some of the new technology can help resolve some issues.

Basically, our goal is to manage flows resulting from supplying electrical power and deliver it to consumers in a safe, reliable, and economic manner.

Power is produced at generating units whenever a load is connected. Power will not flow in an open circuit. The ac voltage produced at the terminal of the generator actually makes power flow. The ac voltage is really the driving force that causes the current, and the power, to flow to the grid. The product of the voltage times the current is what we call "power" — that is what is supplied to the load.

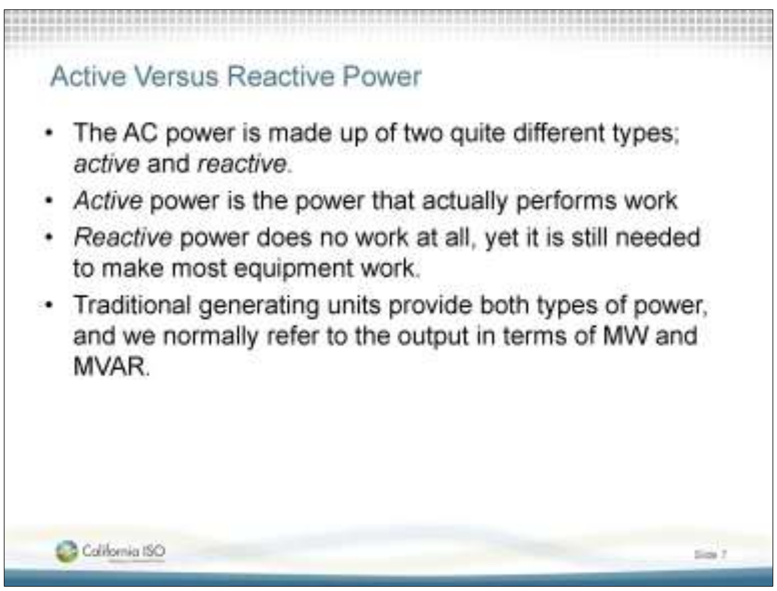

There are two distinct and completely different concepts for power. One is called active power or megawatt or watts. I tend to say megawatts because I come from the transmission side of the grid. A megawatt is one million watts. A megaVAR is a million VARs, which is the reactive power unit of measure. So there is the active and the reactive portions of power. The active power is the useful power. It's what actually does the work for us. When the motor rotates, that is active power. When the fan starts to twirl, the power that makes it work is the active power.

The active power is the useful power. It's what actually does the work for us. The reactive power does no work at all. Yet it is still needed to make most of our equipment work.

The reactive power does no work at all. Yet it is still needed to make most of the equipment work. There is no useful work that is a result of the reactive power, but most of equipment cannot operate without reactive power. Traditional or conventional generating units provide both megawatt and megaVARs.

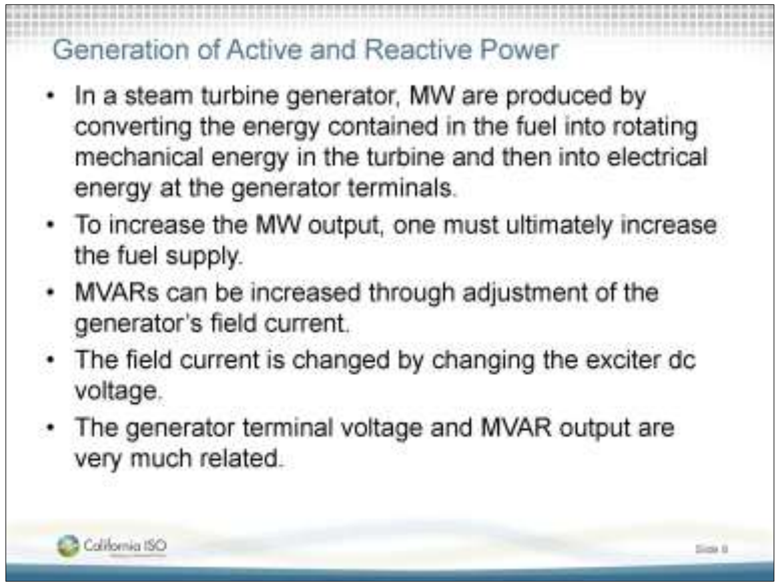

As part of the interconnection agreements with suppliers and generators, we define some requirements for how much reactive power that the connected resource needs to provide when they connect to the network.

How do we generate the active power? Consider the example of a steam generator. Megawatts are produced by converting energy from burning fuel. The energy in the fuel is converted for our use by taking back the mechanical energy in the turbine shaft, then into electrical energy at the generator terminal. Even with solar or wind, it is still using a type of fuel… in this case, a clean fuel.

If we want to increase the output of the megawatt, we have to burn more fuel. The megaVARs can be increased without any impact on the fuel burned. We have to increase the fuel, whatever that fuel is, to get more megawatts. But the megaVARs, we can increase without even changing the fuel we're burning.

Traditional generators produce megawatts and megaVARs.

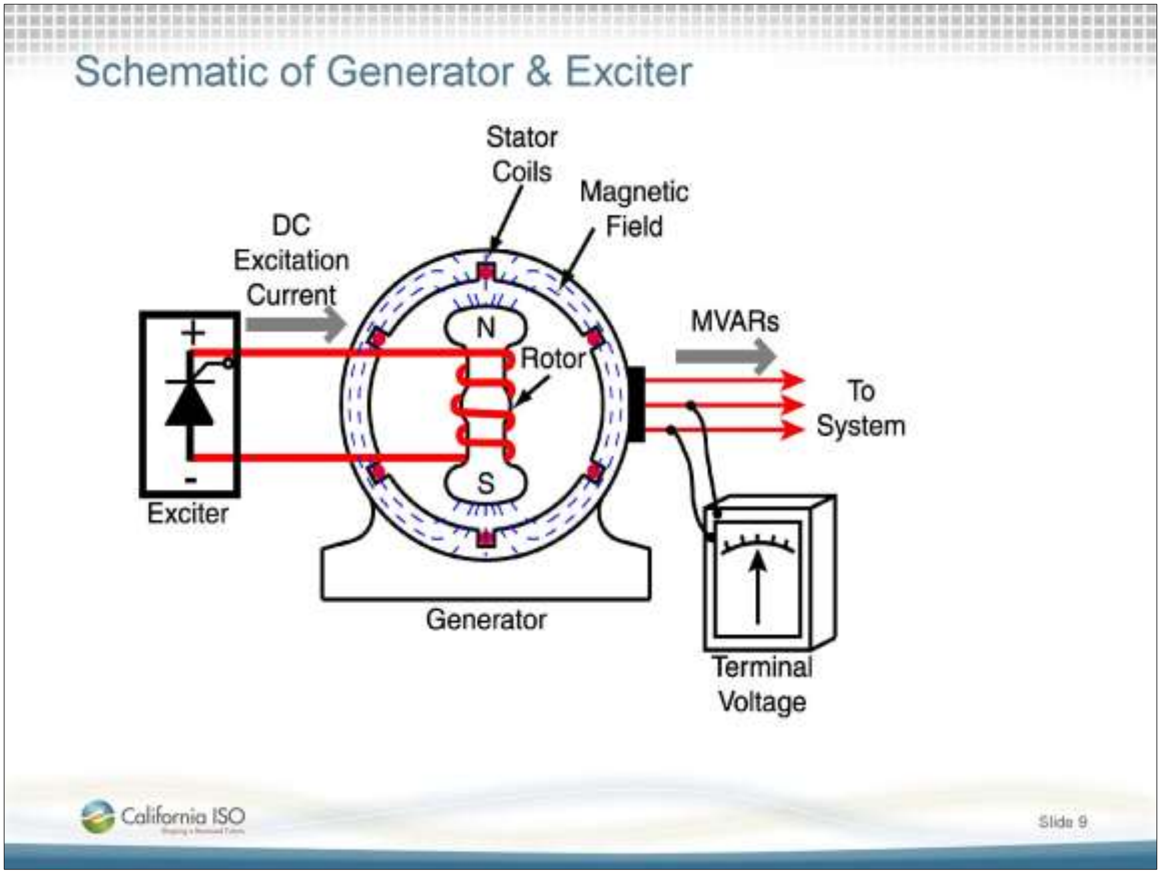

Traditional generators produce megawatts and megaVARs. The way a basic generator works is that there are two components. One is called the stator and one is the rotor. One is magnet and one is coils or wires. The whole idea is that we create a magnet by a dc source, putting dc current through the rotor to create an electromagnet. It's a very powerful magnet! Those red dots in the slide are the stator and represent the coil or wires. The result of the rotating magnetic field and crossing wires is a voltage that will be induced in these wires. That's the fundamental of magnetic induction in simplistic terms.

If we connect these wires of the stator to a load, then the electric current will flow, and we start supplying power to that load. The slide is showing three phase connections.

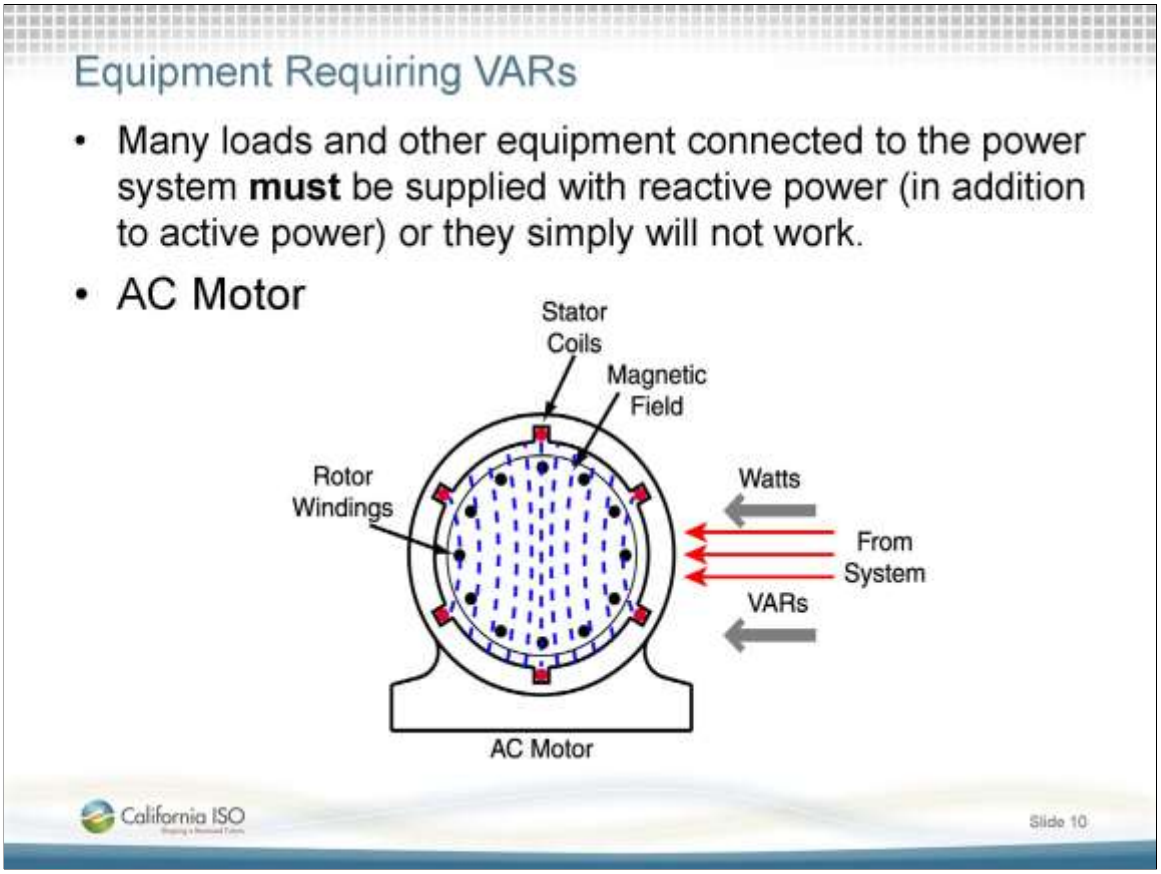

Much of our equipment requires megaVARs to operate. For example, the ac motor, receives electrical energy and converts it into the mechanical energy. When it receives the watts and VARs, the watts are the ones that are responsible for turning the motor.

How is that done? We are giving the electricity to the stator and the work is done on the rotor side. There is an air gap or a gap between these two components. How does the work get done here if there's no direct connection between these two components? The megaVARs actually are creating the magnetic field that crosses the gap. It's like a bridge in the air gap of the motor that allows the watts to be transferred into the rotor. Then the rotor can create mechanical power.

As you see, the VARs don't do work in terms of rotating the motor, but are a means by which the watts are moved. You cannot move the watts across the air gap without the VARs.

Transformers have primary and secondary windings. We provide the electric power to the primary side. The secondary side is not actually physically connected to the primary side, but a voltage appears on the secondary side. The wires are connected from the secondary side to the loads. Power flows from the source into the load. The transformer helps transfer the voltage from one level to another. The transformer can perform its function without the megaVARs. The megaVARs energize the coil of the transformer, magnetize it, and create a rotating magnetic field. That, in turn, induces voltage on the secondary winding.

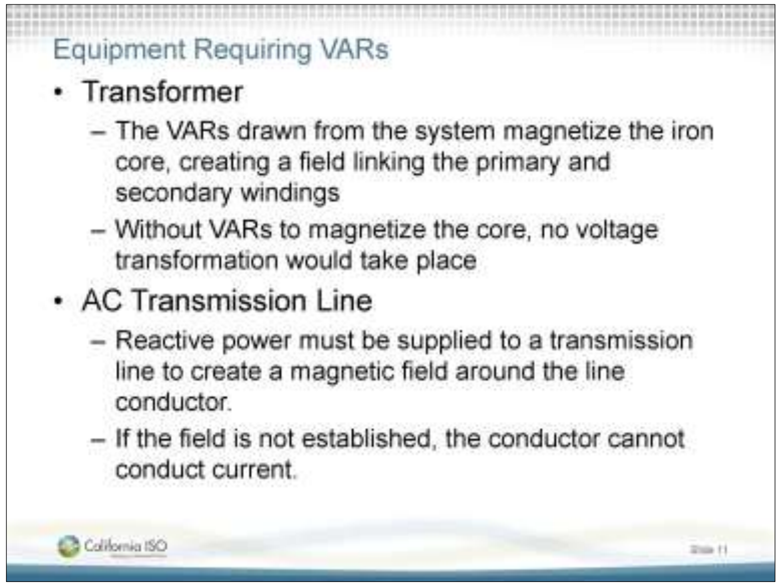

The megaVARs energize the coil of the transformer, magnetize it, and create a rotating magnetic field, which, in turn, induces voltage on the secondary winding.

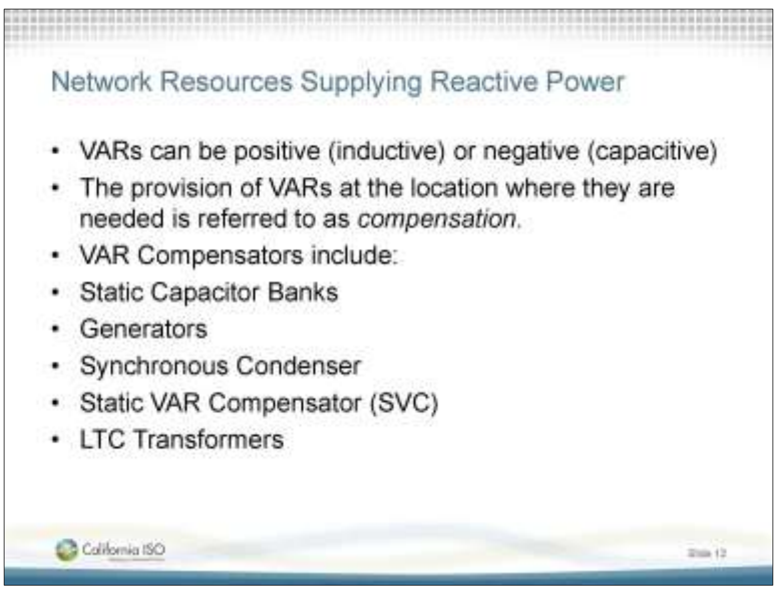

The voltage depends on the number of winding turns of the primary and the secondary windings. There are step-up transformers to increases the voltage levels and step-down transformers to decrease the voltage levels. That’s the basic concept of the transformer. Without the megaVARs, we would not have transformers.

Even the ac transmission line itself needs reactive power because it cannot conduct electricity without absorbing some reactive power in the transmission line itself. We have to create that magnetic field around the conductor to move the power along the transmission line.

Many have problems grasping the concept of megaVARs. Think of it this way: If I want to move a load of sand using a wheelbarrow. The first action is to lift the handles of the wheelbarrow off the ground. Then I exert horizontal force to move it from one place to the other. Lifting the handles of the wheelbarrow off the ground is like reactive power. When I lift, I do not transfer that load from one place to another. Then I push horizontal to move that load from one place to another more easily. Without lifting the handles of the wheelbarrow off the ground, i.e., without the reactive power, I couldn’t do the job. The movement of the load, the horizontal force, provides the active power of moving. Lifting the handles is necessary to perform the job; just as the reactive power is necessary to support the movement of the active power.

Reactive power is necessary to support the movement of the active power.

By convention, VARs can be positive or negative. If a device is absorbing reactive power, as in an inductive load, it is positive. If generating VARs, as in a capacitor, it is negative. VARs can be positive or negative. If a device is consuming VARs, i.e., absorbing the reactive power, as in an inductive load, it is positive because the device is absorbing the active power. If you are generating VARs, generating reactive power, as in a capacitor, it is negative because you are generating reactive power.

What is the transmission grid equipment? I think most of us know, but typical equipment in a network might be compensators, static capacitors, generators, synchronous condensers, and static VARs compensators; all these equipment help supply local reactive power needs.

VARs can be positive or negative. If a device is absorbing reactive power, as in an inductive load, it is positive. If generating VARs, as in a capacitor, it is negative.

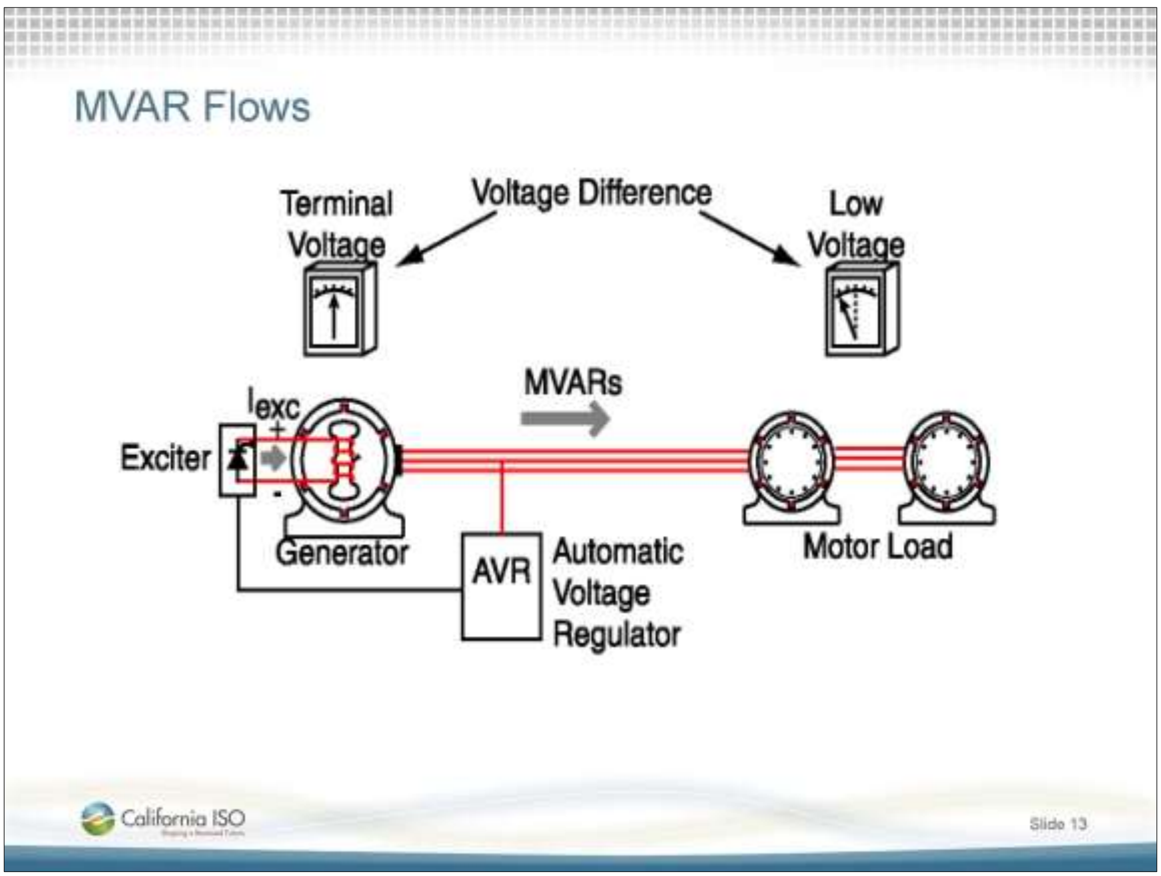

How do the megaVARs flow? From high voltage to low voltage. If we have a load here, which is a typical load model, for example, in our houses; then we need reactive power for this to work. It's absorbing reactive power, so it is inductive, or positive. The megaVARs that go to that load need to have a voltage difference to move. There has to be a voltage difference between the voltage here at the point where we want that megaVARs to go and the point of supply. The more voltage difference between the load and the source of supply, the more megaVARs go there.

Let's say that we had only one model here instead of the two, so don't have to consume that many megaVARs. If it's only one motor, we don't have to consume that much, so there will be excess of megaVARs. The automatic voltage regulator senses a voltage difference because there are no megaVARs consumed. The voltage at that terminal increases because there are more megaVARs than our load is using. The automatic voltage regulator senses that the voltage here starts to increase and it sends a signal to the exciter, the one that produces the dc current here, that creates a magnetic field. It sends a signal here to reduce the excitation current.

By reducing the excitation current, it reduces the voltage across the terminal of the generator. By reducing the generator supply and reactive power, the voltage is reduced. As a result, the megaVARs we are sending are reduced to the value needed at the load. If we have two loads, we have more reactive power to consume. If we're not sending enough, the voltage drops. This signals to the exciter to increase the current and it generates more reactive power. The voltage increases, then we are sending more reactive power. There is a very tight relationship between voltage and the reactive power.

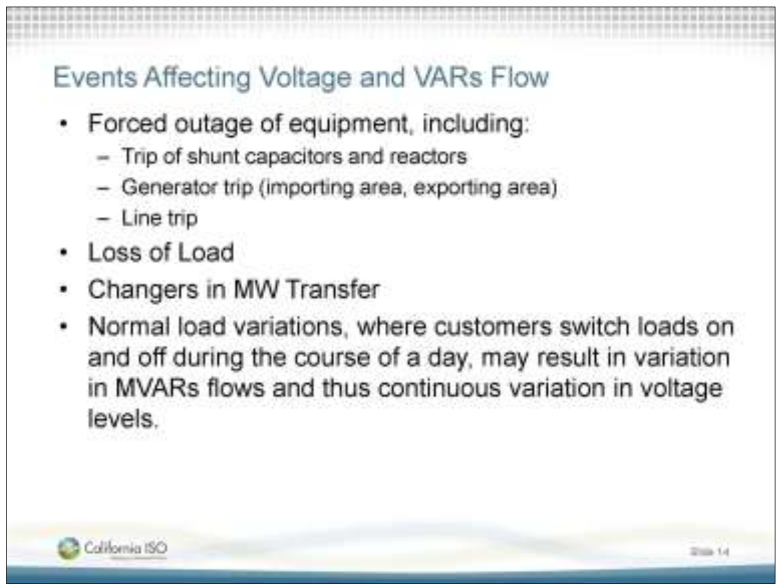

In a normal situation, where everything is settled, there is a flow of the megaVARs reactive power throughout the network. There are certain factors that can impact that flow and the voltage at certain terminals or certain points in the network. Those usually are some capacitors or some supplier of the reactive power trips out; then the voltage goes down because they were supplying reactive power. Without it, the voltage goes down.

Importing and exporting. Let's say we're importing some megawatts and, along with that, reactive power to make the megawatts move. If we lose that input, the megawatts, there is a lot of reactive power lost as well. There will be demand in the area that misses that imported power that was cut. The reactive power that is required doesn't have something to consume. The voltage will go down because we lost that input. We lost the source of reactive power, not only the megawatts, but also the source of reactive power.

The same thing happens when we're exporting on a tie line, from one area to another, and the export is reduced or cut. There is plenty of reactive power to send. The voltage will increase instead of decrease.

That's why voltage and reactive power are important in determining the amount of megawatts we can transfer, the amount of reactive power that we can transfer.

I sometimes use the exchange of active power and real power; it is the same and measured in watts.

The voltage should be within safe margins. We don't want more that 1% drop. Up to 5% is the most that the standards allow with a contingency.

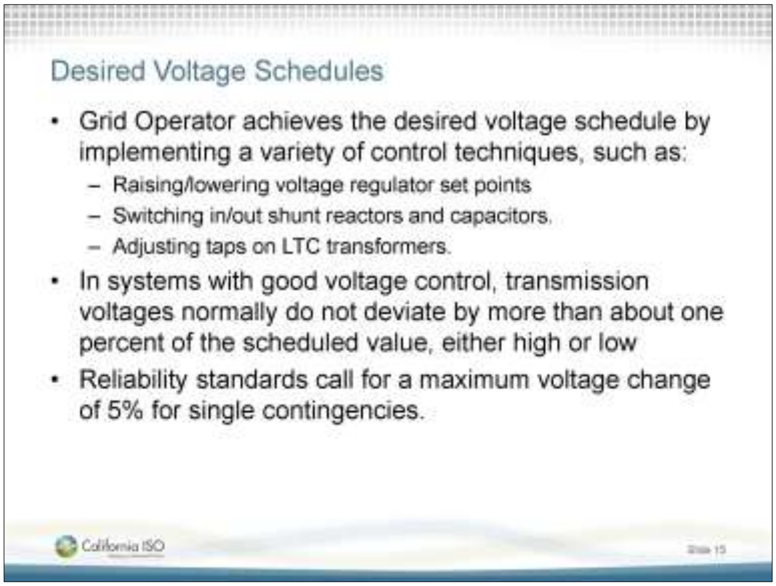

The voltage should be within safe margins. We don't want it to drop more than 1%. Up to 5% drop is the most the standards allow, even with a contingency. It must be within 5%. So we bring it within those margins.

The provision of reactive resources: We are procuring or supplying the reactive power. When we supply reactive power, the equipment that supplies it, is called a compensator. The most important compensator from an operational perspective is the automatic voltage regulation, the AVR. These AVR are usually synchronous machines. A

As we talk about new technologies, we should not forget about the basic needs of the grid.

As we talk about new technologies, we should not forget about the basic needs of the grid. This is one of the things that we need to make sure we don't forget because people usually talk about megawatts replacement. I'm replacing one megawatt with another megawatt (one megawatt of conventional units with one megawatt of renewable units), but there is also reactive power to worry about.

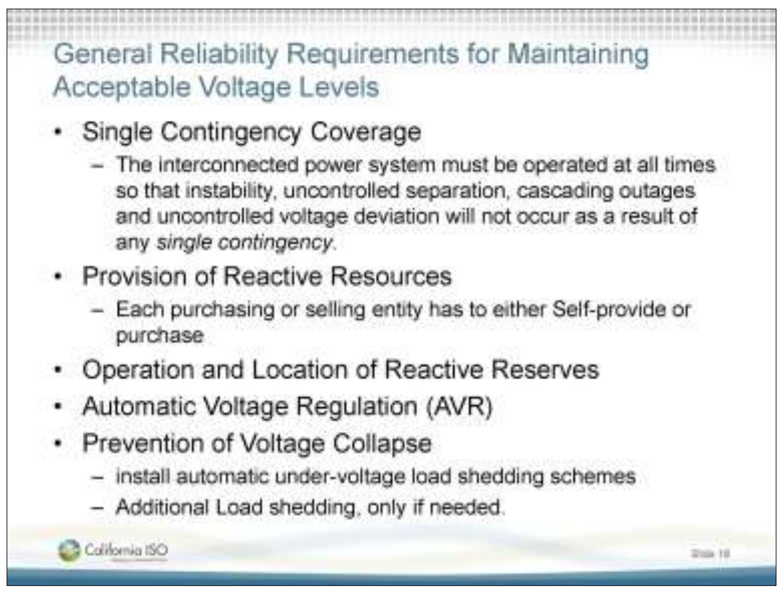

When replacing one megawatt with another megawatt, there is reactive power to worry about.

How do we prevent the cascading voltage collapse and the cascading isolation of certain parts of our network? This is usually a result of losing the voltage or drop of the voltage. We have to make sure that we have enough reactive power sources in our grid to be able to satisfy such a situation.

Over generation or under generation is problematic for the grid. I am not going to go into details, but both are bad because both can create either high voltage or low voltage. Both can overheat the equipment, then some of the equipment will trip. When a generator trips and we have excessive load, the frequency starts swinging, some of the protection devices start opening circuits, and then we eventually start isolating parts of the grid from each other.

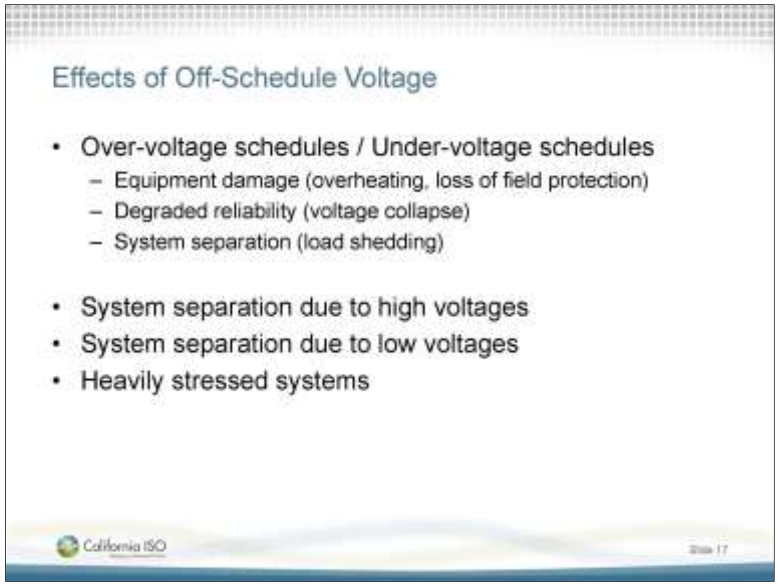

Over generation or under generation is problematic for the grid.

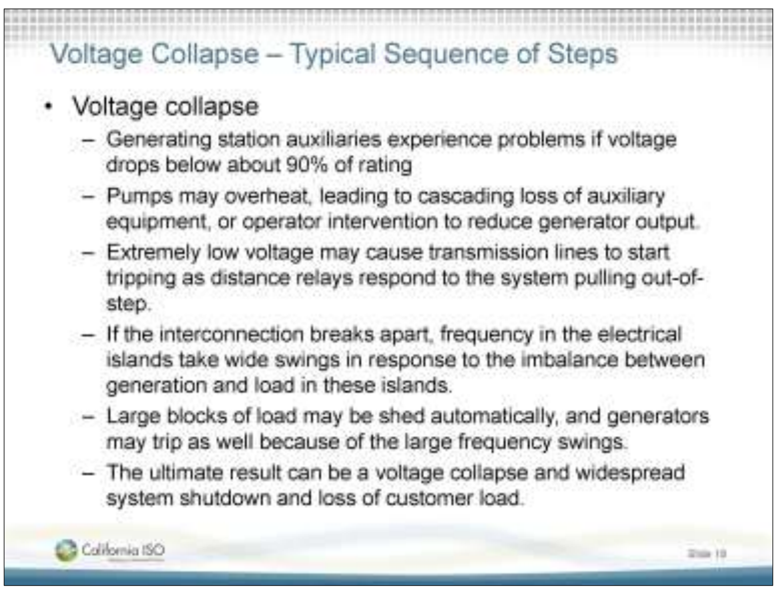

I'm going to go into one example quickly. In 1996, there was a blackout because a major 345,000-volt line tripped due to tree contact. If this was the only event, the system would have survived. The grid is designed so one contingency won't bring it down.

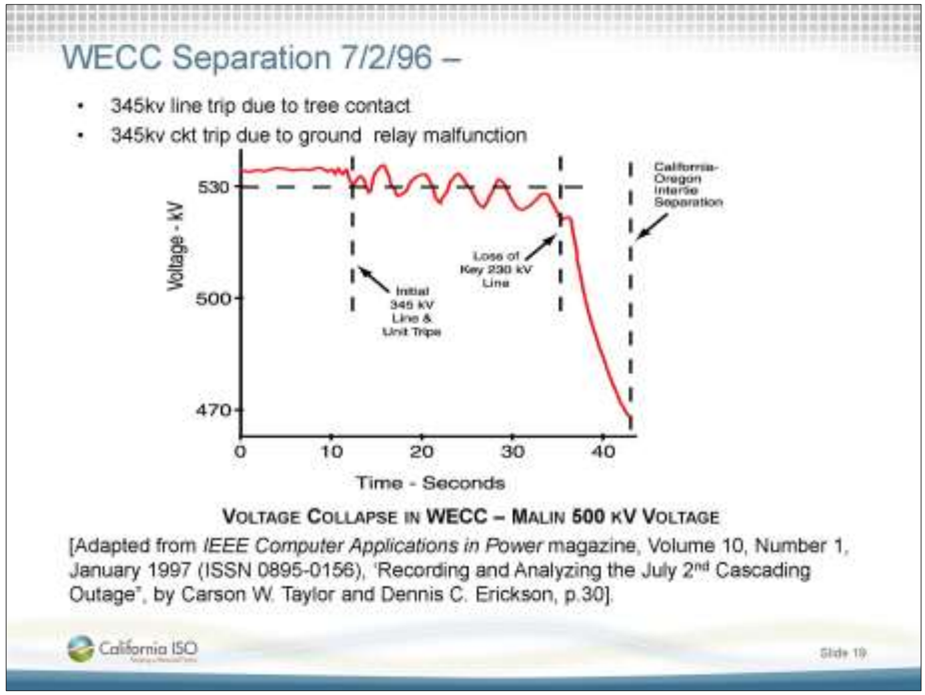

Unfortunately, at the same time, there was another 345-kV circuit that tripped due to a relay malfunction. They were close to each other. After that, the rest of the circuit overloaded, so more current was going into whatever was left from the connection from the particular plant that was generating the megawatts. So the rest of the circuit was overloaded. Then those circuits tripped. Then the generators began tripping because there was an open circuit.

However, the load was still on the network and someone has to supply that load. So inputs from the north started increasing and the network was already in a heavy-demand situation. The circuit was almost, but not quite, overloaded. By transmitting all of that power to that part and trying to get all of that power from there, the circuit was overloaded again. That circuit tripped off and then the whole northwest tie tripped off. California was separated from the northwest when the tie tripped off.

The whole event took almost 55 seconds. It resulted in five different isolated systems, islands. The restoration, bringing service back and connecting those pieces together again, took probably 51 minutes.

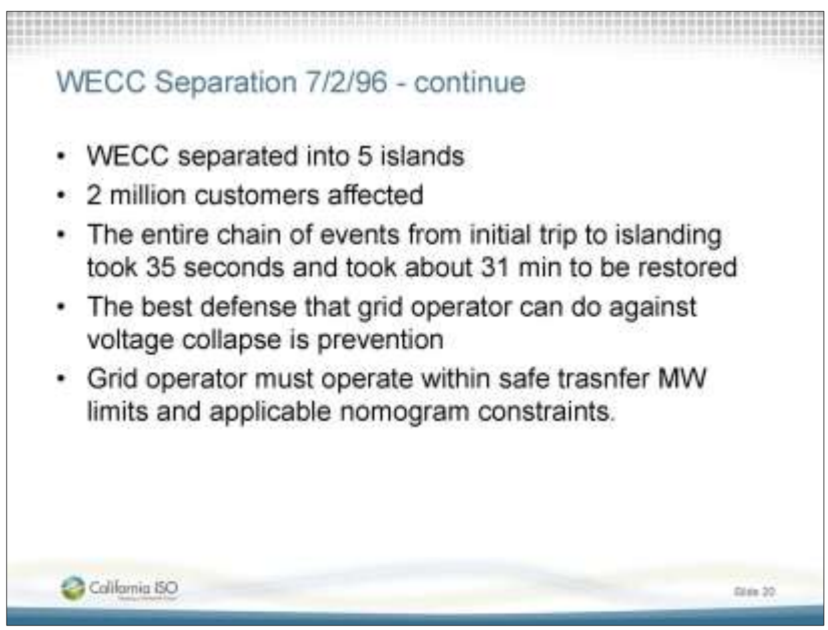

As a result, we had the inconvenience, the amount of load was shed and tripped off, and the blackout occurred. Fortunately, it didn't last long.

Preventing such disturbances is the highest priority for operation. We don't want it to happen at all, so we try to make sure that the grid is operated within safe transfer megawatt limits and respects all network constraints and reliability standards.

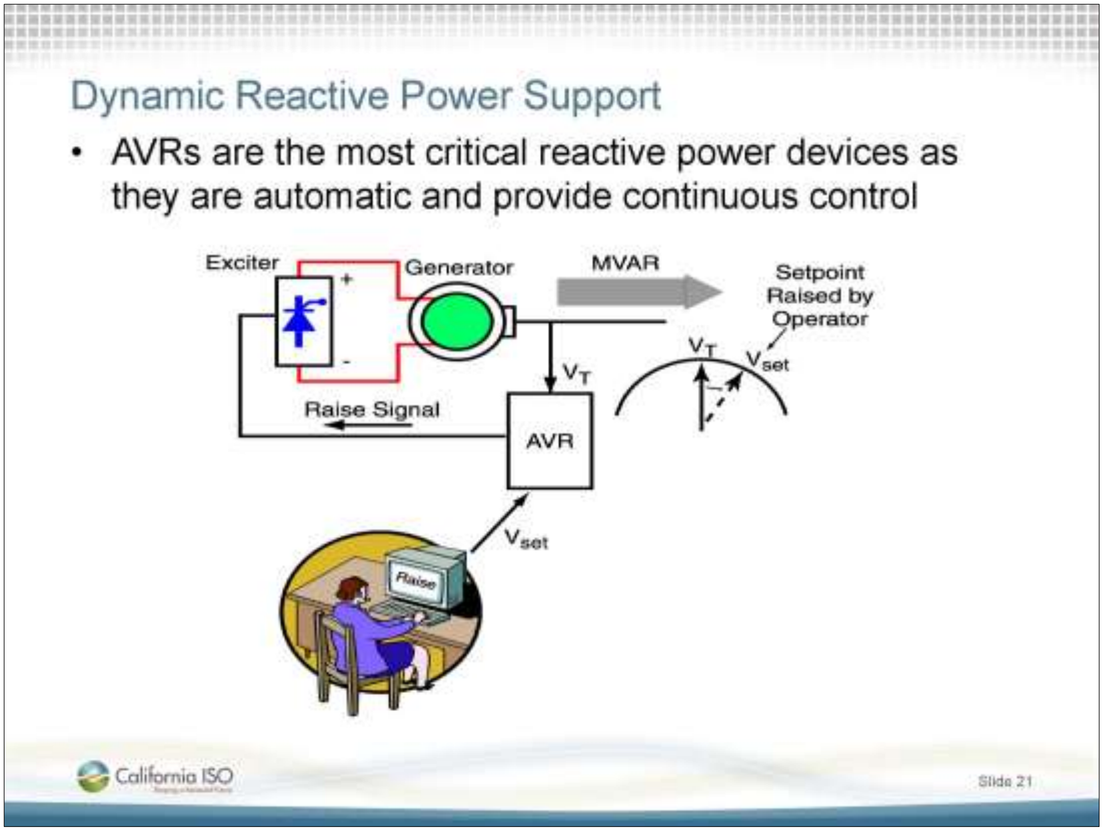

This slide illustrates the improtance of the continuous and fast reactive power control that conventional unit AVR is providing to the grid. The megawatt replacment of these conventional units by renewable capacity should consider the need to ensure reactive power is supplied by the replacement capacity or some other devices connected to the grid.

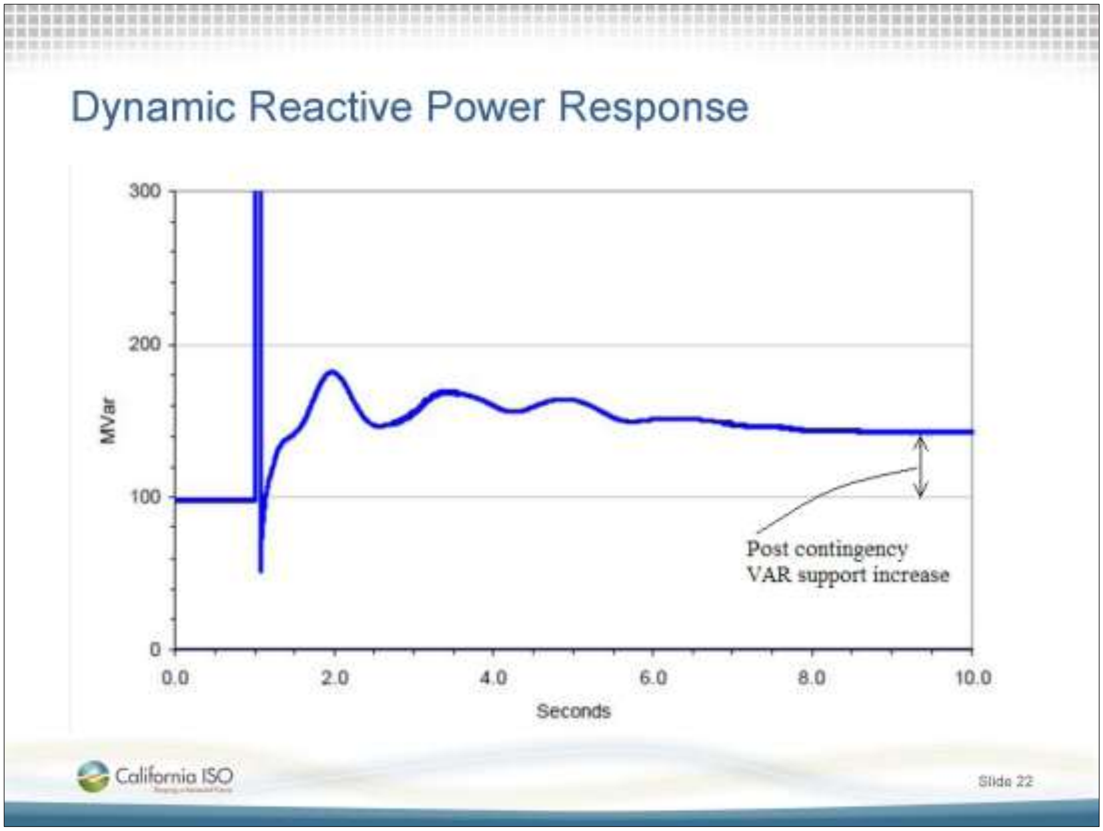

Dynamic reactive power responds quickly.

Dynamic reactive power. I'm stressing the importance of AVR because it normally comes with conventional units and it responds quickly to disturbances on the grid. The blackout of the whole west happened in 35 seconds. There is no time for a static device to switch on or off — that usually takes *more* time. We want something that is automatically and continuously responding to changes in reactive power demand. You see on this slide that a contingency occurred at 1 second. Then that device responded by providing more megaVARs quickly and it stabilized the system. So before, it was providing 100 megaVARs; after, it was providing close to 150 megaVAR. You can see here that, when the contingency occurred, this device was providing 300 megaVARs instantaneously. Then it settled at 150 megaVARs. That's a very important characteristic we need to maintain in our system to be able to survive sudden changes in system conditions.

In the future, we see that much of the generation resource mix is changing from traditional generators to solar and wind. We have to prepare our system to deal with this high penetration of renewable resources because they will replace traditional resources. We must ensure that whatever services traditional resources are providing to the system, these services must be covered by whatever resources that we have in our fleet.

We are seeing that much of the generation resource mix is changing from the conventional generators to solar and wind resources.

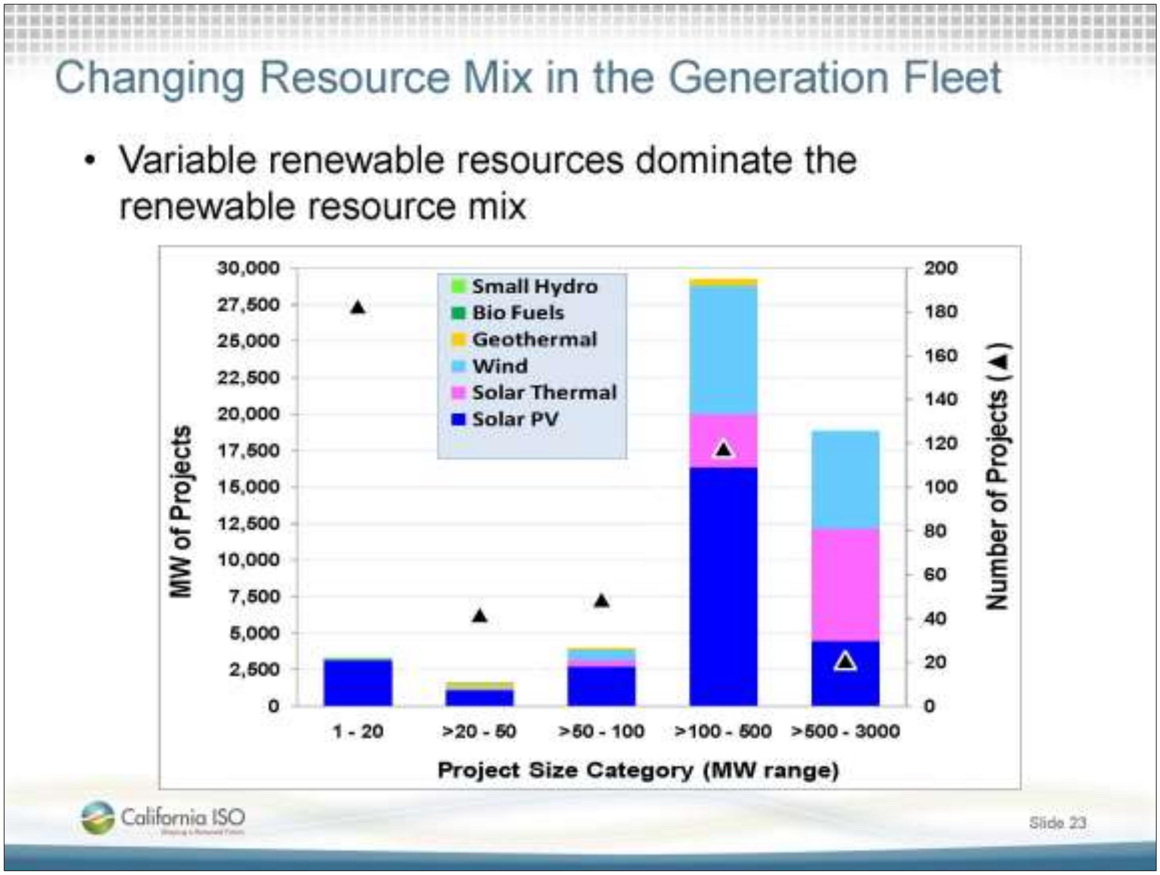

This is important for the grid and for the owners of these new non-conventional resources because they obviously want the system to continue functioning safely and reliably. No one wants a megawatt supply of renewable resources to be interrupted because there is a voltage problem, and grid operator has to open the breaker and interrupt flow of power. We have to worry about how to provide these ancillary services to make business of supplying power continuous and reliable.

November 24, 2011 saw some interesting wind events. The wind, at 10:00 in the morning, was blowing 308 megawatts. It jumped 385 megawatts during seven hours. It peaked at 1972 megawatts here. Then there was a drop from 1972 to 1386, then a sudden drop from 1386 to 66 megawatts in four hours… about 1520 megawatts back down within four hours.

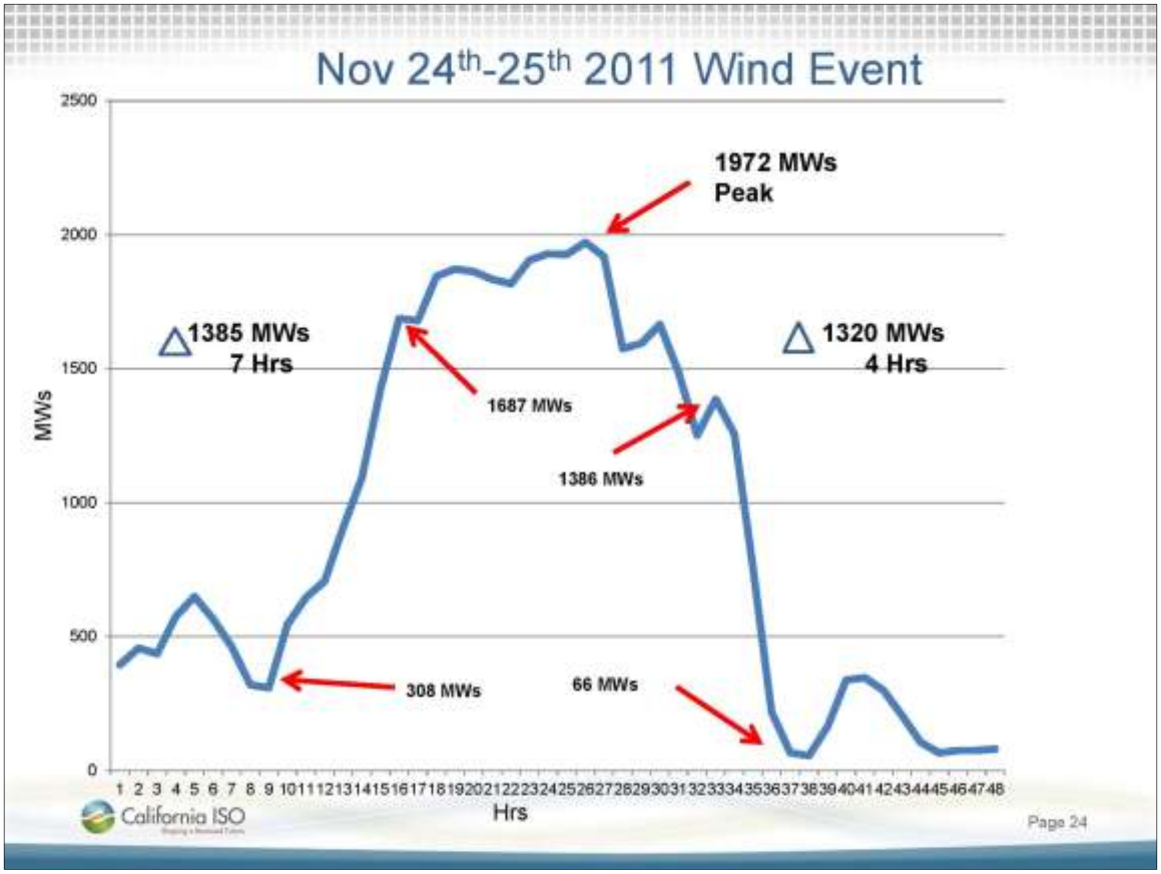

In the beginning of my presentation, I said one of our biggest challenges is to keep the balance between supply and demand. If we don't balance them, the system frequency will start swinging and then portions of the system start breaking apart. It is important to have capability in our system to balance out fluctuations like this.

This is generation variation and uncertainty; this is not a load variation. If the supply drops 1520 megawatts in four hours, we have to have capability to increase our generation elsewhere because the load is the same. The load didn't change in this case: the supply is changing. We have to have some other supply compensate for that drop in wind generation. The same is for solar variations and intermittency, we must provide the flexibility and ramping capabilities to survive such events and balance out supply and demand at all times.

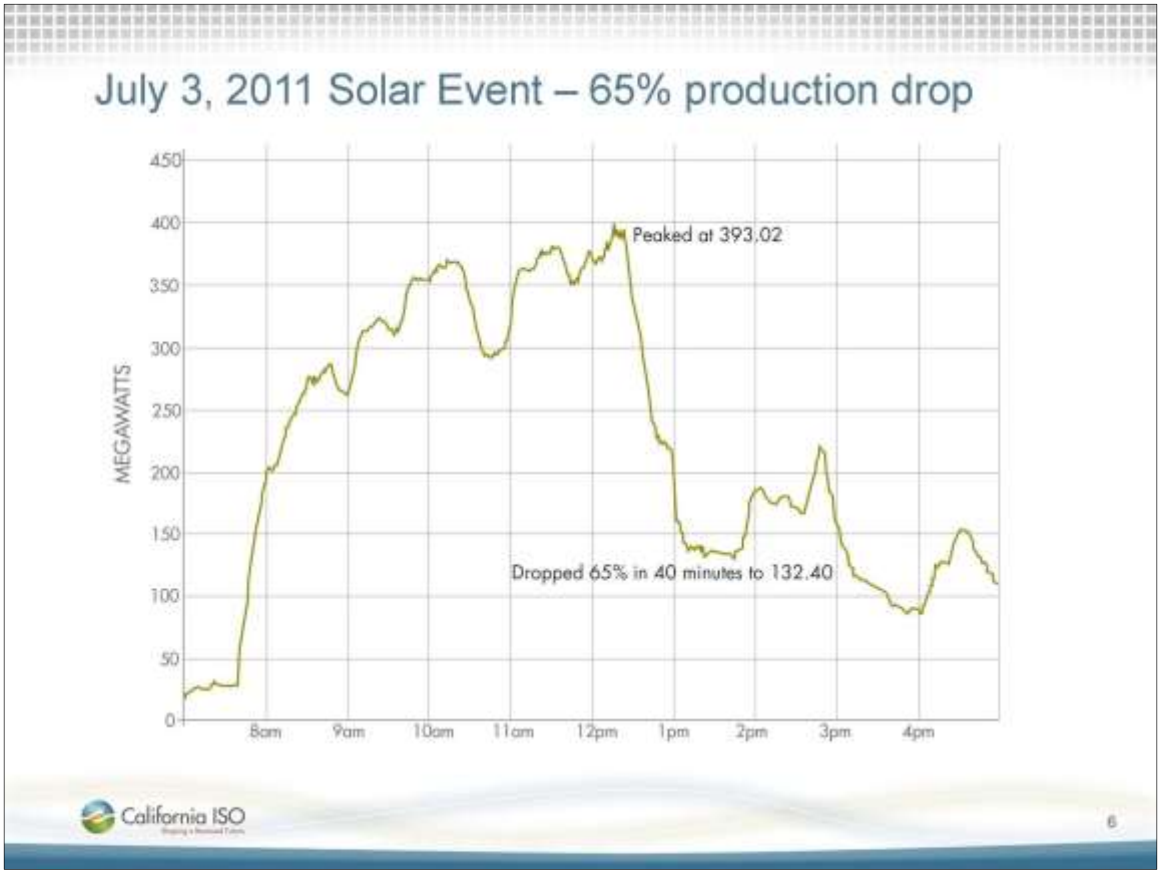

The same thing happened with solar. We had an event in July, 2011. It was about 400 megawatts. [By the way, we recently crossed 2000megawatt solar production in California. It was 400, then it dropped 65% in 40 minutes. We have to make sure that, in 40 minutes, we have the ramping capability in the system to maintain balance.

We have challenges that some of our conventional units are scheduled to be retired.

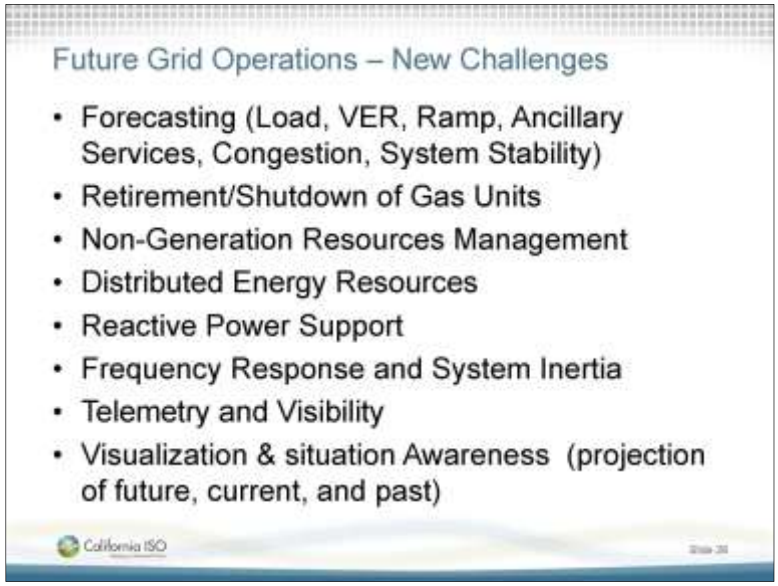

So what are the future challenges? Basically, the forecast. We have to continue improving forecast to make sure we have enough online capacity to meet the most variation of supply and demand. We have to forecast our load. We have to forecast the generation of the variable energy resources. We have to forecast the ramping megawatt and time of ramping that will happen. We have to forecast how much operating reserve we need. We have to forecast the congestion, and when it will occur, and any system condition changes that may result in system instability.

We have the challenge of some of our traditional units scheduled to be retired. Some of those are providing quick-response services. The ramping is actually provided by those units, so if we take these out, who or what is going to provide the ramping capabilities to follow demand and supply changes?

The impact of Distributed energy Resources (DER) on demand forecast has started to show. We see the impact of them directly, but most of it is behind the meter. Our domain is the transmission level, but these DER are mainly on the distribution side. If a lot of them are generating, we see it reflected in our load. We see that load is decreased when they generate. When clouds pass by or it rains, we see our load increases because the generation from them decreases.

Frequency response and system inertia are important.

An important issue for us is frequency response and system inertia. To withstand any shock to the system, such as losing one the major components, we must have some inertia in the system. Usually, it is provided by the big traditional rotating generators. Wind power usually provides a variable-frequency ac. Solar is dc. So we need power converters to supply one frequency, 60 hertz, to the grid. We need power electronics to take that variable frequency ac and dc into an ac at 60 hertz. That’s very important.

Telemetry and visibility. It is very important for us to have visibility at different levels. There are companies developing databases with the information taken directly from the inverter of installed solar panels and store data in a database and sharing this information with the system operators. We have to provide visualization to our operators to enable them to operate the system more efficiently. We have to show them what happened, is just happening, and likely to happen.

Grid Reliability Requires Increased Ramping Flexibility

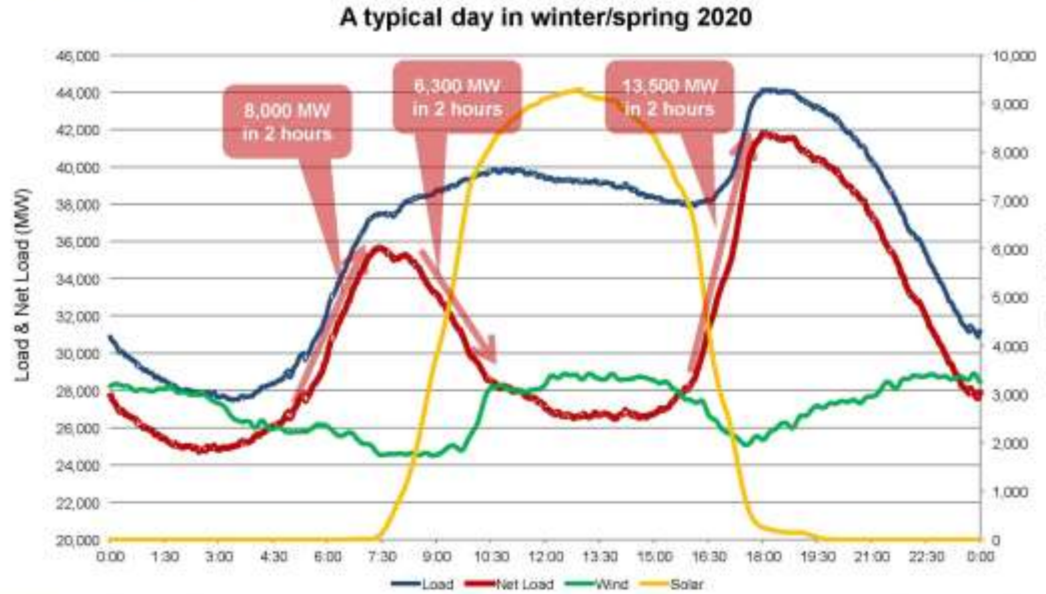

California ISO

Reliable Grid Operations - New Technologies Create New Opportunities

- Demand Response refers to all methods in which demand resources can react in ways more integral to the needs of the power grid (faster, more frequent, bi-directional)
- Demand response are already participating in operating reserves!
- New and more intelligent technologies enable loads to participate in many more forms:
  - Regulation Services
  - Flexible Ramp
  - Localized DR
  - T & D upgrade deferral

California ISO

Slide 28

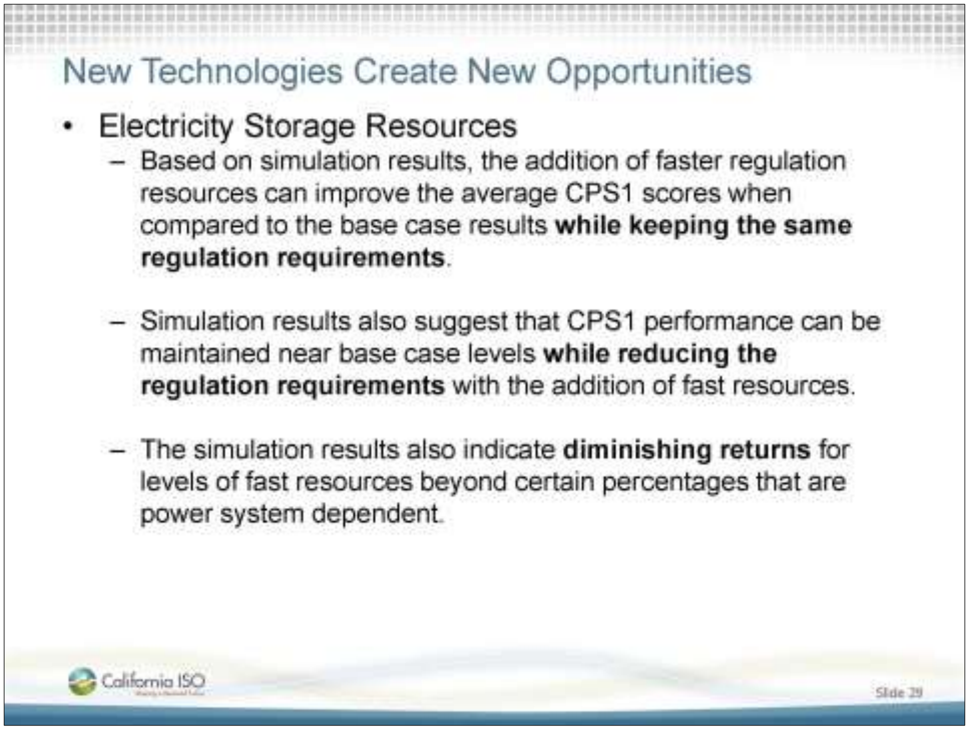

On this graph, the blue line is the original load. When we use storage, it is a new solution or a new technology that can provide us with a solution to these challenges… the opportunities.

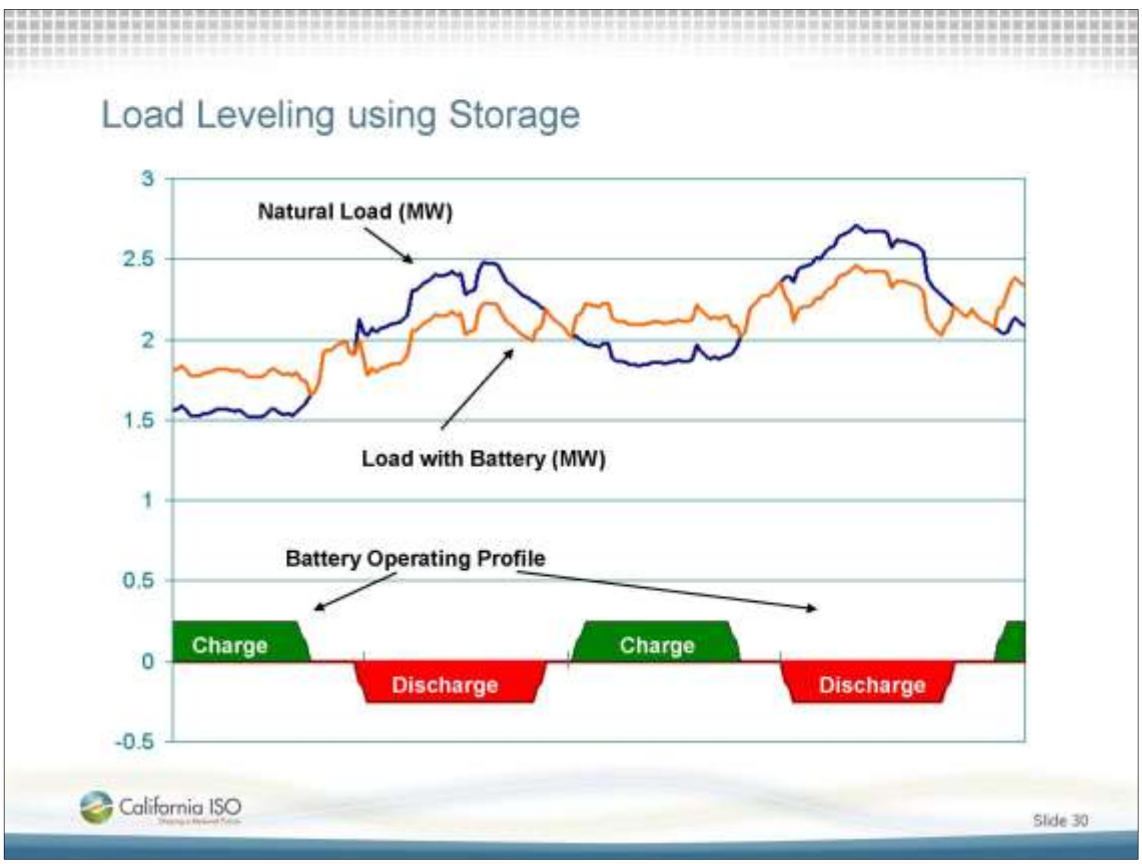

We use storage discharging and charging to provide load, ACE shaping, and renewable stabilization.

If we use storage and start discharging and charging it at certain periods of time, as in the bottom graph of slide 30, it can provide the load leveling, as shown in orange. If we put in storage there and start charging and discharging, our load will look like the orange line. It's almost getting flat. That is a very helpful feature because we've reduced the amount of standby ramping needed.

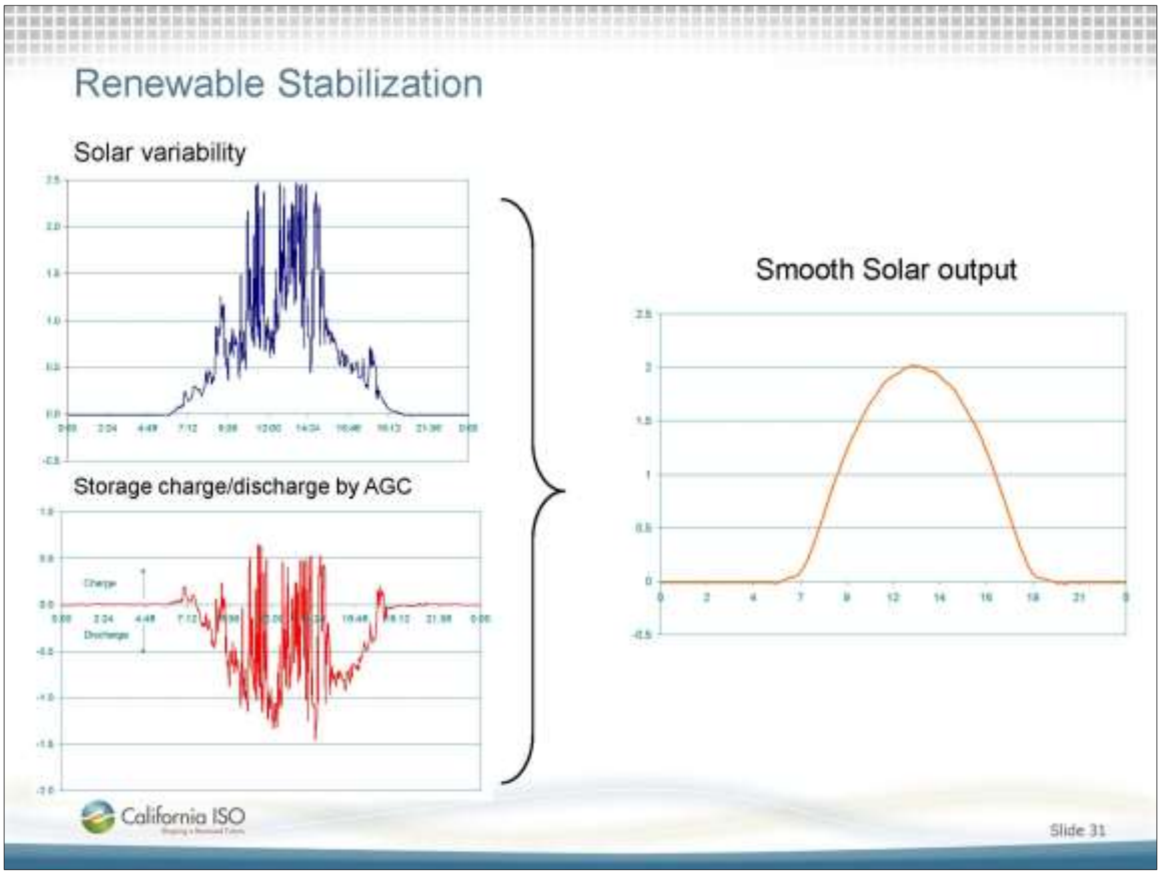

In slide #31, the blue line shows the solar variability. If we look at the solar output itself, within this time period, this is how it looks. When we use storage -- as in one of our pilot projects — its contribution is the red line. The graph on the right is the result when we couple these together. This is a very smooth line, a big change!

ACE shaping is a method of keeping the load and supply balanced.

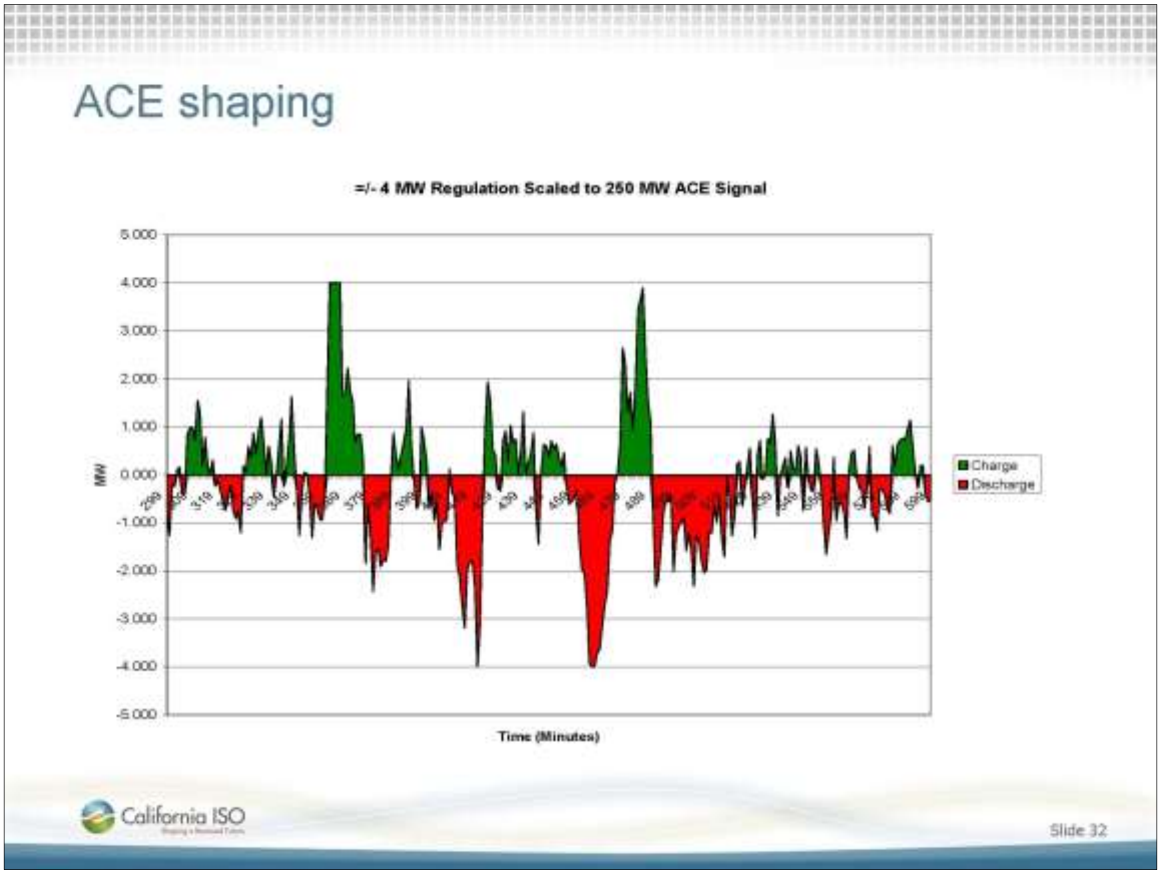

ACE shaping is a method of keeping the load and supply balanced. If we use storage, by charging and discharging it, we can maintain this balance. Zero on the graph means generation and load are balanced. There is a certain margin we are allowed to operate within. When we have extra generation, we can charge the storage just as if it's a load.

In periods where we have less generation, we can discharge the battery of the storage. That gives supplemental sources of energy to bring back the balance between generation and load.

Online voltage and dynamic stability analysis tools are invaluable to evaluate the changing system condition and provide advanced situation awareness.

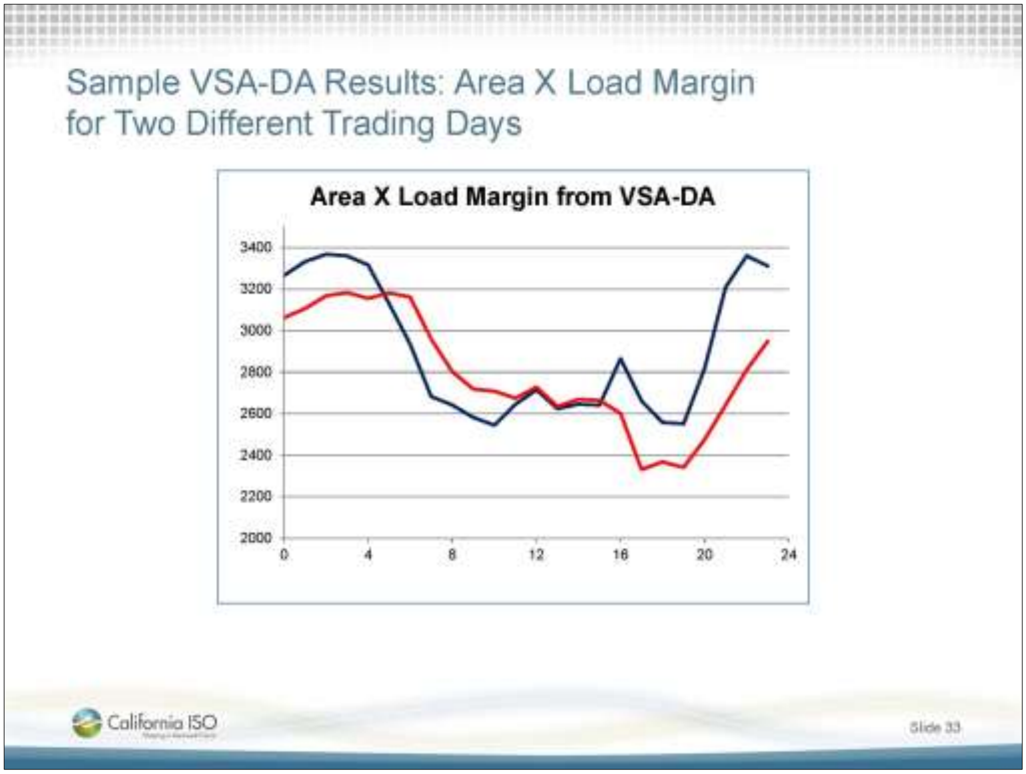

Online Voltage and Dynamic Stability Analysis (VSA, DSA) tools can be used to evaluate the changing system condition and provide advanced situation awareness about the stability margin and how far the system condition from voltage collapse. In California, we also pass the market look-ahead results through VSA and DSA programs.

In summary, there is a strong relationship between VARs and the voltage changes in the power system.

Increasing megawatt loading on the transmission line can cause dramatic increases in reactive power.

Operating systems at severe off-schedule voltages can lead to equipment damage, degraded reliability, and blackout… all devastating consequence.

There are different types of reactive power.

Lastly, fast reactive power controls are needed.

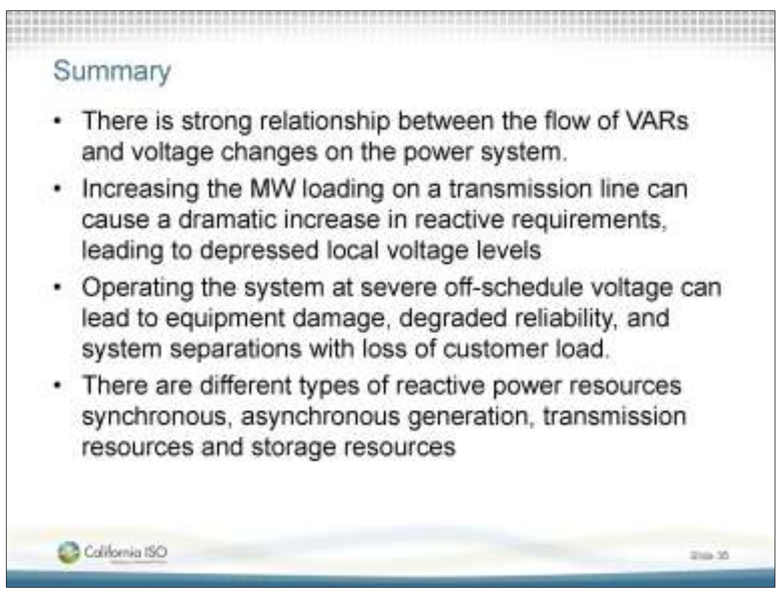

As new generation resources are replacing traditional generation, the new generation resources need to provide the same or equivalent features and ancillary services that the traditional resources are providing.

To the extent that new generation resources are increasing, replacing thermal or traditional generation, the new resources need to provide equivalent or similar features that the traditional resources are providing.

California ISO is totally committed and leading the safe and reliable transition to the clean future.

Thank you.

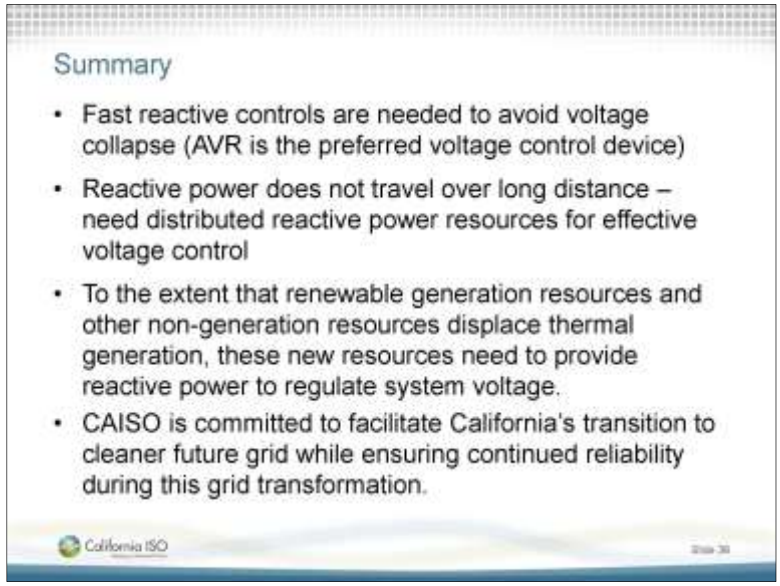

Influences on the Grid

DUSTY BECKER:

Don Von Dollen, our next speaker, is a Program Manager at the Electric Power Research Institute. He is responsible for the IntelliGrid program, which is focused on conducting research and development on communications and data integration for transmission, distribution, and consumer applications.

Mr. Von Dollen has been the Chairman of the IEEE PES Intelligent Grid Coordinating Committee, a peer reviewer for the DOE GridWise Program, and a liaison member of the GridWise Architecture Council. He coordinates EPRI’s Smart Grid activities with DOE, EEI, NIST, NEMA, NARUC, NCSL, and other organizations.

Mr. Von Dollen led a team of industry experts to develop a roadmap for the development and harmonization of smart grid interoperability standards under contract to the National Institute of Standards and Technology (NIST).

Mr. Von Dollen has worked with several utilities on their company-specific smart grid roadmaps.

Before joining EPRI, Mr. Von Dollen was a Research Engineer with the Pacific Gas & Electric Company.

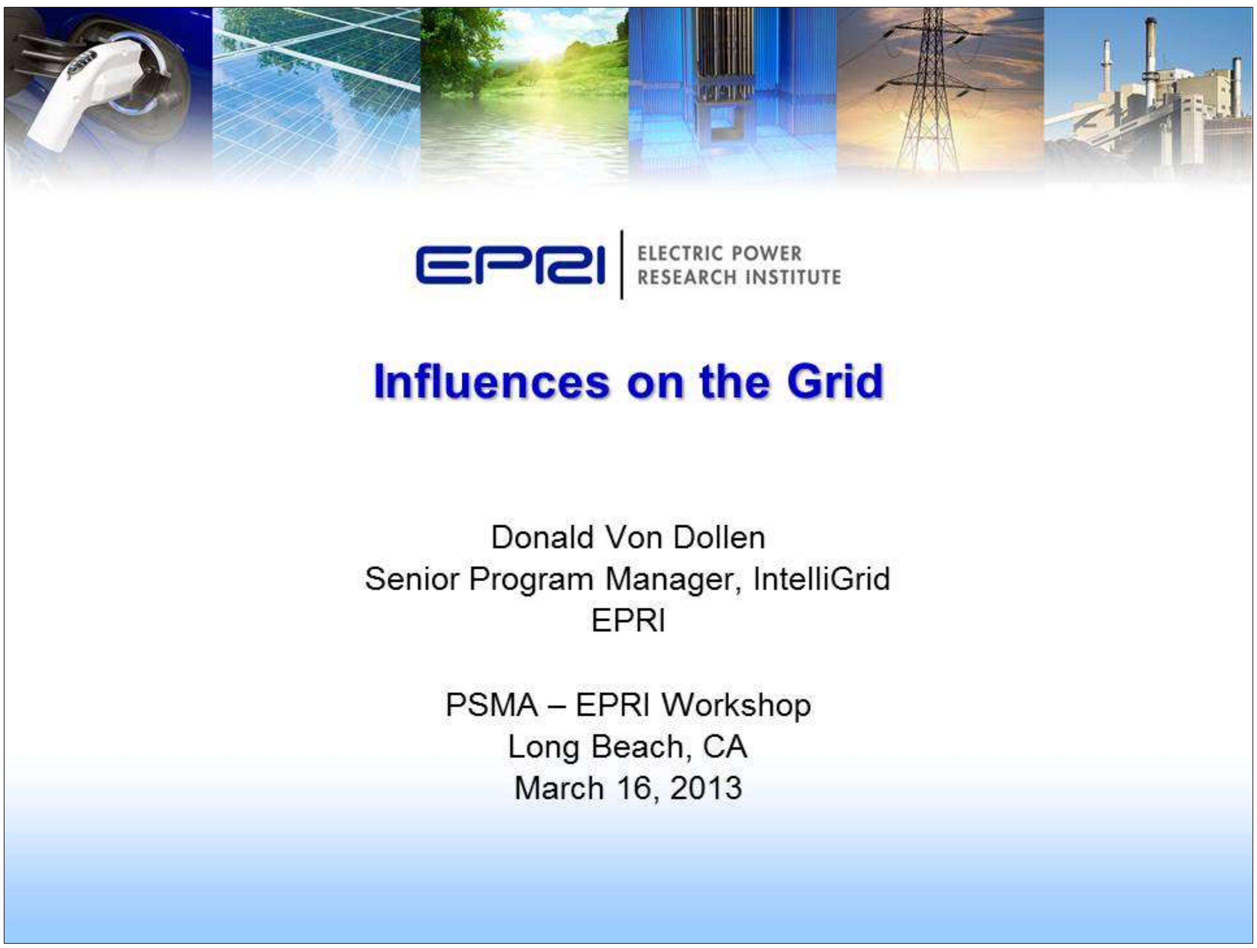

DON VON DOLLEN:

Thank you. Good morning, everyone. Really good presentations so far this morning.

So what we have heard so far from the two previous speakers, really, is about the grid today. What the current power system is. Then, following me, Dennis is going to talk about the smart grid. So I'm thinking that our three presentations are kind of the wind up and Dennis is going to be making the pitch about what a smart grid is. I'm kind of the bridge presentation. I'm going to talk about why.

Why do we want to move from the power system we have today to this "smart" power system of the future? The two previous speakers touched on this as well.

The U.S. National Academy of Engineering decided that the vast networks of electrification of North America were the greatest engineering achievement of the 20th century.

The focus really is: why change? Clark touched on this in his presentation. The U.S. National Academy of Engineering decided that the vast networks of electrification of North America were the greatest engineering achievement of the 20th century. Stop and think about it. That's really remarkable. Not the airplane, not the automobile, not the telephone, not computers… it was the electrification of North America, the electric grid. Clark called it the most complicated machine created.

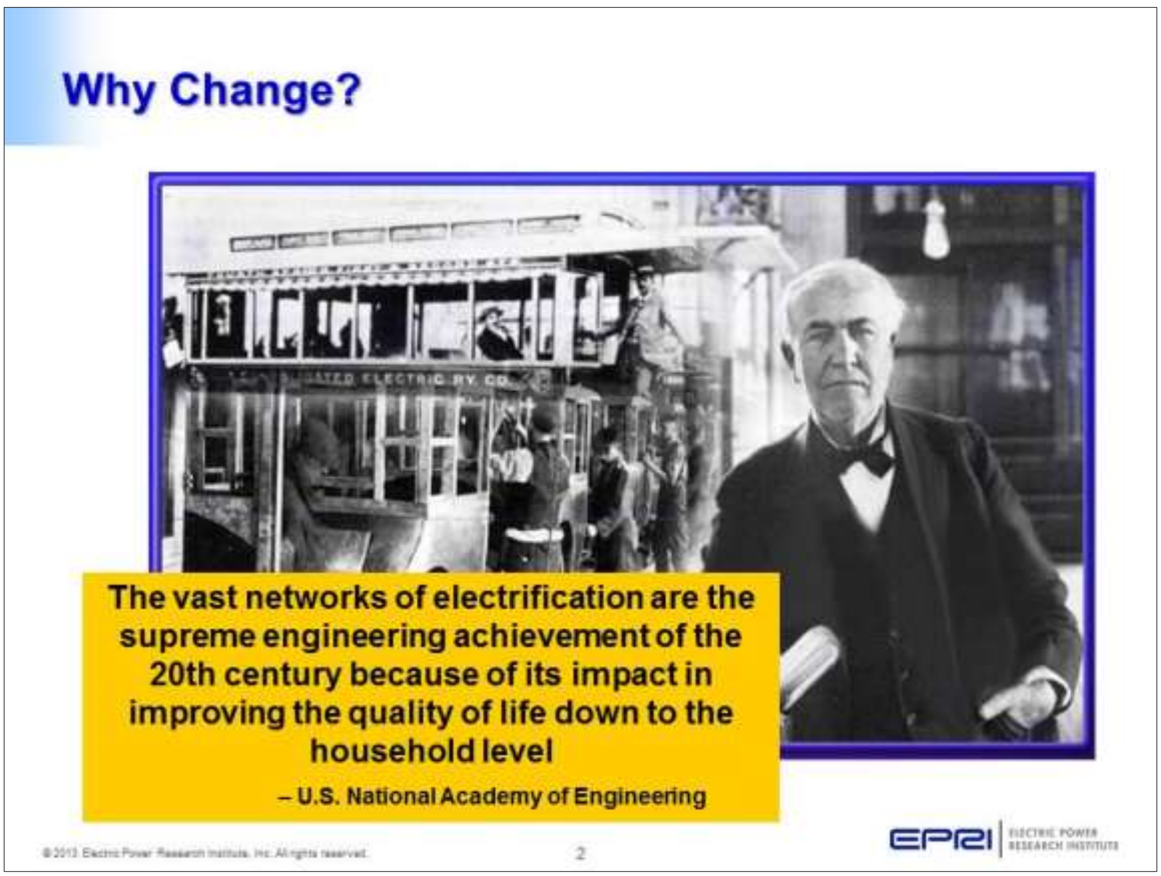

So what's the motivation for change? Why do we want to change what we have today, what's been in place for 100-plus years? Why do we need to do anything different with this great achievement?

That's what I want to talk about in this presentation. I want to talk about some of those drivers for change. Why change? At a high level, what are the new and enhanced capabilities the grid needs to acquire to address these drivers for change? I'm thinking that this should tee up Dennis' presentation, as he starts to talk about the smart grid. The rest of the day, actually, we'll talk about other technologies.

One of the drivers for change is extreme weather events… "1-in-100-year" events are happening more often.

Let me just start talking a little bit about some of these drivers for change. There are a lot of different drivers for change. I'm only going to touch on a few of them.

One is extreme weather events. We were talking about Sandy earlier in the presentation. I imagine some of you might live in the Northeast and have lived through Sandy. The rest of us were just looking at it through the news. We saw just what the remarkable effort that went into restoring the electric service following a large storm like that. It reminds us just how essential the electric power service is to our lives.

This is a slide that I downloaded from the Munich Reinsurance America that shows the number of extreme weather events from 1980 to 2010. The different colors indicate different types of events: earthquakes, tsunamis, storms, floods, extreme temperature events, draughts, and wildfires. You can start to see the trend as far as extreme weather events in the U.S.

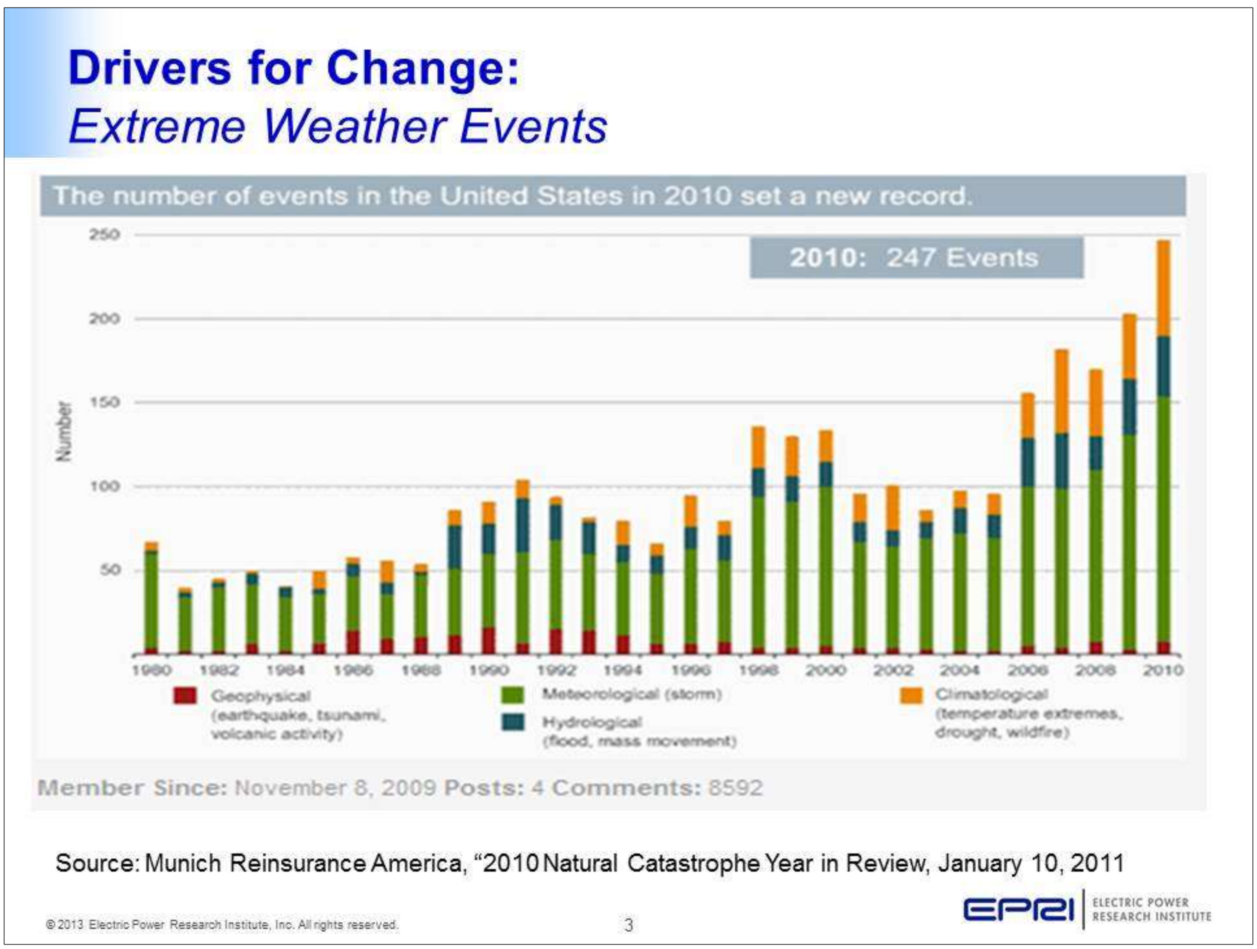

This slide shows the number of electrical disturbances caused by weather in the U.S. from 1992 to 2009. You can see the comparisons. We got this information from the NERC database. When we start looking for causes, anything that was a weather related outage, we started to see those types of relationships.

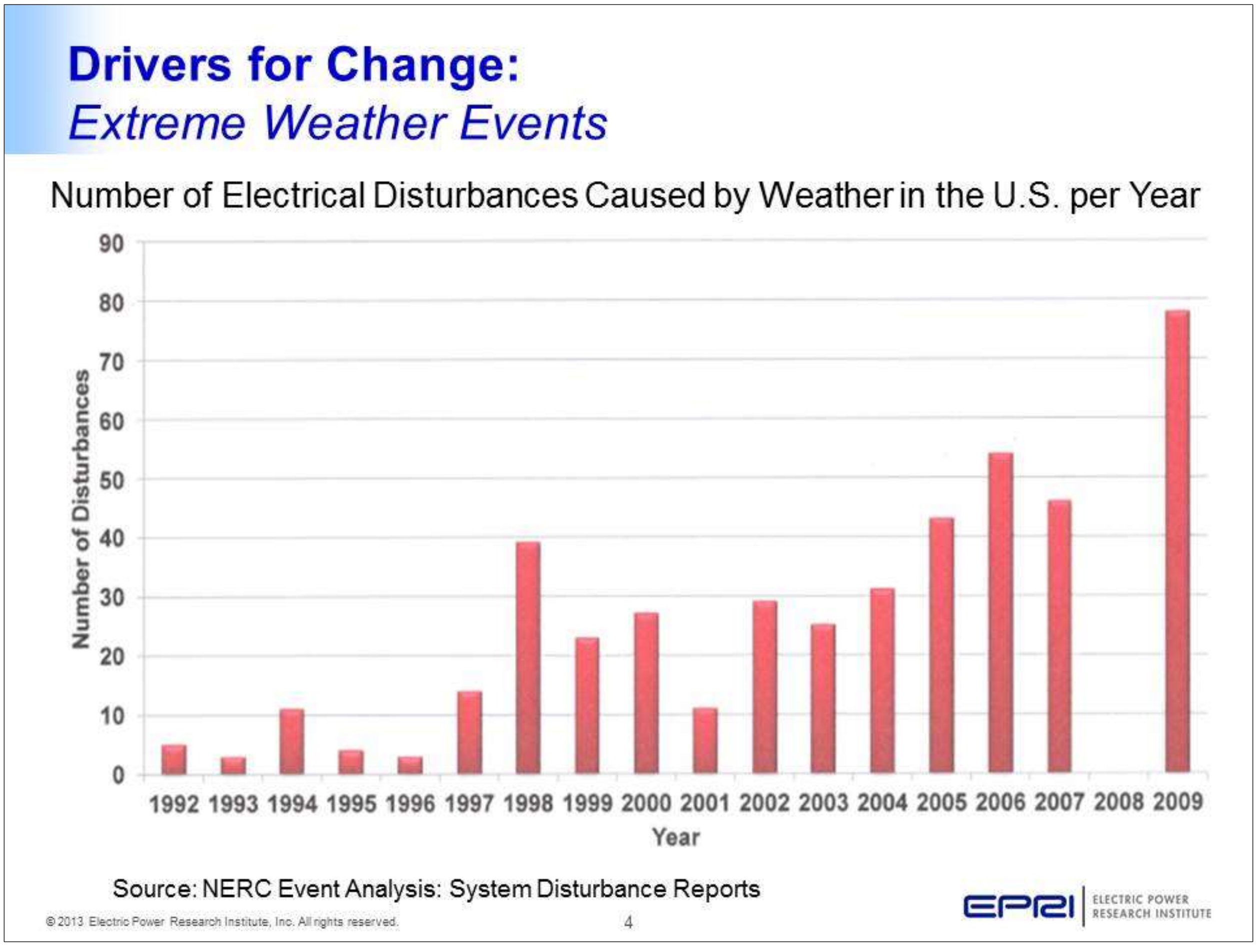

The basic design principles for a lot of power systems go back decades. The focus has been two things: reliability, keeping the lights on, and affordability.

One of the things that we see in the power system is that the basic design principles for a lot of power systems really go back decades. When you think about the electric power system and electric utility operations, for over 100 years, the focus has been around on two things. One is reliability, keeping the lights on; and the other is affordability, keeping the rates low for consumers.

Utilities have actually done a pretty good job optimizing that balance. What we start to see when we talk to design engineers is that they talk about designing the grid to survive 1-in-100-year events for temperature or wind or ice… but we're starting to see those 1-in-100-year events happening more often.

So… drivers for change.

We are looking at new designs for distribution systems and new technologies that can make the grid more resilient to recover from an outage more quickly.

This is one of the things we have to start to do, even if we want to maintain the existing reliability. This isn't saying that we want to get better reliability Just to keep our existing level of reliability, we have to do something different. Clark mentioned that most of the outages occur on the distribution system. New designs for distribution systems or new technologies may make the grid more resilient to recover from an outage more quickly and more easily.

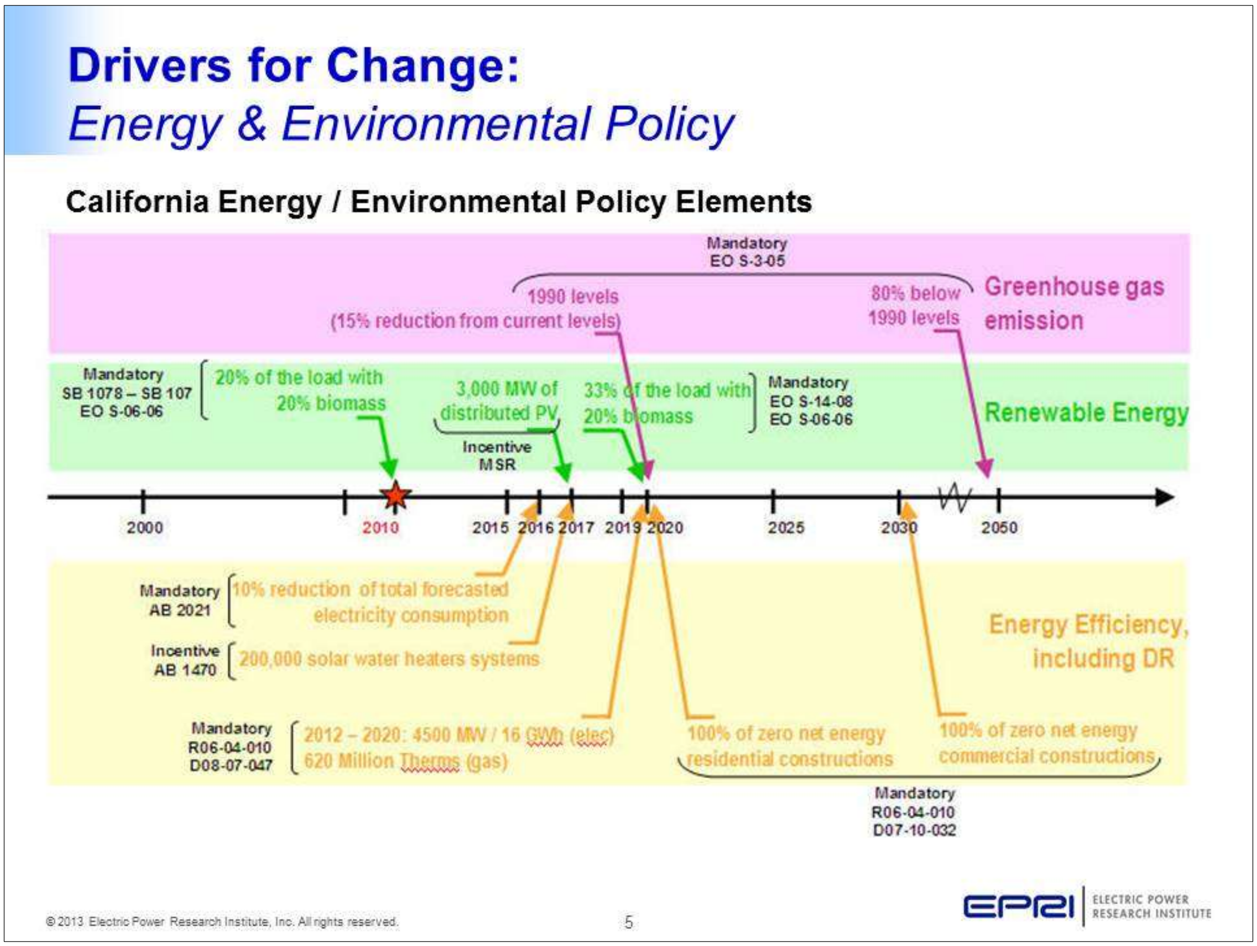

Another driver for change is energy policy. There are policies around greenhouse gas emissions, renewable energy, energy efficiency, and demand response.

Another driver for change is energy policy. How many people live in California? A number of us. You're probably aware of some of the key energy and environmental policies for the state of California. If you live in California, you know we're extreme. We're an extreme state. These are some of the policies around greenhouse gas emissions, renewable energy, energy efficiency, and demand response. We've laid down a timeline for some of the key milestones for California policy.

I'm not going to go through all of them, but I'll point out a couple. A lot happens in 2020 in the state of California. One of the goals, as far as greenhouse gas emissions, is to get down to 1990 levels of greenhouse gas emissions by 2020. That's about a 15% reduction. Which leads to the other policy, right? If you want to hit this target, you're going to have a lot of more renewable resources on the grid.

A lot happens in 2020 in the state of California.

So the policy for California is 33% renewables by the year 2020. We heard very clearly from the previous speaker about the challenges of trying to integrate that amount of variable resources on the grid.

From an operator's perspective, the great thing about a fossil fuel plant is you turn it on and you get a nice, steady stream of electrons. From an operator or planner perspective, that type of predictability is wonderful. Dealing with a lot more solar -- and 33% is an awful lot on the grid -- that much variable generation makes it very difficult to operate the grid -- or more "challenging" -- I should say it that way.

33% renewables is an awful lot.., it makes it more challenging to operate the grid.

When we start to hit these types of greenhouse gas emissions levels, we need to become much more efficient users of electricity. One of the targets is that by 2020, residential construction in California is supposed to be zero energy. That's the policy. That means that a new home is supposed to generate as much electricity as it uses. It's supposed to be net neutral. It ties into California's million-solar-home initiative, which is a goal of an awful lot of homes.

So, like I said, in California, we're extreme... but other states are coming up with policies that are similar in a number of different ways as well.

Another very big challenge is trying to change the electric service we have in the United States. This slide is difficult to read, but it’s good information on another driver for change: the ageing infrastructures. This is data from one utility we work with, but it’s representative of data from across the industry. It shows a number of different devices: power transformers, breakers, wood poles, relays; and the service life for each of these components.

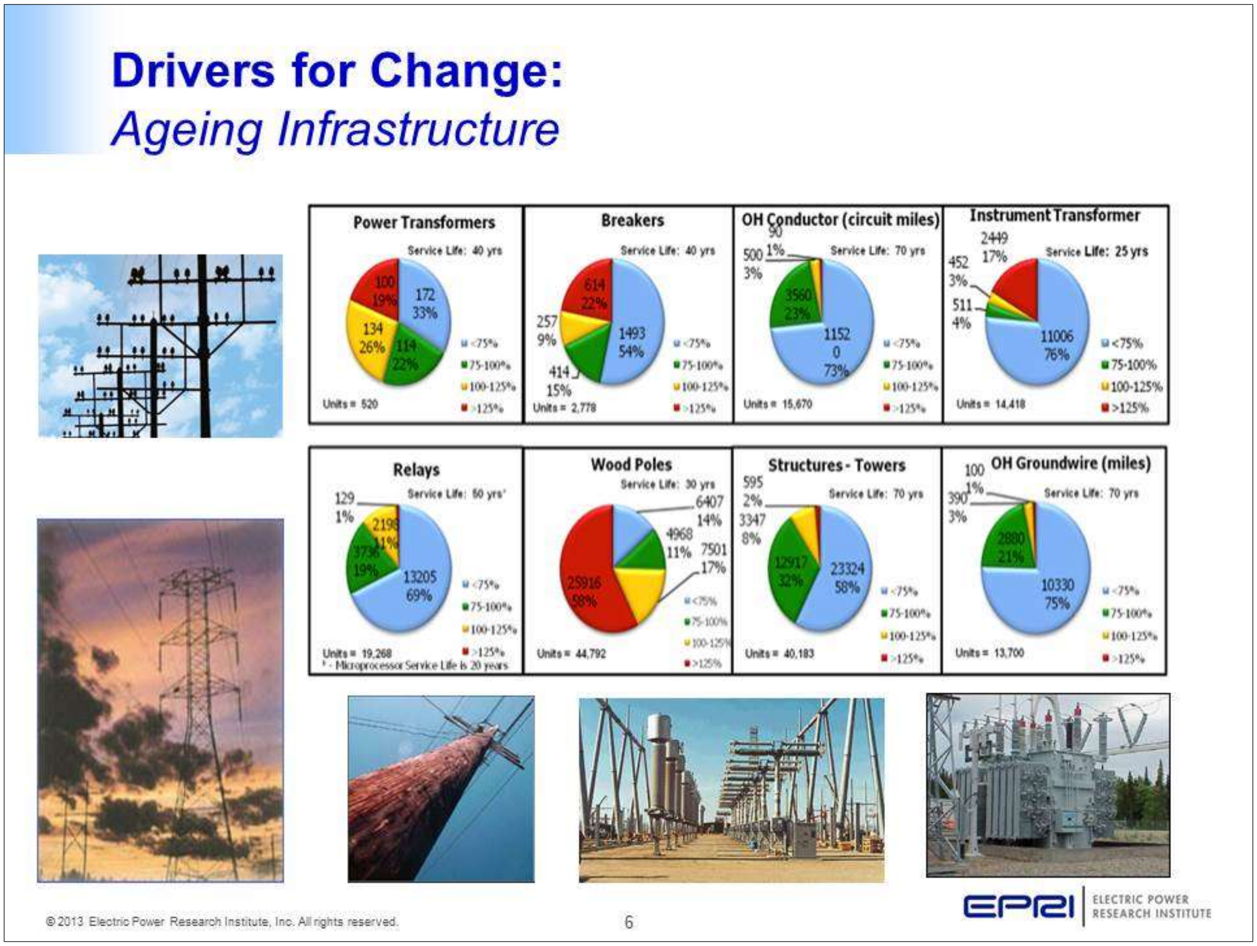

I'm going to pick power transformers as an example. Forty-years of service life for one of those big transformers that Clark showed pictures of. The blue slice of this pie means that the assets that are less than 75% of that forty-year life. Green is between 75 and 100% of its service life. Yellow is 100 to 125%. Red is greater than 125% of its service life. When you start looking at managing these assets, yellow and red are old assets.

Forty years is the service life for a big transformer: 45% are over 40 years old.

Age isn’t the only thing that determines how much life is left in a piece of equipment. We’ve dissected cables in our labs back in EPRI that we’ve taken out of service. We’ve taken 70/80-year-old cables out of service and they can look as new as the day they were installed. We’ve taken cables that were 20 years in service and it looked like they'd been through a war. So age isn’t the only factor… but it is a factor.

Another driver for change is actually an opportunity for us to grow our connectivity. This is data from Cisco that shows data usage from 2011 to 2016. Some factoids: Each U.S. resident now has, on average, five network devices. Tablets are at 34% market penetration. Smart phones are at 45% market penetration. I thought this was particularly telling: By the end of 2013, the number of mobile connected devices will exceed the number of people on earth.

Another driver for change is an opportunity to grow our connectivity for sensors for communications for data acquisition.

That's a great opportunity for us as an industry because one of the things that we're hearing about is the increased situational awareness that the grid needs to have to better management the ageing asset fleet. This is a great opportunity for sensors for communications and for data acquisition within the industry.

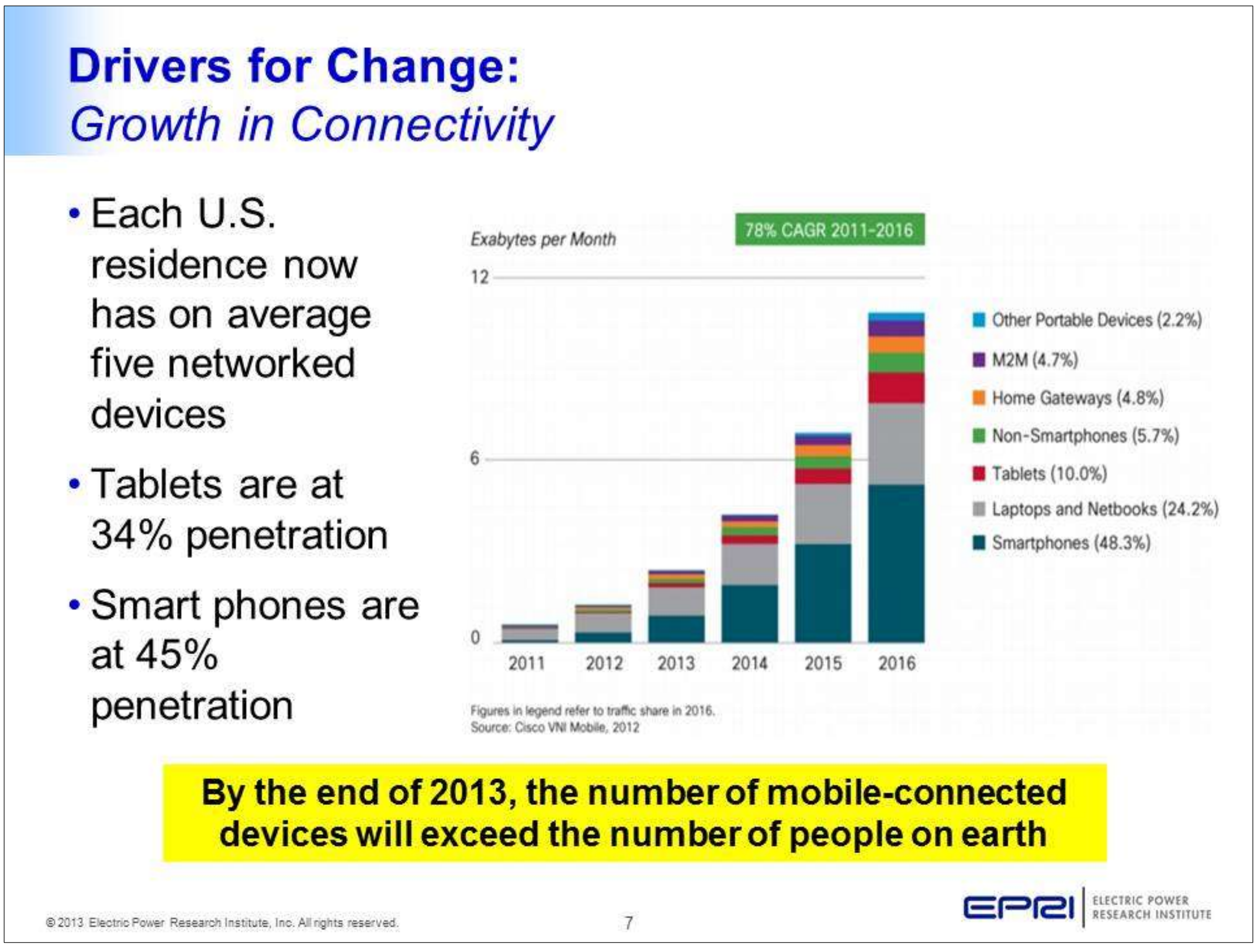

Drivers for Change:
Changing Consumer Expectations

- Consumers expect more information faster
  - Information on electricity consumption
  - Information on service restoration
- Expect utilities to know if their home has service
- Want information and services to conform to their lifestyle

Changing customer expectations take us into a different direction. For our industry, it's always been said that the two touch points that utilities have with customers: once a month they send a bill when there's a power outage, the customer picks up the phone and yells at the utility.

Changing customer expectations: they expect more information faster.

Customers expect more and different service from their service providers. That's what a utility is: one of many service providers to the customer. Customers expect more information faster. A lot of customers now want information on their electricity consumption, not just monthly information. They want information on service restoration progress. If there's no power at home, they want to know what happened and how much longer before power comes back on.

They expect utilities to know if there's service in their home. Traditionally, the way utilities know of an outage is when people start calling and complaining. That's how they've known if there's power flowing to a certain area. It's very interesting. Some utilities have been installing smart meters, so they know if there is electric connectivity to a customer. As information comes back, utilities have been surprised by the number of outages that nobody makes a phone call about.

Again, it is the expectation of customers. Customers expect that when they call a call a customer service center, they'll know if there's power in their home or not. They want information and services to conform to their lifestyle. They want information coming to their smart phones and their tablets. They want information coming to them the way that they like to get information.

So when you start looking at all of these influences, all of these drivers for change; what does that add up to? What are the capabilities? How does the power system need to change?

The grid needs to change. Number one is resiliency.

Not only are we going to need to be able to harden the system so it can withstand more and more extreme weather events, but we need to have better situational awareness. We mentioned that earlier. We need better monitoring communications. And as we get more and more data, we need the analytic statute of being able to interpret that data, turn it into the actual information.

What about this idea of more self-healing systems? This means more automated bulk detection, isolation, and restoration. We need to be able to recover from failures more quickly. We need to make use of technology that will make the system more resilient.

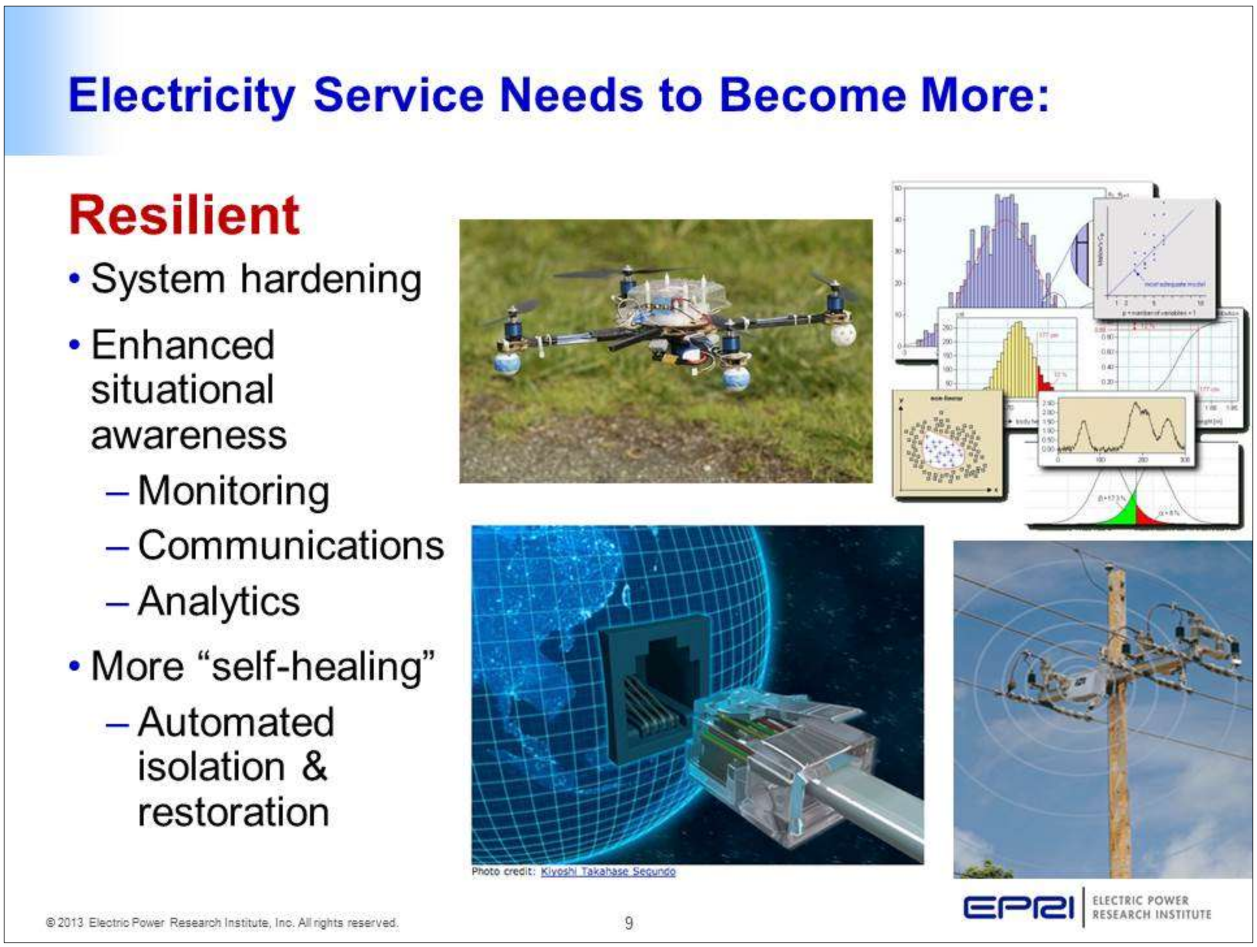

This is one of the technologies we're developing at EPRI. This is a drone. This was actually used during Sandy to look at damage, storm damaged areas, and to do a very quick assessment of the extent of the storm damage. We're also developing pattern recognition algorithms that can tell what's different, the difference between what the system looked like a week ago and what it looks like after the storm. You can, in a very automated way, start to determine what's changed and what's been damaged by a storm.

As the grid has more distributed generation, the grid needs to become much more bi-directional.

Flexible. The grid has always been designed -- and Clark touched on this in his presentation – so power flows in one direction. It flows from that central generating station down to the point of consumption. As we move to more and more distributed generation, the grid needs to become much more bi-directional. Power needs to be able to flow upstream as well as downstream. As we start hardening distribution circuits, we’re going to go from a hub-and-spoke concept, where power is on the radial lines that go out from a substation, to more looped configurations, which are more resilient. If we have an outage on one radial, we don’t lose power to all of customers outside of that radial. If we loop it, we can feed power from the other direction as well. If you’re generating power, you can feed power back in more easily.

Greater amounts of demand / response, enhanced situational awareness, and automated control for flexibility.

Greater amounts of demand / response. If you reduce the load, reduce the consumption on a peak afternoon, it’s very helpful to the grid operator. That’s a good tool for the operator.

Being able to accommodate electric, plug-in electric vehicles, varies in different parts of the country. Even within California, it depends on what parts of the state you’re in. Some distribution networks can accommodate electric vehicles more easily than other electric distribution systems. It depends on the kind, when they were built, and how they were designed.

Efficiency. Clark mentioned this as well. The grid really needs to become more efficient. We have about a 7% loss in available power in the transmission and distribution networks… so even small changes that reduce that 7% can have a very large impact.

Maximizing the useful life of an asset base directly ties into rates.

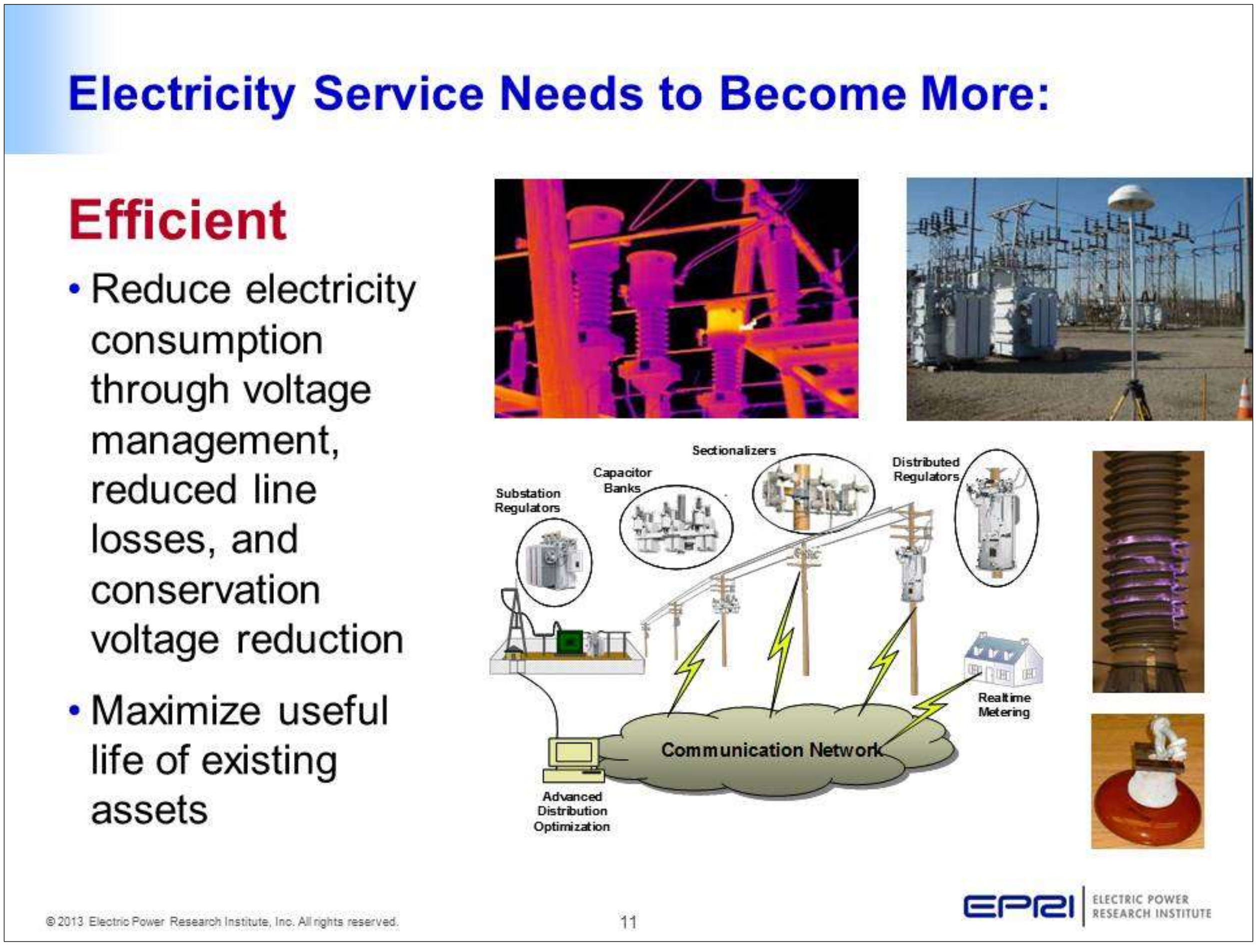

Efficiency is a way of maximizing the useful life of the asset base. It becomes very important because it directly ties into customer rates. In a regulated environment; if we have to replace those assets, that goes right directly to rates.

Our electric service needs more consumer focus. Customers need to have more information, tools, and technologies that help them better manage their electricity usage.

Finally, consumer focus. Our electric service needs more consumer focus. We need more efficient in-home and in-building devices. Customers need to have more information, tools, and technologies to help them manage their electricity usage, to become more efficient users of electricity. Value-added services and greater choices can flow from that. I think Dusty made the comment that a vendor is probably going to give me a call. Very likely; very likely!

Some of these devices can be the combination of smart meters and in-home displays and energy-efficient devices. This one… this is the logo for green button, a White House initiative to develop a standard for consumer energy usage data; to structure and store that data in a very standardized way.

Three California industrial utilities have implemented this already. Go on the PGE or Southern Cal Edison website and push this button. This button is on their website and you can download your own data in a standardized format. You can share that data with third-party service providers. Now there are all kinds of developers and dozens of apps that can make use of this data in this standardized way. It presents information however you like to see the information.

The Challenge

- The challenge is to not only create more resilient, flexible, efficient and consumer focused electricity service that spurs economic growth & promotes environmental responsibility, but to do so in the most cost effective and timely manner possible.

 13 EPRI ELECTRIC POWER RESEARCH INSTITUTE

Okay, so these are just the characteristics that the grid is going to really need to adopt over time. The challenge is; not only to create a more resilient, flexible, efficient, and consumer-focused electricity service; but to do it in the most cost-effective and timely manner possible.

To be cost effective, we can't just go gold plate the grid. We need to create new and enhanced capabilities, but we need to be very mindful of how it's going to affect rates; how much it's going to cost.

Timely is not a one-size-fits-all type of smart grid. In every part of the country, every utility has different needs. They're going to have to adopt capabilities at different rates. The smart grid must be able to bring those capabilities on when they really need to. That becomes the real challenge.

Thank you very much for your attention.

What is the "Smart Grid"?

DUSTY BECKER:

Our final speaker for this segment is Dennis Symanski.

Dennis is a Senior Project Manager at the Electric Power Research Institute. His responsibilities include investigating ways to make data centers more energy efficient; supporting Southern California Edison's Irvine Smart Grid Demonstration Project; representing the U.S. on International Electrotechnical Commission technical committees, and writing electromagnetic compatibility standards.

Before joining EPRI in March, 2009, Mr. Symanski worked on power quality standards and data center operations at Sun Microsystems and Data General; worked on relay coordination, electric vehicles, and windmills at New England Electric System; and worked on electric power generation, transmission, and distribution at Exxon Research and Engineering.

Mr. Symanski holds a Bachelor of Science degree in electrical engineering from Northeastern University and a Master of Engineering degree in electric power engineering from Rensselaer Polytechnic University. He is a registered professional engineer in California and Massachusetts.

DENNIS SYMANSKI:

Thank you, Dusty.

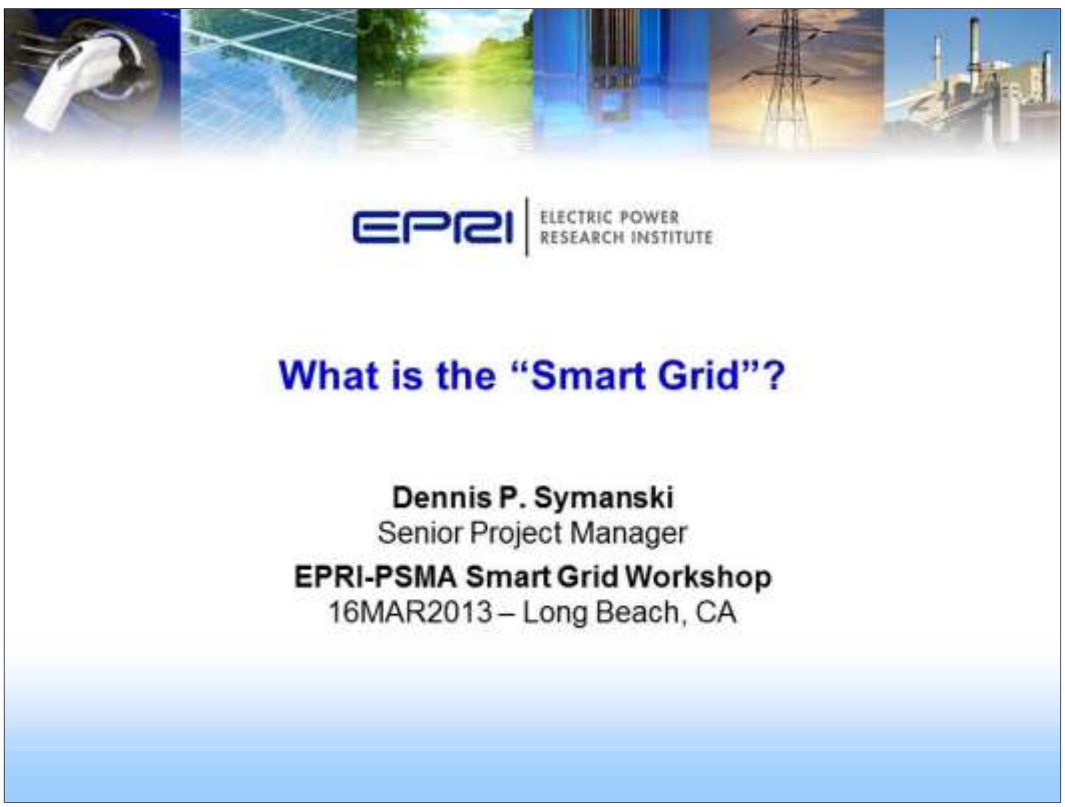

So I'm going to have to explain what the smart grid is. When I went to college, watts were called "real" power and VARs were called "imaginary" power. That just shows you we've come a long way... but we still have a long way to go.

Secretary Chu said, "America cannot build a 21st-century energy economy with a mid-20th-century electricity system."

Secretary Chu is no longer at the DOE, but when he was there, he made a comment that we can't keep using the mid-twentieth century grid to actually feed the economy of the twenty-first century. We need to start transforming the grid from a one-way system, which has worked very well, but we need it do more things now. We need it to do things like integrate generation scattered everywhere. We need it to take a look at plug-in electric vehicles when everybody comes home at 6:00 at night. Everybody plugs their Tesla in. We have to make sure that all the transformers are manageable and voltage is adequate to provide all of the other needs that we have in our homes.

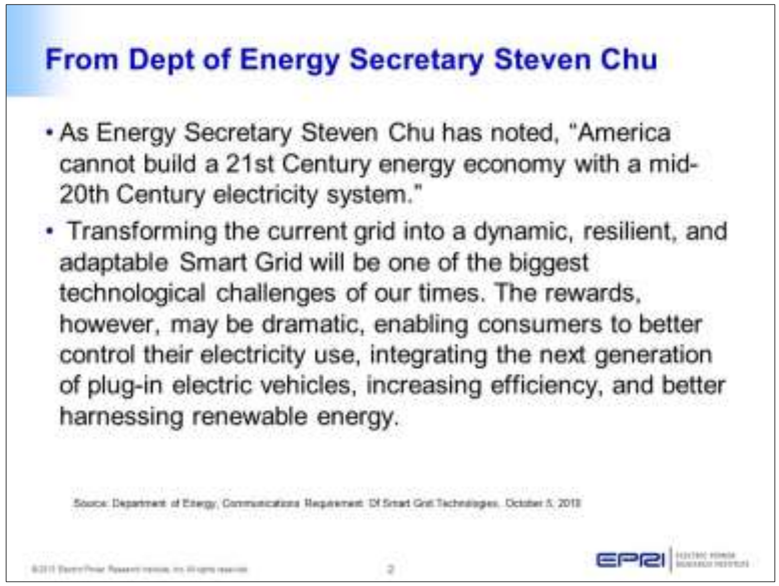

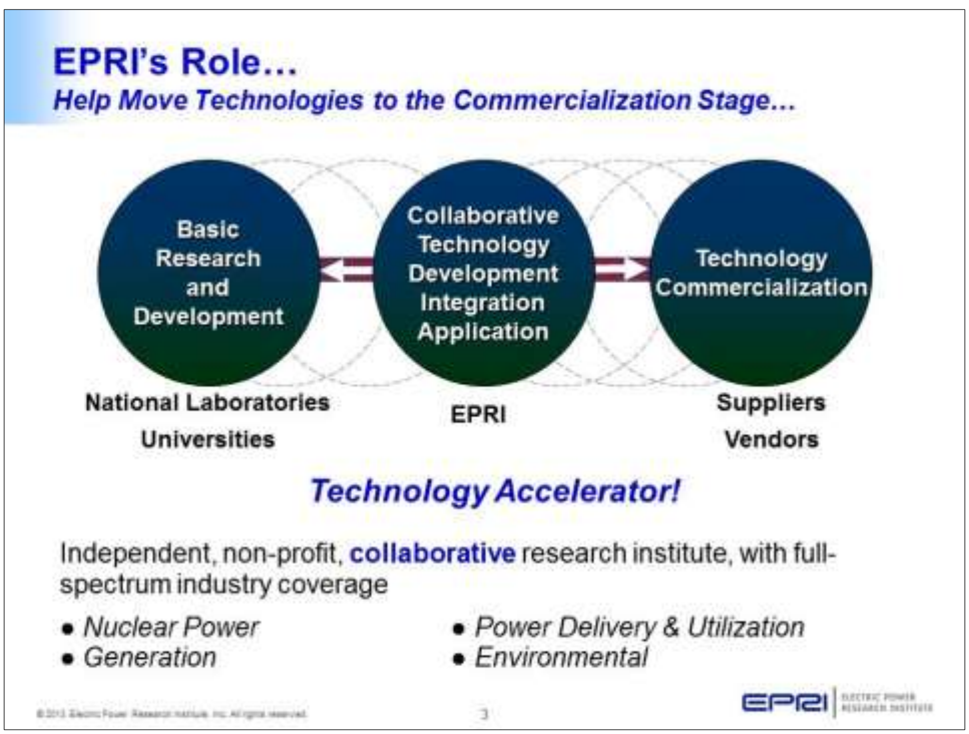

One thing about EPRI, I just want to say, is we consider ourselves a technology accelerator. We work with early research, with places like National Labs, and we work with suppliers and vendors to take technologies we think have an opportunity to make an impact on the electric system and accelerate their adoption.

EPRI is a technology accelerator.

We deal with everything from nuclear power plants to fossil plants. We look at the environment. We look at the power delivery and utilization, the plug loads.

We've been talking about the smart grid for quite a while. President Obama even mentioned it in one of his State of the Union addresses. So everybody is saying: "Oh, the smart grid? It's here. We're doing things." There's an expectation that the grid is going to be smart tomorrow. Then we have things like Hurricane Sandy comes along and, all of a sudden, the grid in New Jersey and New York was not smart enough to keep everything running. Some people were out of power for over a week.

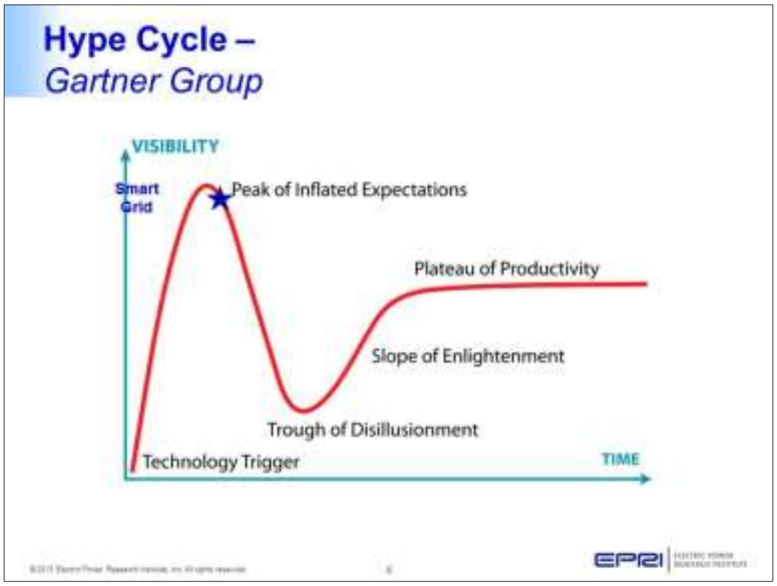

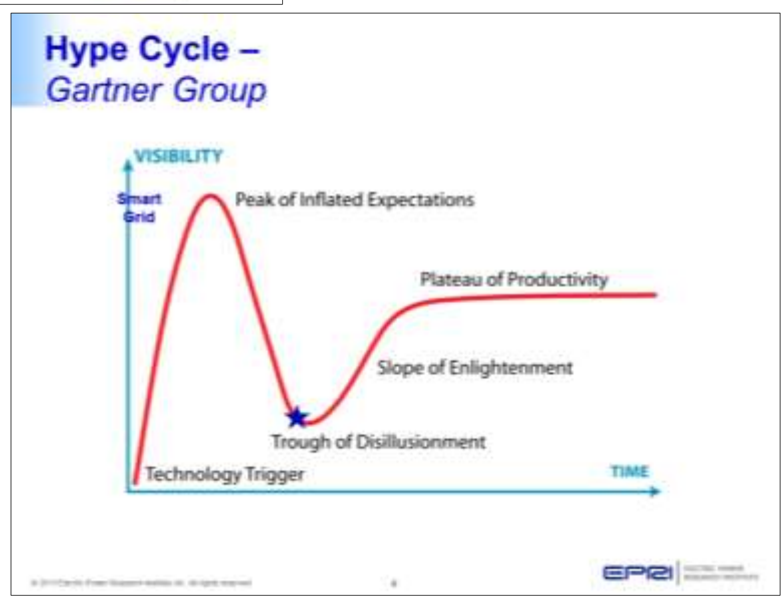

We're in the early stages of making the grid smart — it's not here yet.

There are things happening, but we're in the early stages of making the grid smart. There's going to be a realization that some of these technologies will work and some of the technologies will not. We'll eventually get to point where we have a smart grid, but it's not here yet. We're doing research to try to make it come to fruition.

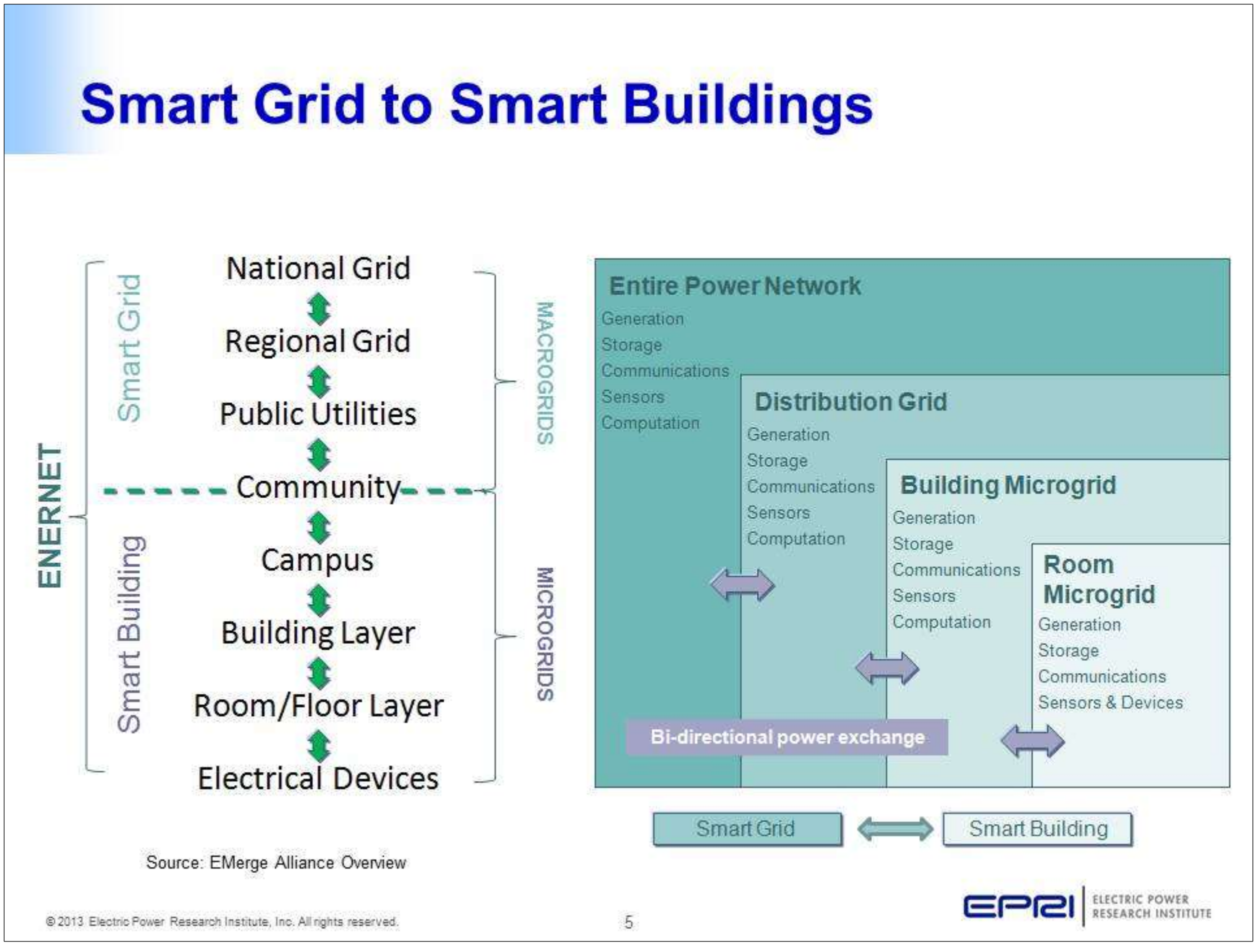

If you ask ten people what their definition of a smart grid is, you'll get about 12 definitions! Everyone has a different view. What is the smart grid? If you look at CAL ISO and FERC, you're looking at the entire North American continent. It's a huge interconnected grid. If you look at it from the green dotted line up, that's what the utilities are looking at as far as the grid.

Okay, so we're trying to make it smart. It used to be, as Clark said, a one-way street. You used to have huge power plants way out in the boondocks, transmission lines that bring power to the substations, which then feed it to distribution circuits into homes, commercial buildings, and industry. If you look at it, the grid is becoming a bi-directional power exchange already. It used to be one way and now we're getting down to grids called "campus" buildings. Even a room like this... this could be called a microgrid. We're trying to do things through the EMerge Alliance and other organizations to make this room more efficient and more intelligent.

We're trying to talk to devices for bi-directional communication from the utility to the plug load.

We're trying to be able to talk to devices so that it's bi-directional all the way, back and forth, from the utility down to the plug load… and we're getting there.

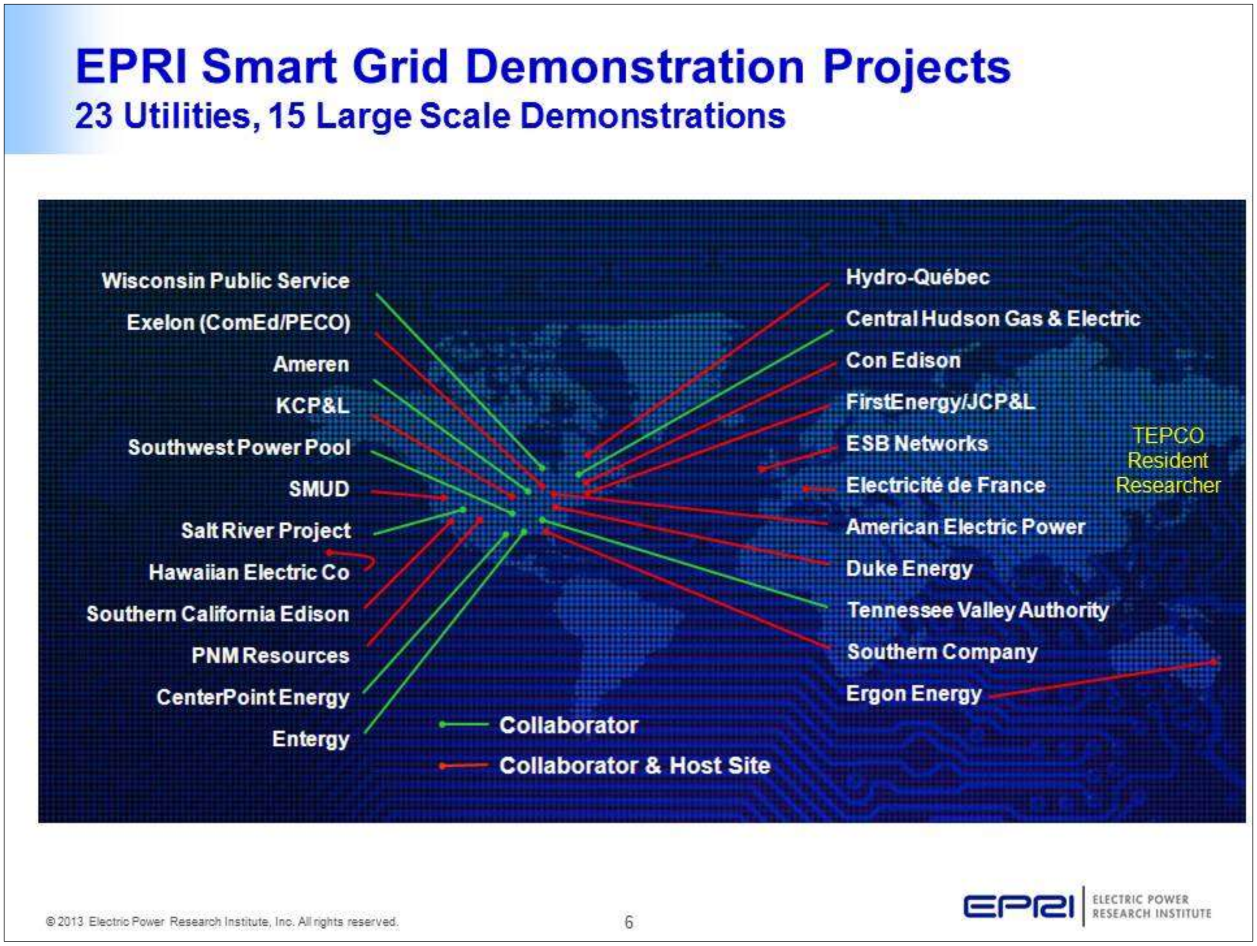

Right now, EPRI has a smart grid demonstration initiative going. We've got 23 utilities involved; 15 with large-scale demonstrations. We're talking about millions of meters and all sorts of end uses worldwide. This is not just in the United States. We've got Hydro-Québec, Électricité de France, Australia's Ergon Energy, etc. We're trying to coordinate all of these different demonstrations, all of these different technologies; to find out what works well; what doesn't; what works well, but is too expensive to implement; what technologies cancel the benefit of another technology. We've got to make sure that all of these technologies play well together, which is not an easy thing to do.

We need to find out what works well; what doesn't; what works well, but is too expensive; what technologies cancel the benefits of other technologies.

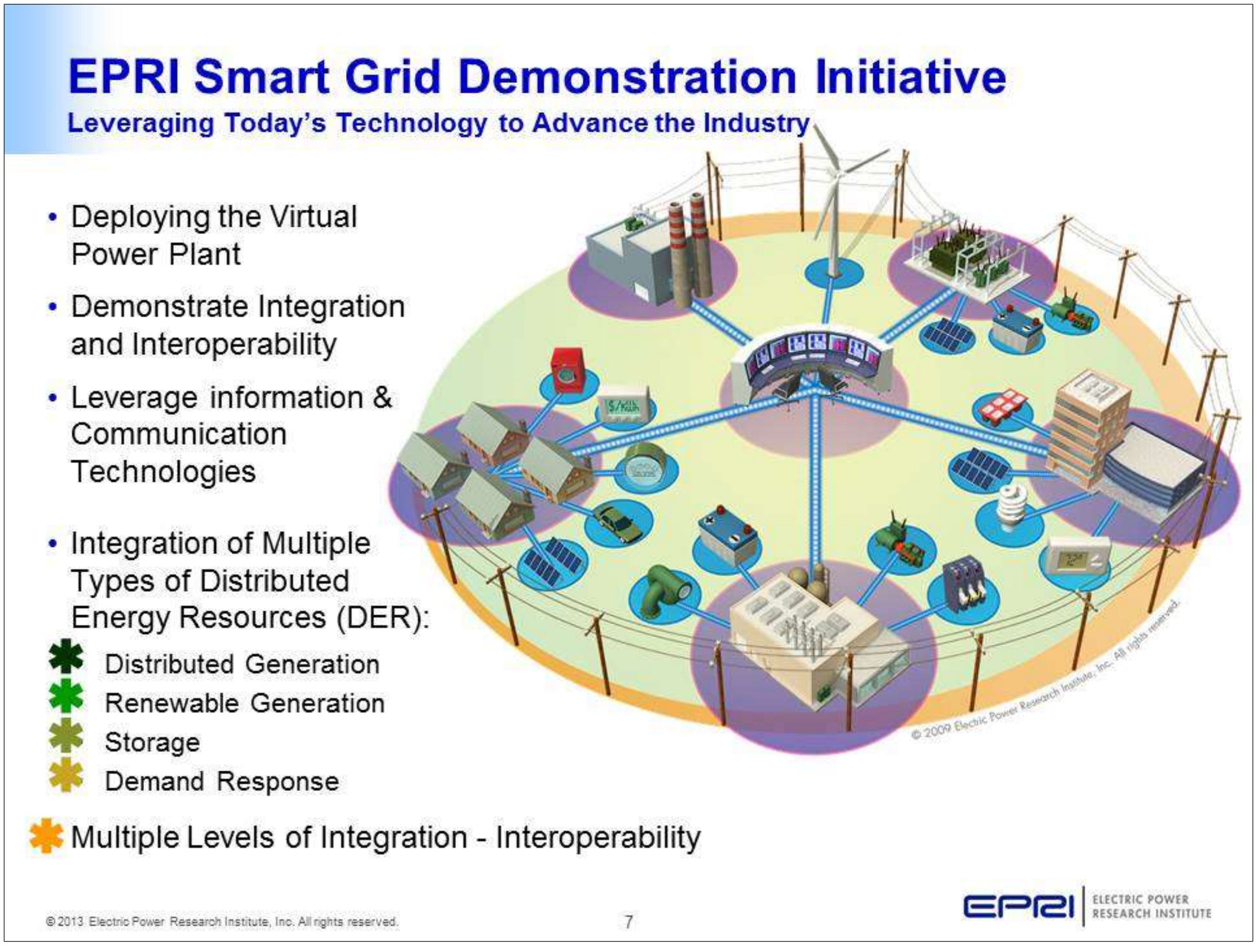

The initiative is trying to do things like create a virtual power plant by putting all of these technologies out there instead of building another hydro plant, nuclear plant, or fossil fuel plant. We're trying to see if these technologies can be modeled together so that, instead of building another power plant, we can call on devices; whether it's a plug load in a home, a pool pump, a programmable communicating thermostat; to make adjustments. We could change the settings of your air conditioning system, shut off your refrigerator for 15 minutes at a time, etc. to reduce load. It would be controlled by people like CAL ISO so they'd know that they can get a certain number of megawatts worth of load reduction when they need it to prevent building a new power plant or switching on a peaking plant, which is the least efficient plant in the territory to supply the last few megawatts.

We're trying to create a "virtual power plant" by using these technologies instead of building another hydro plant or another fossil fuel plant.

We're going to do a lot of things: distributed generation, renewable generation, storage. Satish will be talking about what type of things we can do with storage to try to make the grid smarter. And demand/response... when there's a critical peak demand, can we go out and tweak some of the plug loads to reduce the load and avoid needing to crank up the last, least-efficient power plant on the grid?

One of the demonstrations is by Southern California Edison at UC Irvine. We're taking a bunch of faculty housing and trying to do what the California Energy Commission and public utility commission is telling us to do by the year 2020, any new construction residential building must be zero percent energy. Over the course of one year, they must generate as much power as they use.

We’re taking four blocks of faculty housing with one block representing a control block. We’re doing nothing but putting in smart meters to measure how much power they’re using. As we go to back from the control group, we have a maximum block with full ZNE capability. This is what would be considered zero-net-energy. We’re giving them plug-in electric vehicles. We’re putting energy storage into the garage, something about the size of a small refrigerator with 4 kilowatts worth of power capacity. We’re replacing all of the appliances with hyper-efficient appliances, putting in triple-pane windows, putting in custom PV panels on the rooftops, creating a home area network where they can see how much power they’re using and the price of the power that varies hour by hour for critical peak pricing.

If a heat wave comes through, there’s going be a very high level of demand. You can actually go in and pre-program through your home area network. You can say “Okay. I’ll increase my thermostat setting from 72 to 78 degrees for this four-hour block of time. I will allow Southern California Edison to go in through the home area network and, through the smart meter, control some of the devices, such as a pool pump, a refrigerator, a washer/dryer.” These appliances have to become smart and become part of the demand-response solution.

Let’s allow Southern California Edison to go in through the home area network and smart meter to control devices; such as pool pump, refrigerator, washer/dryer; so these smart appliances become part of the demand response.

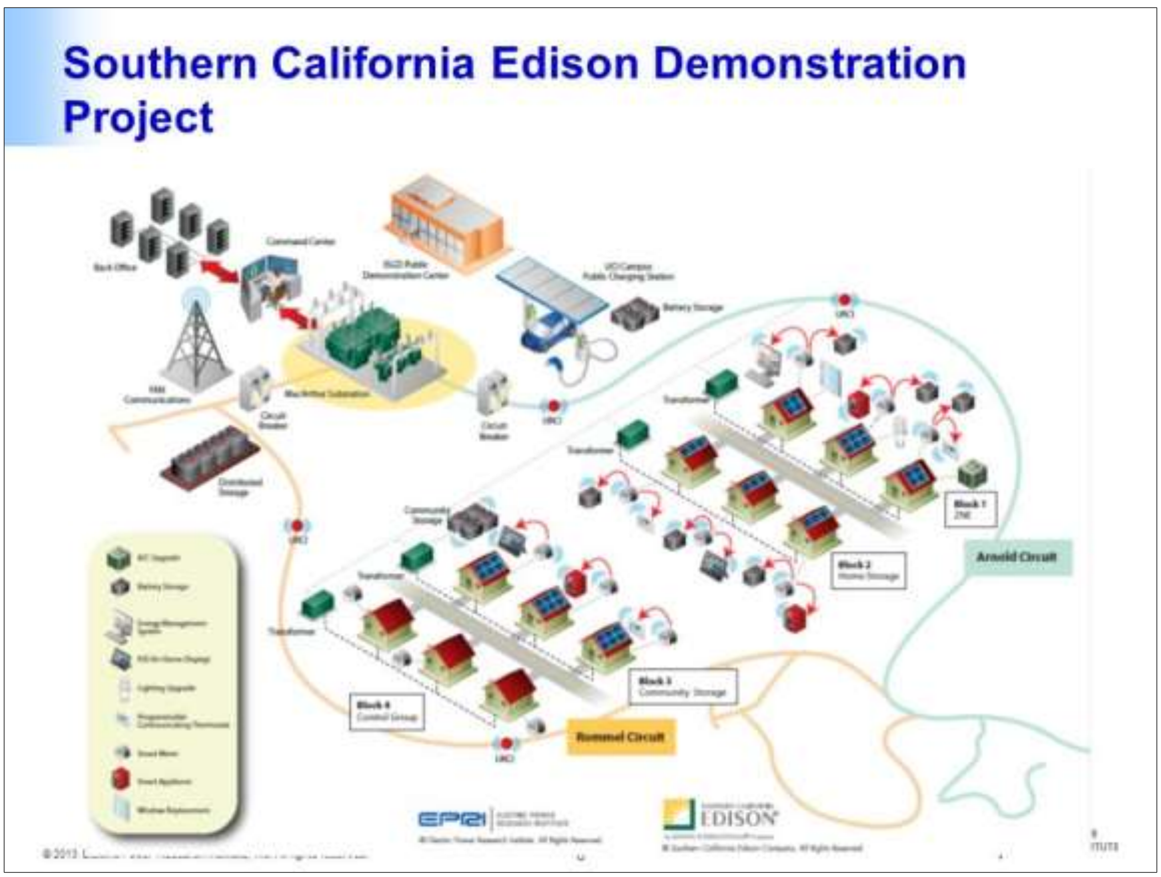

They’re going to talk about things like self-healing circuits. Traditionally, distribution circuits have been one way; from the substation to a load. And they’re radial. What they’re doing here, if you look at the diagram, is one circuit coming out of (the orange one) being tied together with another one (the green one) to create a loop.

They were never intended to be tied together. This was always supposed to be radial. But this is going to become a loop circuit. They're putting in circuit breakers with wireless communication along the circuit here, here, over here, and here. They're going to talk to each other. If there was a fault somewhere in the line, the entire circuit would normally trip out. Everybody was out of power until the trouble truck went out, found the problem, cleared the fault, and reenergized the entire circuit.

Loop distribution circuits will reduce the area affected by a fault. Communications will tell the trouble truck where to go.

If there's a fault somewhere here now, these two circuit breakers are going to talk to each other and say the fault is somewhere in between those two. Those are the only two that will trip. All the other homes on the rest of the circuit will stay energized. Not only will that minimize the number of people who have an outage, but it will also tell the trouble truck to go out there and only look at this little section of the circuit, not the entire circuit to figure out where the problem is.

There's also a seven-story parking garage on campus. They're going to put PV panels on the roof and a big community energy storage unit, one of these containerized lithium ion batteries with about 20 charging stations below. They're hoping to figure out the usage. How much energy can you really produce during the day to charge a vehicle? How much can you store in the battery to charge when the sun isn't shining? How much do you use the grid overnight if somebody wants to charge their electric vehicle overnight?

Electric vehicles are a big question mark in the smart grid.

The electric vehicles are a big question mark in the smart grid and how they will affect the grid as we're operating it. In the past, we would need to make changes to ensure a one load doesn't interfere with the other loads are on the circuit. There are a lot of different technologies we're trying to test within this Southern California Edison demonstration project. We'll see how it all works. We're in the process of introducing this hardware into the field. By the summer, it should all be operational and we can start collecting data to find how it's working.

This is part of our tax money at work.

This is American Recovery and Reinvestment Act of 2009 (ARRA) funded project — this is our tax money at work. It's going to teach us a lot. There are other utilities doing very similar projects. In Sacramento, there's a Sacramento Municipal Utility District project. As I said, we're involved with 15 total projects.

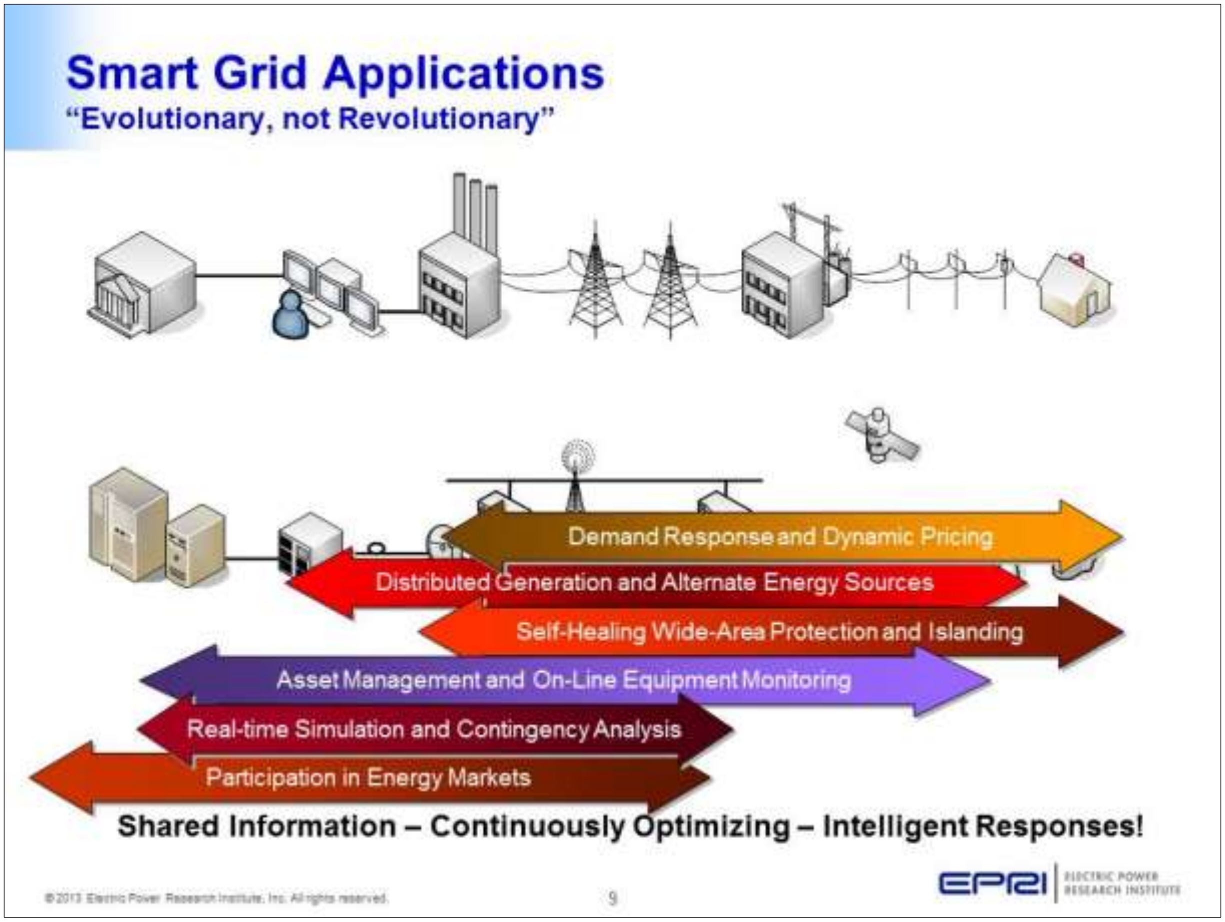

This is what Clark was describing at first. This is the grid the way it was developed in 1950, a one-way street, one-way delivery of power from the generator on the left to the consumer on the right.

Now, we're trying to take what we've got with the Internet, which is communication in two directions, back and forth, with a lot of intelligent hardware and software involved. We're trying to merge these two systems; a huge piece of the smart grid is communication back and forth. We need to know what's going on to adapt. We need to know what we are doing in locations everywhere from the generator to end use. The benefits include utilizing things like demand response and dynamic pricing. You can actually tell the consumer: "By the way, the price of electricity is going to be doubled tomorrow because everyone has their air conditioning on and their electric vehicle plugged in. There are many different things we can do once we have information going back and forth.

> *We're trying to incorporate what we've got with the Internet, intelligent hardware and software and two-way communication; a huge piece of the smart grid is communication back and forth.*

Some utilities have already started doing this and some customers are saying, "I am sorry, I don't want that. I don't want to because of a health condition," for example. They don't want someone else to be able to increase the temperature of their home. That's okay. They can opt out... but they're going to have to pay extra for the utility power they use during that block of time.

That's just one of the challenges for the utility, the consumer, the regulators, and everyone else to figure out. What's the best way to do this so that we don't interfere with people's ways of life, yet still provide power as everyone needs it.

Some customers are saying, "I'm sorry, but I don't want that."

So there are a lot of different things happening. There are a lot of different technologies being used. The utilities that are part of the demonstrations projects get together three times a year and try to figure out the strategic topics to work on to make their part of the grid smarter.

Some are things like conservation voltage reduction, distribution management systems, sending information back to CAL ISO or the other system operators so that can see what's going on in the system. They know the capabilities they have and the situation, so they can make adjustments to keep everything operating properly.

We know we can do a lot of energy storage. How expensive is it? Where do we put it? How do you operate the storage?

And then you've got, as I said before, consumer behavior. Some people will be very willing -- they want to save the polar bears and everything else. They'll say, "Okay. I'll do it." By allowing the utility to control some of their devices through the year, they may get a cheaper rate on their electricity.

There's going to have to be give and take. Regulators are trying to figure out how to do this, how to structure the rates so everybody gets a benefit. We keep the system running efficiently and everybody gets to use power when they need it.

There are also things like cyber security and electric vehicle charging. There are a lot of things going on that we need to figure out: how to handle it and how to ensure it works properly.

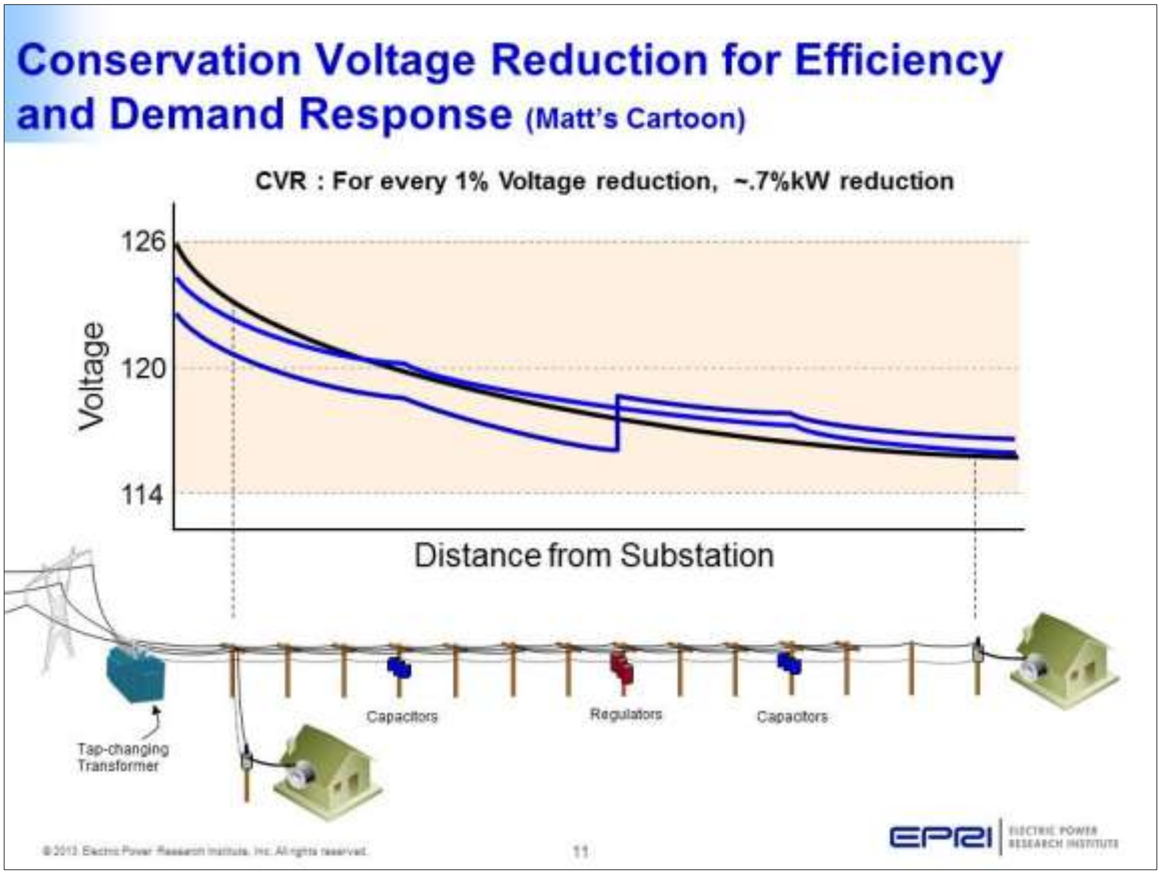

Conservation voltage reduction. Everybody wants 120 volts in their home. But because of distribution, it's actually a plus or minus 5% tolerance on that 120 volts. You can have as high as a 126 or as low as 114 volts in your home. In years past, the only measurement the utility had was the substation and it had to make assumptions about the voltage at the end of the line. They knew what it was at the substation because they have metering there and the telemetry goes back to the control center.

Now that we have smart meters, we can read the actual voltage at the customer's location.

Now that we've got these smart meters at the beginning and end of the distribution circuit, we can read what the voltage actually is. We don't have to make assumptions. We'll have a lot more information as these meters are installed. If you can reduce the voltage by 1% on the circuit, you get about a 0.7% kilowatt reduction in the load. That's basically what conservation voltage reduction is.

In times past, if there was a tap-changing transformer at the substation, that's the only control the utility had for the voltage all the way down to the end of the circuit. Now we can use that tap changer to actually reduce the voltage so we're saving energy, but still keeping the voltage at the end of the line within the margin that's regulated.

Now we can do things like put in capacitors along the line, which produce the VARs needed instead of generating them back at the generator or substation. We you can actually generate the VARs where they're used, at the distribution level.

When you produce VARs locally, you boost the voltage at the distribution circuit.

That allows the voltage to increase. When you produce the VARs locally, you boost the voltage down at the distribution circuit. That's great! To keep the end at 114, you can reduce the voltage at the beginning of the circuit, saving even more energy!

There are other things you can do, like putting in voltage regulators part way down the circuit. If it's a long circuit and there are lots of long circuits in the U.S., you can now boost the voltage at the back end of the circuit. There's no problem at the front end of the distribution circuit, but the voltage was lower at the tail end. With the regulator, boosting the voltage at the back end allows you to save more energy by dropping the voltage again.

So you can save energy that way. You can also save energy by having community energy storage. This is something not in the home; if your service is underground, you've got a par-mounted transformer somewhere down the block. It's that little green box that sits on a concrete pad. We're taking a look at having another green box right next to it or underneath it with either a lithium-ion battery or some high-tech battery we can use for the benefit of the customers as backup power.

You can save energy by having community energy storage.

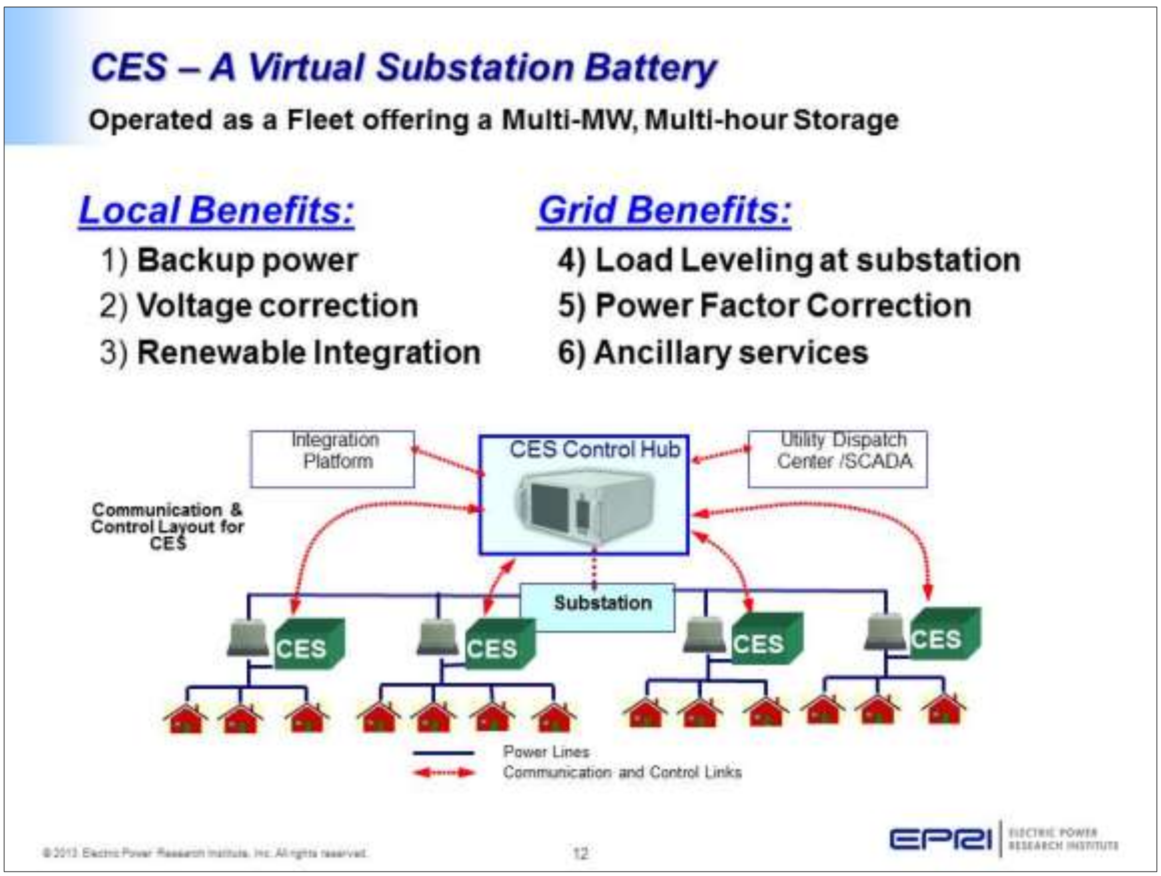

If there is an outage, you can use that energy to back-feed the transformer and feed the home or the commercial building, even though the utility power is out. You can use that energy for watt and VAR control. You can also use it to "smooth" out renewable energy variance, whether it's wind or solar, in the home or commercial building. When the clouds come over and the power output of the PV panel goes down, you can supplement it with the community energy storage.

Satish will be talking a little bit about this. There are other technologies that we can use for that.

There's also the benefit to the grid in that now you can use that energy stored out on the distribution circuit to level the load at the substation. You don't have to worry so much about power correction. You don't need as many capacitors because you can do a lot of the watts and power control through community energy storage. With the smart inverters, you can actually produce VARs or watts or a combination of whatever you need.

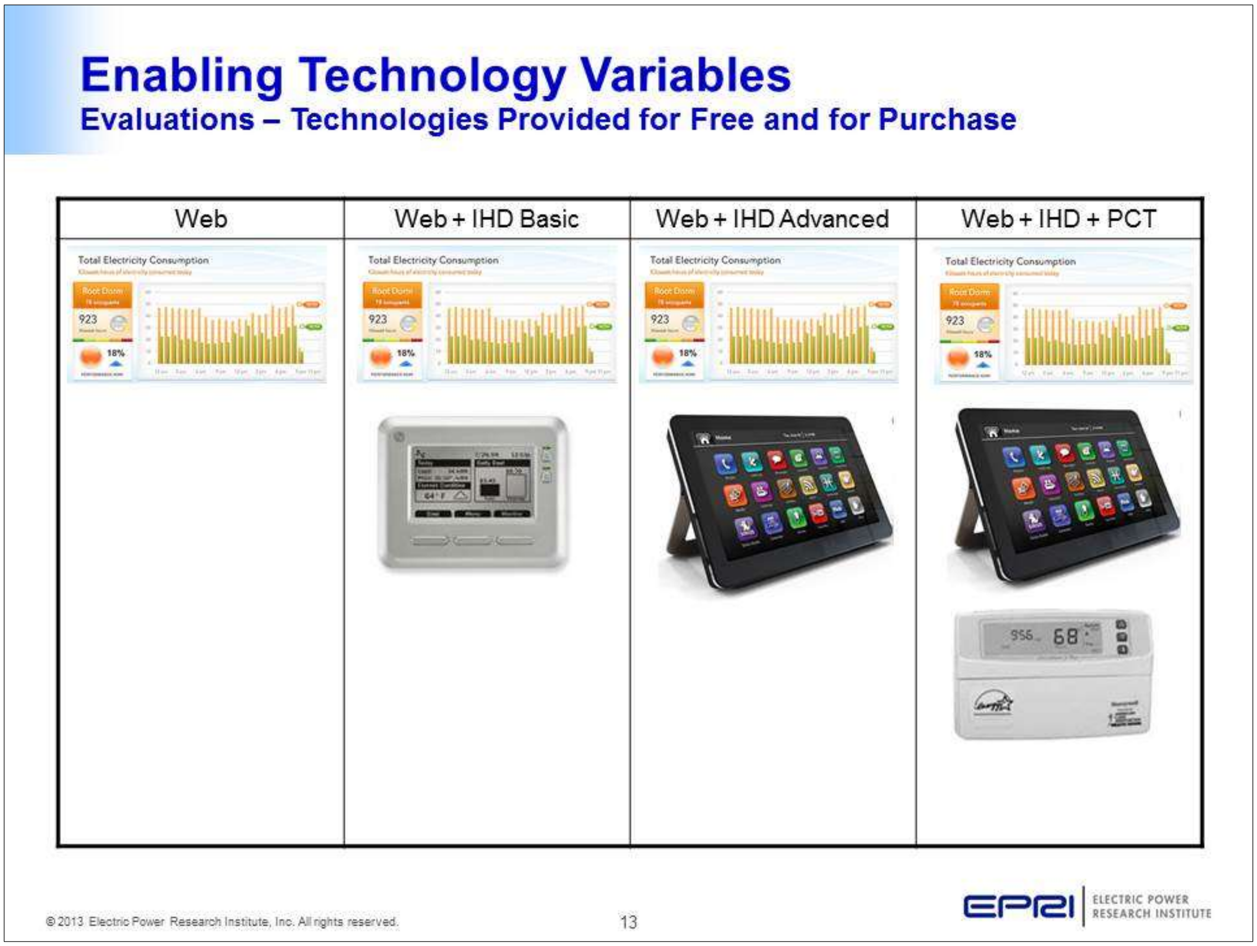

These are some of the technologies that are going into homes in these demonstrations at these 15 major utilities. Right now, there's basically nothing in the home. You have a meter; you get a bill every month. There's not much information. But you can have an in-home display, a basic display to show your current usage. It will tell you usage and sometimes it will tell you the outside temperature and anything else. It will tell you, possibly, what your rate is, exactly how many cents per kilowatt hour you're being charged right now. This is information you can use to reduce the load in your house.

A programmable, communicating thermostat will allow Con Ed, or whatever, to go in and control the thermostat setting. This is demand/response.

Things like an advanced in-home display also gives you things like Internet access. You can actually communicate back and forth to the utility. Then you start putting in smart devices, such as a programmable, communicating thermostat. This will allow Southern California Edison, Con Ed, or whomever, to control the thermostat setting. If you allow it, they can control how much energy you're using in your air conditioning, and they could give you a cheaper rate during the rest of the year because you allow them, during the peak time, to reduce your load. This is a demand/response.

Now we've got communication back and forth and we've got information being stored. Now there are issues with cyber security and personal security.

There are a lot of different pieces. Now that we've got communication back and forth and information being stored, there are issues with cyber security and personal security. If, all of a sudden, the load in your home goes way down for a week, you're probably away on vacation. You're not there using your television and you're not opening your refrigerator door and having that cycle often. We need security on that information.

Where is the data? Who owns the data?

All this information, all this communication…. where is the data? Who owns the data? There are issues being legislated in California; leading again, probably. We want to make sure that data is secure. The right people need to have access to it and nobody else should have access to it.

Why are we worried about plug-in electric vehicles?

I'll give you a thirty-second tutorial on electric vehicles. There are two different types: hybrid and there a total electric vehicle.

The old vehicle was a gas engine driving the wheels. Now you can either drive the wheels with a gas engine or a battery you charge.

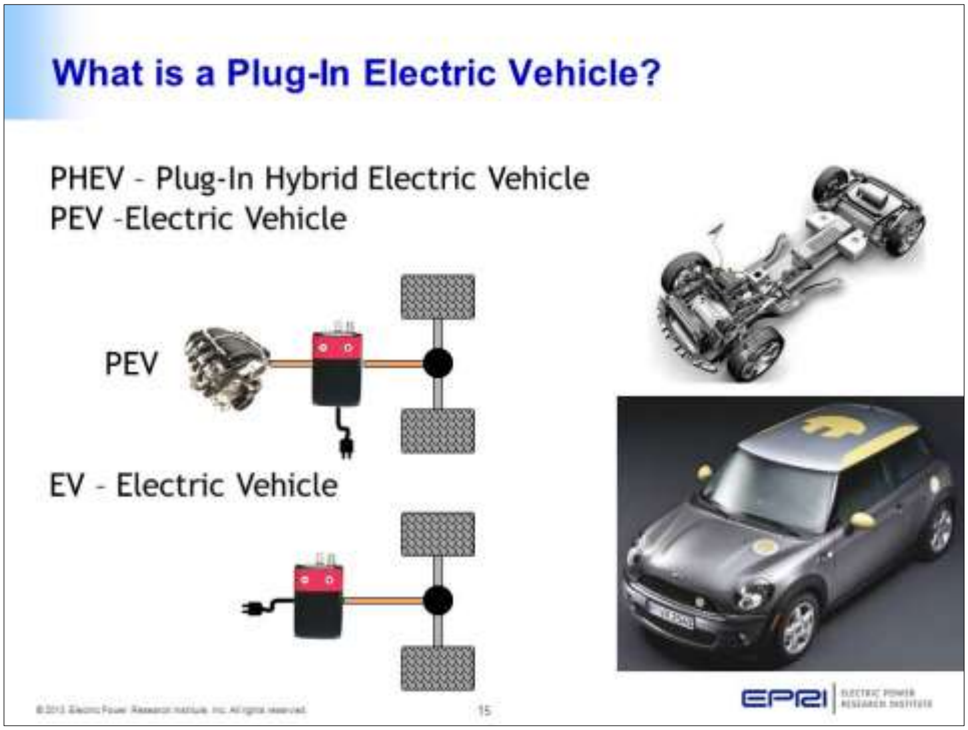

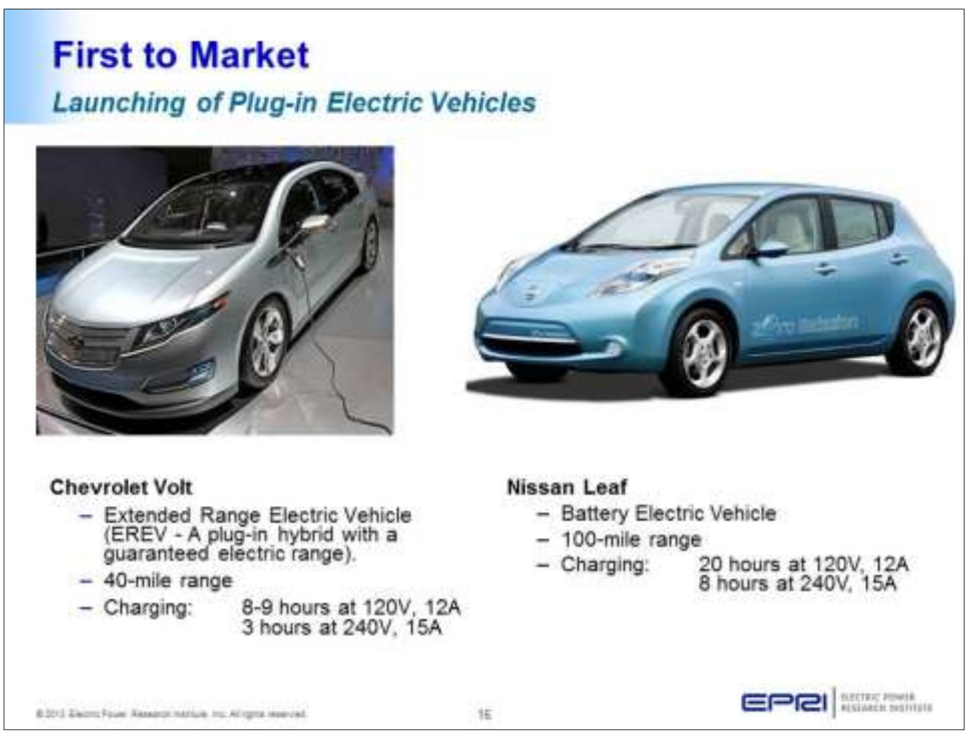

Now there's also a different type where the wheels are only driven by the electric battery, and the electric motor. You plug it in and charge the battery. If you are driving for forty or a hundred miles and you run out the battery, there's a gas engine that charges the battery that drives the wheels.

The total electric vehicle is just a battery, like the Nissan Leaf. Once you run out the battery, you're stopped until you find a place to plug it in.

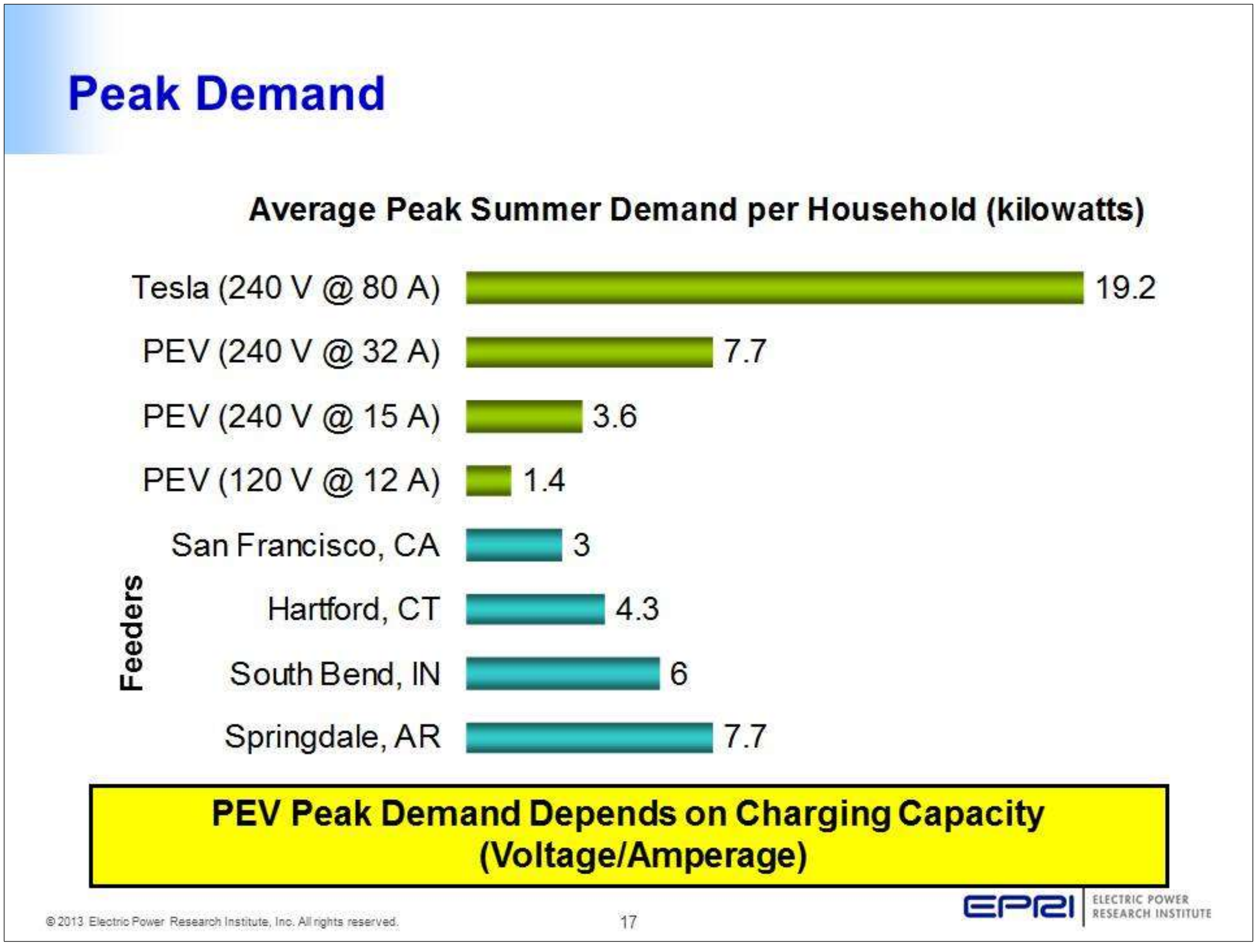

Why do the utilities care about this? These are the two leading electric vehicles right now. We worry about them because, if you look at the bottom half of this diagram, you see the summer peak demand in San Francisco, Hartford, South Bend, and Springdale... 3 to 7.7 kW.

An electric plug-in vehicle that you're charging at 240 volts exceeds the entire house in most of these areas. Even in Springdale, you're doubling the load of the house just by coming home and plugging in at 6:00PM. When you come home, you plug in your electric vehicle, and the house load doubles... or triples or quadruples, if you're in San Francisco.

Electric vehicles are going to have a big impact on the utility transformer, the utility distribution line, and everything in between them.

So these are substantial. If you're fortunate enough to own a Tesla, it's even worse. Not many of us have Teslas right now, but there are going to be more and more plug-in electric vehicles. This is going to have a big impact on the utility transformer, the utility distribution line, and everything connected between them. This is not a trivial thing for the utilities to plan around.

The smart grid. I've got three slides to summarize what's going on with the smart grid and why.

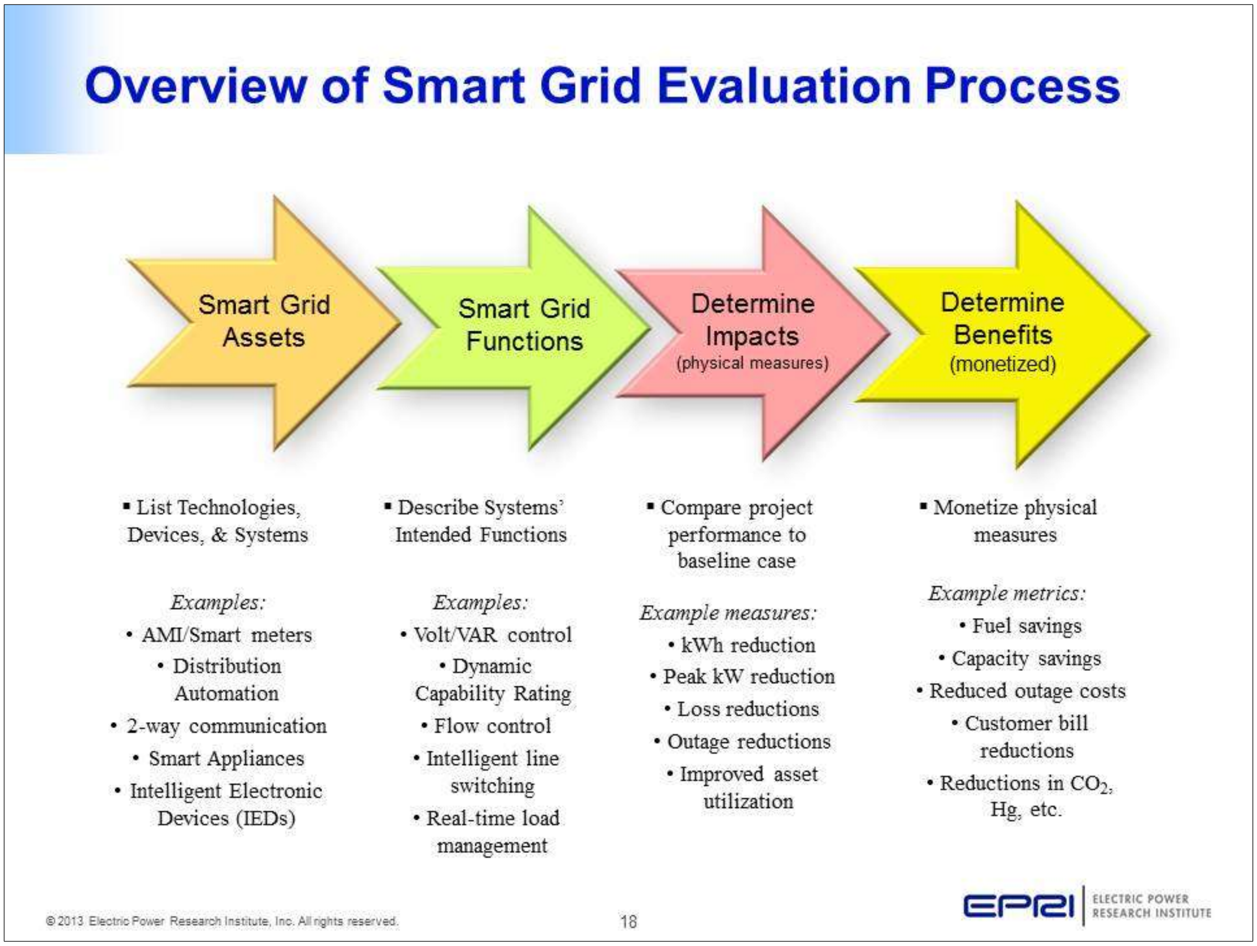

There are a bunch of smart grid assets. The meters that have been on our houses for 40 years are being replaced by smart meters that can communicate and send telemetry back to the utility every 15 minutes to say "Here's the voltage, how much usage, and kilowatt hours they use."

The meter that has been on your house for 40 years is being replaced by a smart meter that can communicate and send telemetry back to the utility every 15 minutes.

This is right now. This is a big change for the utilities. The utilities used to have a meter reader go out once a month to read the meters. Now they're getting 7,000 times more readings than they used to get every 15 minutes. Now we've got two-way communication. We've got smart appliances being installed in a lot of homes either by choice (or by regulation in California).

Smart appliances are being installed in homes by choice or by regulation in California.

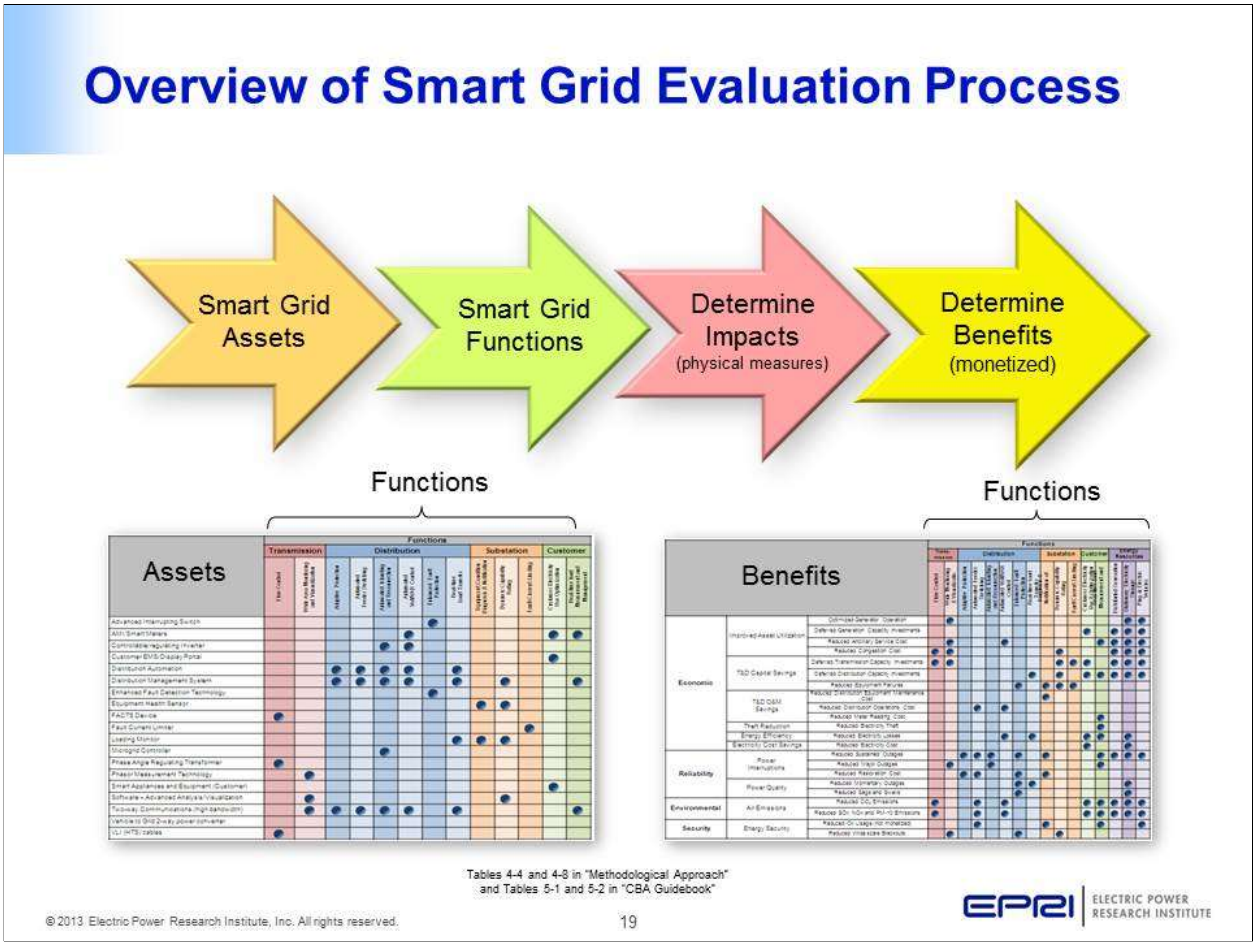

There's a lot happening and many different functions. Some can switch on and off at the utility's request and there are different things that affect VAR control. There are certain impacts on the home and the actual entire grid.

Hopefully, there will be kilowatt hour reduction. That's the plan. Peak load reduction, load losses, load reductions… what are the benefits? There will be benefits both to the utility and the consumer. The consumer will get a smaller bill and the utility will have more control over the loads, the distribution circuit, and everything else.

There will be reductions in greenhouse gas emissions. With so much going on, how do you figure out which things to do? There are a lot of assets to go and install in the field. There are different benefits to the utility and the consumer. We have to figure out exactly what's going on and where the benefits are.

As we do these demonstrations and get data back, we have to analyze that data to figure out if it did what we thought it was going to do.

So as we do these demonstrations, as we install this equipment, get it operational, get data back; we have to analyze that data to figure out exactly what happened. Did it do what we thought it was going to do? Sometimes you install it and it doesn't quite give you the impact you expected or it gave you a bigger impact than you predicted.

You may hear more customers saying, "I don't want this" or have a lot of people saying "Yeah. Look at the extreme weather events we've been having because of global climate change." Some places are colder than they used to be. Some places get more or less rain than they used to. So people are going to be more and more apt, as we figure this all out, to be involved in how they use electricity.

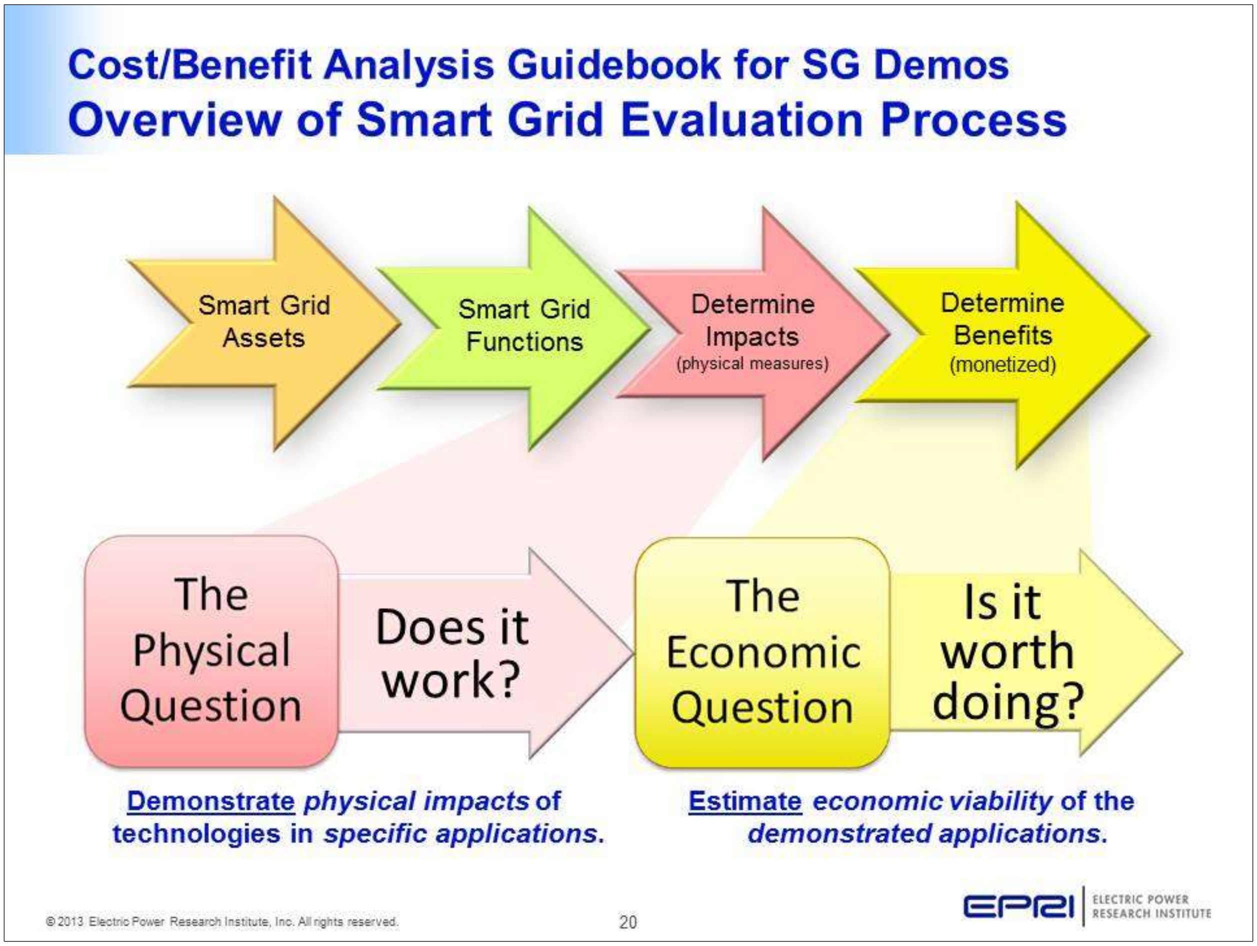

Then we're going to try to figure out how we did on all those things work. Some of them worked very well and were relatively cheap. Some of them worked very well, but were extremely expensive and it's not worth continuing this technology because we're not getting enough benefit from it. What kind of combination of technologies work well together? Which cancels the effect of another one?

Which cancels the effect of another one?

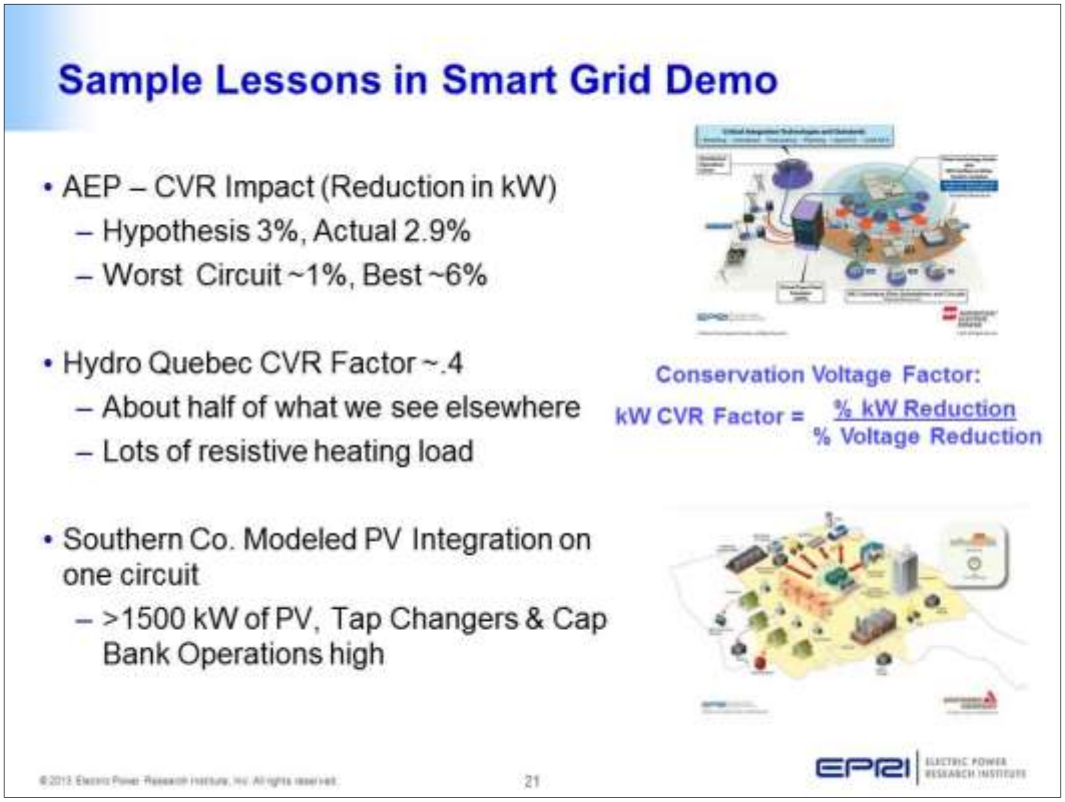

We've done conservation voltage reduction in a variety of utilities: one of them is American Electric Power and the other is Hydro-Québec. With AEP, we figured we'd get maybe 3% energy savings and we got about 2.9 in the demonstration. What we found out that some of the circuits only got 1% and some got 6%... the average was about what we thought.

Why is there such a wide variation of energy reduction? We got what we thought, but not the way we thought we would see it. Some savings were tremendous.

Hydro-Québec got about half of the conservation voltage reduction everyone else got.

In Hydro-Québec, because they are so electric energy intensive, there's not much gas heat. They use electric heat in the winter, and they have a long winter. They got about half the savings of everybody else. We're trying to calculate and assess the data. We believe it's because they have resistive loads for electric heating.

We're learning a lot from all this. We modeled the PV penetration and integration in one circuit in Southern Company and we're finding that, with 1500 kW photovoltaic installations; due to rapid fluctuation in the output, tap changers and capacitors in the tanks are switching on more frequently. These tap changers may last for forty years. In these substation transformers, these huge transformers are about half the size of this entire conference room. These are huge transformers. Sometimes we build them and expect them to last for forty years. If we can't find a way to smooth the output of the wind and photovoltaic output, it may have negative effects we have to consider. Can we get rid of those rapid fluctuations through energy storage, for example, and have the batteries provide the smoothing instead of having the tap changers switch on and off constantly and wear themselves out?

If we can't find a way to smooth the output of wind and PV output, it may have negative effects.

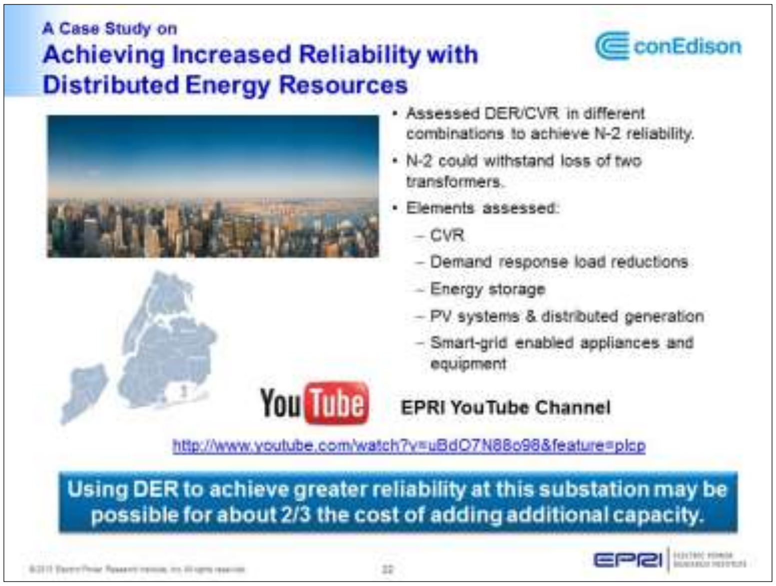

It's not an easy situation. Unfortunately, we don't have Internet access. There's a YouTube video you can access here in the conference room to when you get the report at the end of this.[1] We see what they've done, through utilizing some of these new technologies. They've been able to, instead of building a new substation or building a new generating plant — by using distributed energy resources, all of these demand/response load reductions, storage, PV installations, smart appliances, having control of everything -- they were able to achieve the same reliability and the same power output using only two-thirds of the cost of adding capacity.

Adding capacity to places like Con Ed is very expensive. Real estate is very expensive. Construction costs in Manhattan are very high. They're looking at ways of doing the same thing with distributed resources within commercial buildings, residences, and everything else. They're working very hard on that.

Commonwealth Edison found that, by having a lot of the control over appliances and pool pumps and everything else, they can get up to a 20% savings in energy usage.

Commonwealth Edison is working on things like home energy networks (HANs). Here's the programmable thermostat; there's the home area network. They're finding out that, with the normal load, the peak comes up in late afternoon. They're finding out that, by having a lot of the control over these appliances and pool pumps, they can get up to a 20% savings in energy usage in some of these facilities by utilizing things like peak-time rebates. If people allow control of some of their devices, the utility gives a rebate.

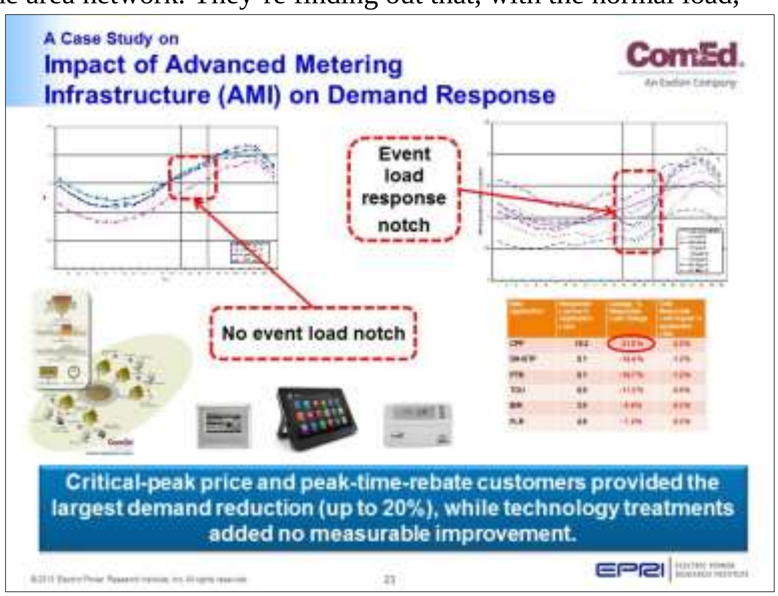

Things like this can help. If we can rely on this, people like CAL ISO, if they know that by calling on this, the demand response, there can be a reduction in the KW usage of all these plug loads -- and there's millions of plug loads out there in homes -- if they can get that, there's a very real benefit to both the utility and the consumer.

[1] http://www.youtube.com/watch?v=uBdO7N88o98&feature=plcp

So we try to figure out what kind of research plan we need. Do these technologies work? Do they give us the benefit we expect at a reasonable cost? All of this has to work together or the grid's not going to get very smart very quickly.

All of this has to work together or the grid's not going to get very smart very quickly.

There are a lot different things that we're tracking, everything from temperature and weather-related events to whether customers are willing to utilize a peak-time rebate. Can we get customers enough information for them to make reasonable decisions? Will they allow the home to go on completely automatic control by the utility? Maybe, maybe not. We'll see. Customer behavior is a huge piece.

Will customers allow the home to go on completely automatic control by the utility? Maybe, maybe not. We'll see. Customer behavior is a huge piece.

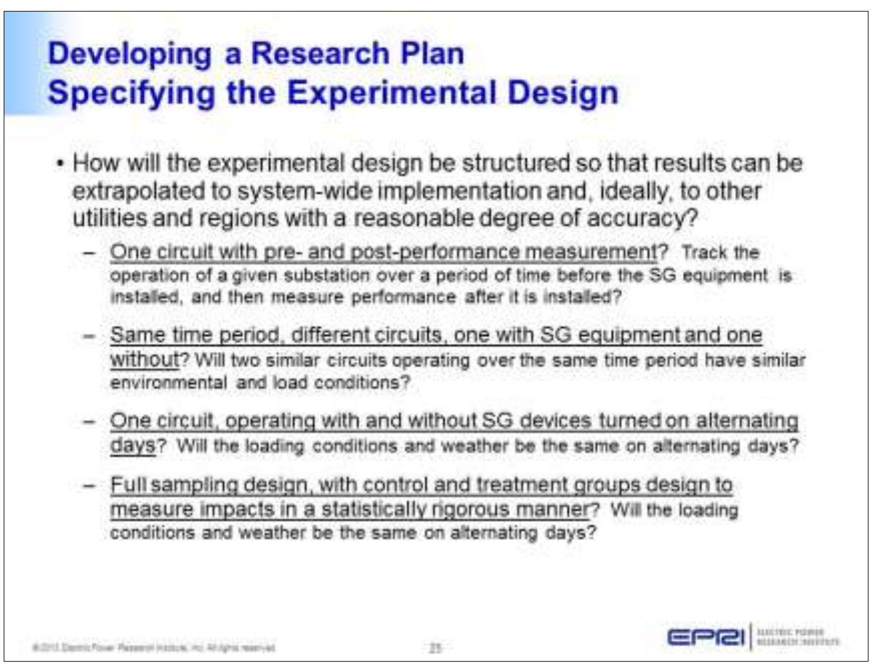

So what can we do? There are a lot of smart grid devices. As we control the voltage down the distribution circuit, if you remember the conservation voltage reduction, we're trying to reduce the voltage along the entire distribution circuit to get energy savings. Instead of everybody operating at 120 volts, everybody might be operating at 114 volts. This means that any voltage dip is going to be even more severe to the plasma TV or refrigerator or whatever. We have to make sure that the power supply you guys design and develop understands what’s going on with the smart grid and can work well with the utility and in the most cost-effective and efficient manner possible.

We have to make sure that the power supply you guys design and develop understands what’s going on with the smart grid.

We’ve just started these research projects and some of them are a year, year and a half, or two years down the road. Some of them are just starting. Southern California Edison is installing their equipment right now. They’re trying to retrofit that faculty housing and figure out what kind of technologies are working well, which ones are not working at all, and which ones work together well.

So that’s the end of my presentation. I think we’re going to have a question and answer period now.

Thank you.

Section III
First Panel Discussion

First Panel Discussion — Dusty Becker, Clark Gellens, Dr. Khaled Abdul-Rahman

DUSTY BECKER:

I'm not going to MC this in any fashion, but the guys are up here, so if anyone has any questions, comments, observations, or anything, throw them out there. Remember, this is a workshop. The idea is to exchange ideas. It might even be something new, maybe we need to study something, we have to look at something different, or there is something else we need to do. So feel free.

QUESTION:

My question is: are you studying what people are doing in other countries? A lot of the talk is very US-centric. In my networking with people, this guy from Germany was astonished that we were having rolling power outages in California. We still have planned power outages and he was like, "In Germany, we never have any power outages."

DUSTY BECKER:

Then again, there's India.

CLARK GELLENS:

It's a very good point. It's actually quite surprising when you learn that we're pretty far along. The assumption made is that we haven't when people ask: "Why don't you go look here? Why don't you go look there?"

The most advanced is probably Europe. In terms of the grid stuff, the European Union is driving the wagon and forcing the formation of something called ENTSO-E, the European Network for Transmission System Operators Electric. There's also a gas one and they have a fairly comprehensive research program... some elements of which are exactly what my colleague was referring to in terms of understanding best how to integrate things like wind. What are the elements?

He couldn't cover everything, but the issue of forecast. You see those rapid ramps both up and down. Well you can imagine a weather front coming through and, boy, he wants to know about that. He wants to know about that in very short order. He wants to know an hour ahead and minutes ahead and several times a minute. So that's a good place to look.

Another good place to look is Australia, which is, I wouldn't say ahead of us, but certainly has some pretty wide-ranging market rules that they've put into effect. We've been keeping an eye on that.

Japan is fascinating right now. Why are they fascinating? Because until a year ago, when you talked about renewable energies, they were very passive. But all of a sudden, there's a huge push as a result of the incidents that the earthquake and so on caused. There's a concern to move away from nuclear, a big push for solar, a huge push for efficiency. In Japan, you can buy a water heater that gives you five times more energy output, hot water energy out, than electricity… in using a heat pump with CO_2 as a refrigerant.

Then there's China and India, where high-voltage systems are being deployed. It's not that we wouldn't know how to do them, but the political will isn't here right now. Nor is, necessarily, the space. China just doesn't seem to care. If they want a line to go from point A to point B, somehow they arrange to confiscate all the property. We would have a more difficult time doing that.

DR. ABDUL-RAHMAN:

I just want to add that from a grid operation perspective, we have also been looking into those experiences of other countries. We have people going and visiting the control centers of Germany and Spain… and many of the challenges that we are talking about, they are also facing.

The solution to the problem is a little bit different. Some of the countries tried putting requirements on how buildings connect to the grid. They end up putting those requirements in after they experience issues.

I think we're trying to leverage those experiences and we're seeing we need those things identified up front so that we don't have to repeat the same tragic result -- we don't want another hard way of learning how things are working.

Now Germany, for example; they have a renewable fleet. But every country has its own special kind of situation — where resources are located, the political views of the country, and the diversity of the resources. Germany has the advantage of being in Europe. I think there is a driver for all of Europe to create power. That means you have a bigger area in an over-generation condition when you have too much wind, for example, that is above your load. You cannot continue to have over-generation. So they export that power to the neighboring countries. Having a bigger area reduces the variability of the resources.

CLARK GELLENS:

At a negative price. During the peak day in Germany, they have to pay the countries in the south to take the wind energy… then they sell it back to them.

DR. ABDUL:

That is correct. So these kinds of things are actually triggering our thinking here. For example, in the west, we are thinking that if we want to do more penetration of renewable, we have to think regionally. Maybe inter-regionally. We have to think outside.

How do we poll? Not a small area; it has to be a bigger area; the biggest it can be. That's why, for example, ISO is another great operator in the Midwest, Midwest ISO; they have a bigger footprint. They are still not seeing that much variability because, overall, you have over-generation on one side and under-generation where they need to have generation. So overall, balanced. As I said, the objective is to balance the supply and the demand.

QUESTION:

Most of the attendees at the conference are well aware that; if you have an efficient load that draws constant power, it will have a negative input impedance. Then, furthermore, in the grid, any source that is constant power is going to be high impedance. Those two things together look like it's a major instability. I'm wondering how many people, who need to know that, do… and do they understand it?

CLARK GELLENS:

Not very many. One of the things I need to do next week is go talk to a bunch of senior executives about what's going on in end use. In fact, it's under the "What's The Future Demand of Electricity from Grid-Related Services" heading.

Among the many variables in that equation are some you already heard mentioned, photovoltaics, increasing efficiency, and the like. One that is often ignored is that we have, as an industry, traditionally served analog loads. These loads are relatively passive. You change voltage and you get more or less a corresponding reduction in demand. Dennis showed some data like that. It's not perfect, but more or less.

Then you put a smart device in the middle, like the devices you folks all are very interested in, and those devices react to changes in voltage. For example, the simplest example is an adjustable-speed drive mechanism in front of a motor. In an inductive motor, years ago, if you changed voltage by 1%, you're probably going to change the power draw by 1%. Now with a smart device in between the adjustable-speed drive mechanism says, "Oh, you've reduced voltage? I'll maintain constant torque and adjust that for you. At the same time, I'll change the input impedance in that device." Not well understood.

Section IV, Session 2
The Evolving Grid

Uncontrollable Sources and Controlled Loads

DUSTY BECKER:

Our next segment involving the grid is presented by Dr. Jian Sun. Dr. Sun received his Ph.D. from the University of Paderborn in Paderborn, Germany, in 1995. He was a Post-Doctoral Fellow with the School of Electrical and Computer Engineering, Georgia Institute of Technology, from 1996 to 1997. He worked in the Advanced Technology Center of Rockwell Collins, Inc., from 1997 to 2002, where he led research on advanced power conversion technologies for aerospace applications. In 2002, he joined the Department of Electrical, Computer, and Systems Engineering at Rensselaer Polytechnic Institute in Troy, NY, where he is currently a professor. Since 2010, he has also been the Director of the New York State Center for Future Energy Systems (CFES), which conducts research in the broad area of energy; including wind, solar, energy storage, smart gird, and smart buildings. His research interests are in the general area of power electronics and energy conversion, with particular emphasis on modeling and control, as well as applications in renewable energy and aerospace. He has published more than 150 journal and conference papers on these subjects and holds nine US patents.

Dr. Sun currently serves as the Editor-in-Chief of IEEE Power Electronics letters, and was the guest editor for the IEEE Power Electronics Society (PELS) and an AdCom Member of both PELS and the IEEE Systems Council. He also served as General Chair of IEEE COMPEL '06 Workshop, Chair of PELS Technical Committee on Power and Control Core Technologies, Co-Chair of the IEEE 2012 ECCE Technical Committee, and Vice Program Chair of ECCE 2009-2011.

DR. JIAN SUN:

Thank you for the introduction.

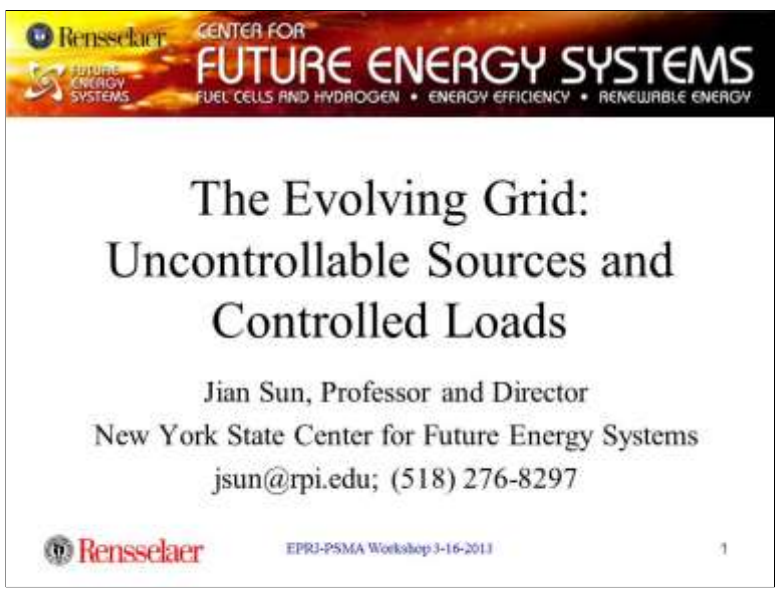

It seems that I'm the first one to speak on the power electronics side. I think we heard a few questions also related to the subject of source impedances issues that come with power behavior. I will try to cover some of these.

Just a general comment, I think it's a great idea to bring the power electronics community and the system community together.

My own background is in power electronics. A few years ago I started to dig into the power system aspects. I attended some of the IEEE and PES meetings. The first couple of times, I saw, even though we are all actual engineers, I saw two total different schools. We all understand power, right? Power systems, power electronics. We all understand power. We all use "MW"...but "MW" means two different things to the two communities. For the power electronics team, it's often milliwatts; and, for the other side, it's megawatts. So I think it's a great idea to bring these two communities together.

I also saw quite a big gap between the two subject areas. Part of my research in the last few years has been to identify the gaps and try to fill some of the gaps. That's sort of the quick introduction.

The system of synchronous generators will synchronize to each other as long as you operate within the limits of the capacity of the system.

We've heard quite a few discussions about the traditional grid and the generation centers. I say fully controllable in this slide, although it's not exactly true. We have the grid and it can generate to this grid. The control is centrally coordinated, in what we call dispatch. The system of synchronous generators have the magic ability to synchronize to each other so you don't have to worry about the synchronization as long as you operate within the limits of the capacity of the system.

On the loads side, the loads are variable. They are unregulated. We talk to our system engineers and the only term they use is megawatts. For the system analysis, we lock them in, aggregate those, and we classify into resistive, inductive, and constant-power load. I mean, that's sort of how we deal with it. We didn't really look to the individual loads to try and regulate them.

In today's power system, imbalance is absorbed mostly by the kinetic energy storage generators.

What about transient imbalance? Stability is all about balancing the power. You can never have perfect balance between the loads and the generation. In today's system, the transient imbalance is absorbed by most of the kinetic energy storage generators. The nice thing about this single generator of kinetic energy storage is that we don't need to worry about it. Any unbalance, within limits, is absorbed automatically by the storage in these big generators. There's no control required. I'm talking in principle, not in extreme cases when you run out of the boundary and you have instability issues.

Evolving sources and loads. That’s sort of the assignment for my presentation and these are pretty trivial after a few presentations this morning.

Renewable energy sources are increasingly variable and distributed solar and wind. They are like negative loads because they are uncontrollable and variable.

With renewable energy, the sources are increasingly variable and distributed largely because of renewable energy, solar, and wind. Look at it from the power system standpoint: they’re more or less like negative loads in the traditional sense because they are uncontrollable and variable. They’re more like loads, but it’s inactive. You still want to balance the power, maintain power balance, so you now have to manage the loads. Loads have to be managed and they sort of behave like… well, ideally, they behave like matching sources in the traditional way so you can regulate them and control them.

Storage is important because you can never match generation and consumption perfectly. With less and less kinetic energy storage, we have to add other storage into the system.

Storage is important because you can never match these generators and consumption perfectly. In any given moment, you must have storage with increased use of renewable energy. We have less and less kinetic energy storage; we have to add other storage into the system.

For a battery to respond to the grid, it must be controlled all the time.

I think what you call the difference is the battery, the other means of storage. They’re all storage, but they’re not the same because these storage devices, to actually be able to respond to what’s going on in the grid, must be controlled. You have to control them all the time. You can’t just leave them and expect them to work.

Constant power loads is another question and I think we heard about some misconceptions about that. So I want to touch on that, as well, toward the end of the presentation.

So just quick look at how we control the system today and in the future. In this case, black is the consumption, the load; red is the production. If you look at the power balance, demand versus generation, the difference will have been absorbed by the storage, mostly in the form of kinetic energy.

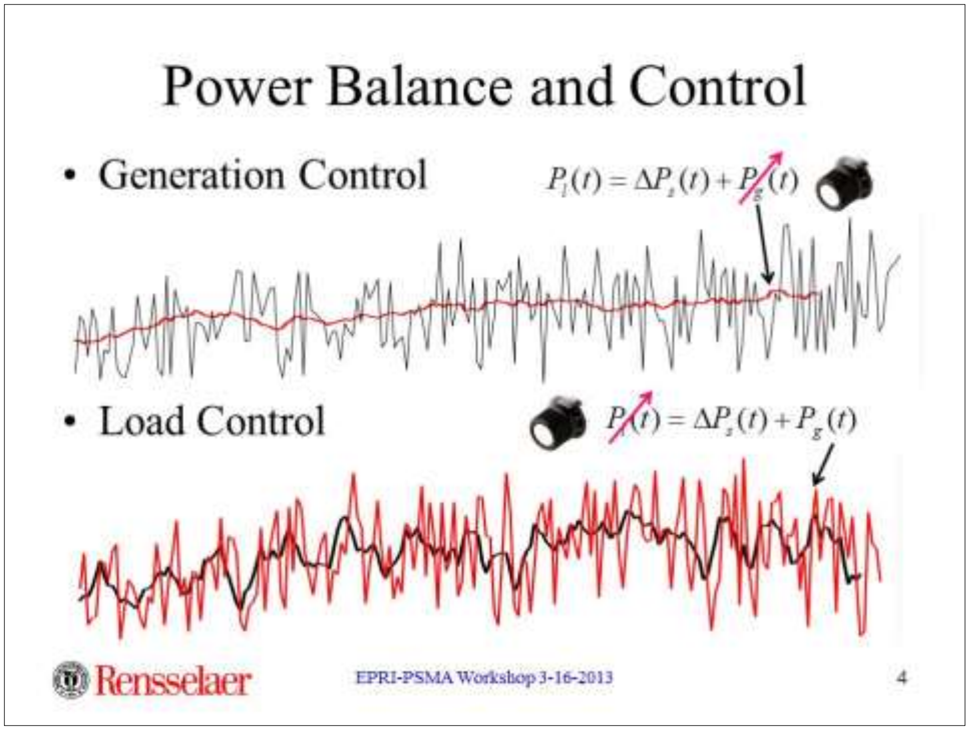

Today, we regulate production and track average consumption. The instantaneous imbalance is absorbed by the kinetic storage.

The regulation we have now is on the generating side. We put it out there, we regulate the production, sort of track the average consumption, and the instantaneous unbalance is absorbed by the kinetic storage.

That's how we regulate today. We don't have the means to regulate the generation with renewables and other variable generation. We need to put the nut on the other side, or maybe on both sides, to maintain the balance.

So put a control on the load side and the production becomes more volatile. The red is the production. Now try to regulate the load to follow that. This is sort of an extreme case, obviously. In reality, you always have the control on the generator side and the difference, again, has to be absorbed by the storage.

Mathematically, this is the same equation. You just move one term from one side to the other.

Nothing has changed mathematically. You just move one term from one side to the other. But there's some fundamental difference when you try to regulate the generation from a few sources, as we do today, to the consumption on the other side of the equation.

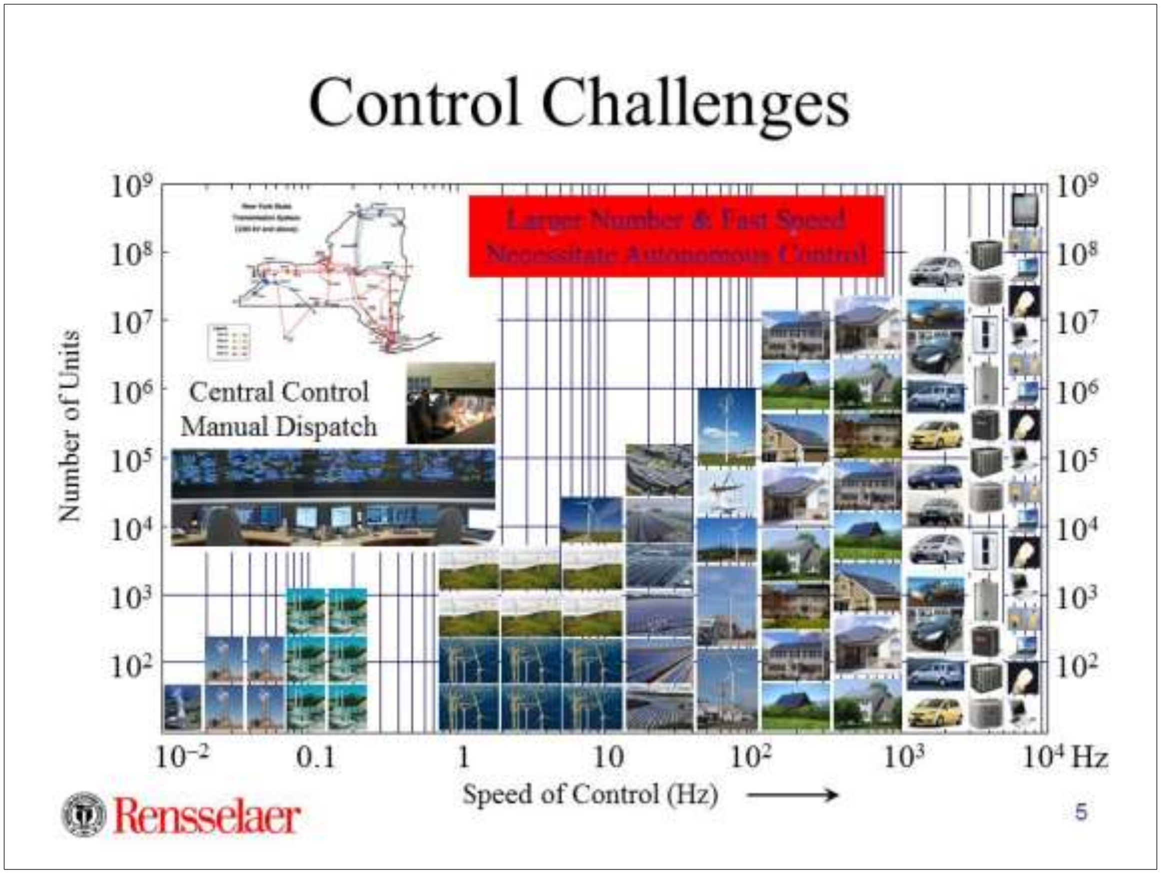

To illustrate that, I put this diagram together. The horizontal axis shows the speed of control, the speed of reacting on different devices. The vertical line shows the number of units within a system.

Today we have a few, and when I say few, we're talking about hundreds of those. Still, a relatively small number compared to the loads. A few generators, power plants, and the control is slow because you're dealing with large mechanical generators and the bandwidth is normally in the sub-Hertz range. Because they are slow, and we can control them from a central location, we can do the magical dispatch of production and try to balance the load with the generation. We consider future renewable generation from wind farms, individual wind turbines, solar farms, roof-top solar panels; then into electrical vehicles and load/demand response of all different kinds.

As you move up this slide, the number of units goes up exponentially, while the speed of control required also goes up. The lower the power, the faster these devices can respond.

We cannot control distributed power sources in the same way we do today, because of the large number and the speed of response.

Look at this picture. Obviously, we cannot control them the same way as we do today because of the large number of units and the speed of response. What this entails is that a larger number of fast-speed response, would necessitate autonomous control. We can't do this central management anymore.

A large number of fast-response units necessitates autonomous control. We can't do it with central management anymore.

Most specifically, let’s compare, again, control of traditional generator versus a wind turbine generator.

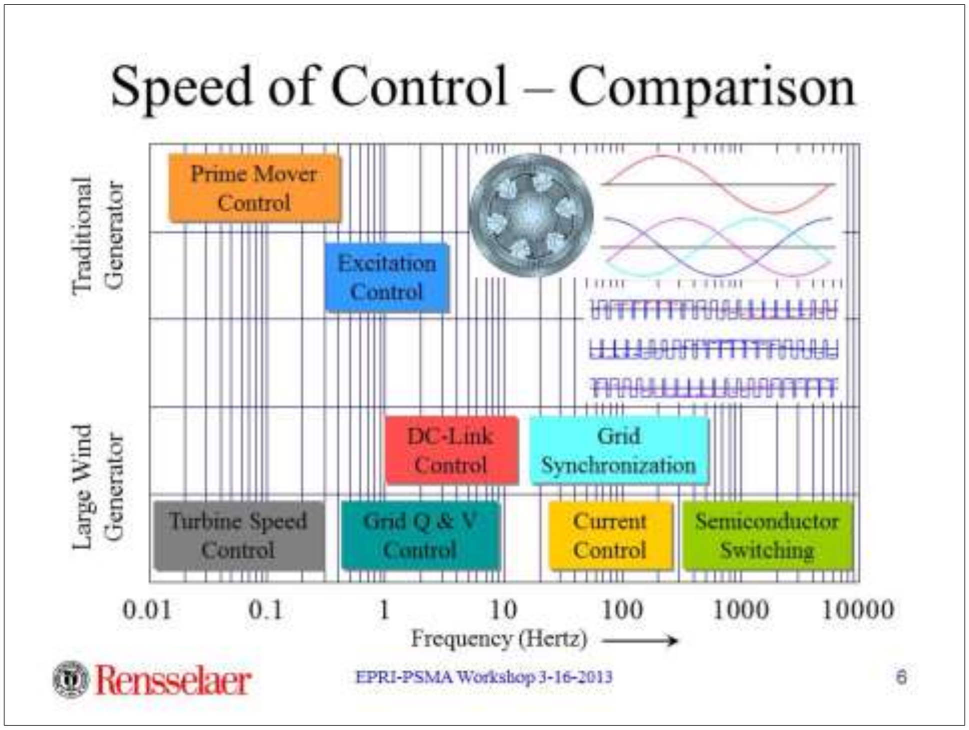

Here are the different control functions in each case and then how fast each control function is designed. In the traditional generator, the frequency is regulated by the speed of the prime mover and the bandwidth is in the sub-Hertz, milli-Hertz range, because you have a larger mass.

Excitation control. The loading control can go a little bit faster. It’s electromagnetic sort of behavior there. You can be faster, but you don’t want to do that any faster than 60 hertz because; if you do that, you change the magnetic field significantly within every cycle and create distortions. So those control functions are rather slow as well.

The difference with larger turbines and other renewable sources is that you have compatible functions in the low frequency ranges to regulate the speed for maximum power tracking, controlling the dc link voltage in the generating four turbines, the reactive power, and the voltage at the grid interconnection. These are comparable to existing generating controls. A little bit faster because these are lower power.

More importantly, we have high-frequency functions above the fundamental frequency to synchronize with the grid and to regulate the current to follow the 60 hertz of the grid voltage. Then semiconductor switching is at a much higher frequency.

More importantly, we have these high-frequency functions above the fundamental frequency to synchronize with the grid 50 or 60 hertz to regulate the current to follow the 60 hertz of the grid voltage. Then the semiconductor switching of interest to PSMA is switching at a higher frequency.

What I’m pointing out here is; if you look at traditional generating versus these renewable sources, the high-frequency part is a very important difference that has not been given attention. In the past, in this system, what’s going on here is typically considered power quality problems. They are considered to be sort of a nuisance. You don’t worry about them because you have no control of that frequency range.

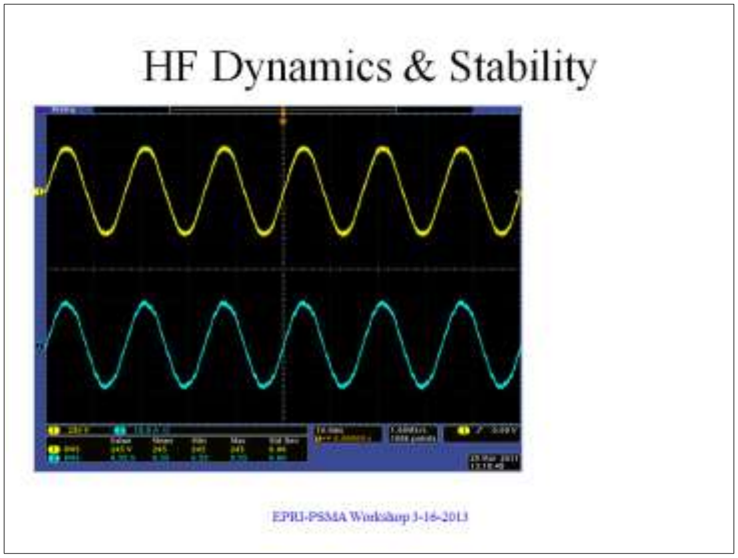
HF Dynamics & Stability
EPRI-PSMA Workshop 3-16-2013

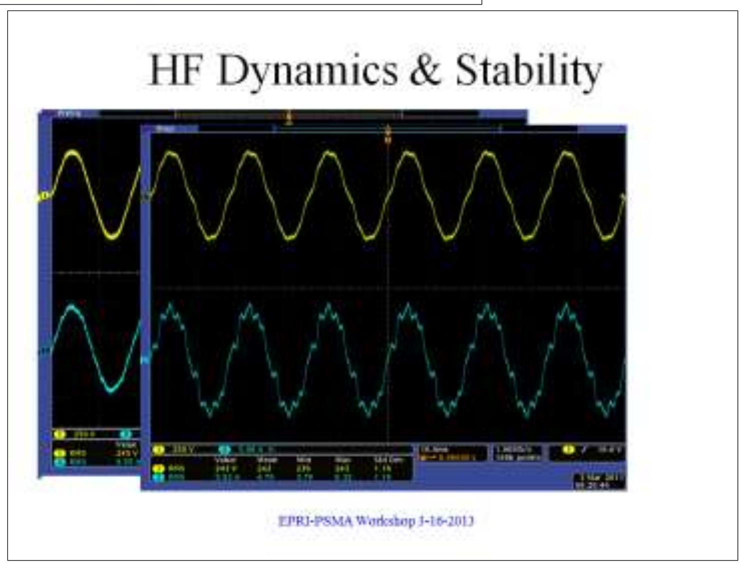
HF Dynamics & Stability
EPRI-PSMA Workshop 3-16-2013

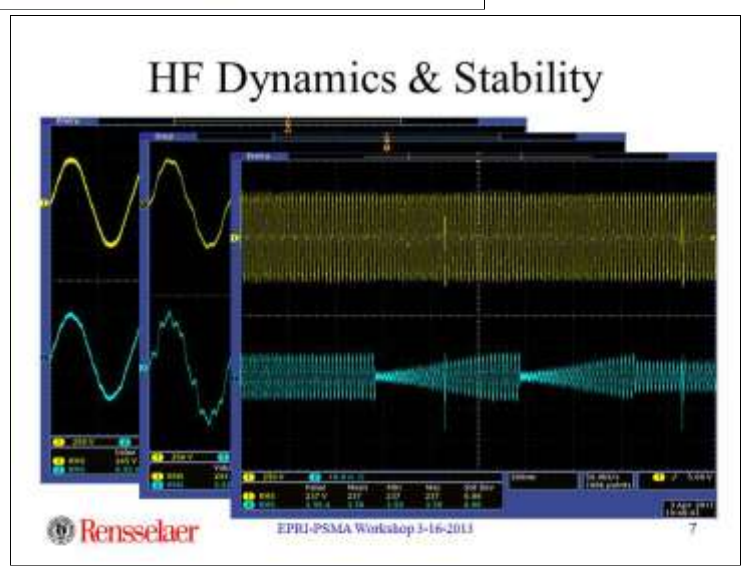
HF Dynamics & Stability
Rensselaer
EPRI-PSMA Workshop 3-16-2013
7

So what's going on here? Just some resonance? That will eventually die out. There is no active control. Stability is mostly about this frequency range and these are separate.

If you look at the future system, you see all these renewable sources. Resonance in the narrow range and different components will couple into the active control of these turbines. Then you have stability problems. That's an area we try to address in the interface between power electronic and power systems.

Just for example, I have a few images to show. The first is a single-phase solar panel, the solar panel connected to the grid in my lab. The initial graph shows a recognizable voltage grid voltage and current into the grid, with virtually no harmonic distortion.

In the lab, we have the ability to increase impedance of the grid just by inserting passive elements and doing some active control to simulate changing impedance of the grid. In fact, we have a programmable grid in the lab. As we increase impedance in the grid, you start to see harmonic distortion in the current of the inverter, and that also distorts the grid voltage.

Now this is the same device that we have. Nothing changes in the circuit. The only thing that changes is the impedance of the grid. You know, we're making the grid weaker in this case. If we keep doing that, the converters eventually start to become unstable.

This slide shows the same voltage and current over a longer time scale. You can see that the current cuts off after a while because there is too much distortion. The current loop cannot stay stable. So after a while; it disconnects and shuts down. We didn't unplug the converter, so the motor is still connected, still alive. So after a couple of cycles, the inverter is going to start up here, ramp up the current, and then it can only run for few cycles before it goes into this instability again.

What was considered a power quality problem in the past has become a stability problem. This is really where the power electronic and power system people have to work together.

So what I'm highlighting is that what was a power quality problem should be considered a stability problem. This is really where the power electronic and power system people have to work together.

So the implication on the power electronic side is, certainly, that we interface all these sources, storage as well as loads, to the grid. This has been a shift in focus.

Initially, we only thought about functionality for grid service and it can perform that function. For grid, it's certainly mostly ac-dc converting. In the last few years, the focus has been more on performance: efficiency, cost, power stability, and sources because we know how to perform these functions for best performance.

When these devices work in the power system, we can't just optimize it locally. They need to have predictable global effects on the power system.

I think that, in the future, we have to move one step further to focus more on the system impasse. With the device work in the power system, we can't just optimize locally. I think we have to sort of follow these and get more attention on the system impact.

Obviously, this is important because all of these power electronic devices provide the interface to the grid. They are the key to converting the grid-connected devices into an asset that you know you can control. That points to the need for local controls. Controls that reside locally on the individual units, but give you the predictable global effects so the power system people know what to expect from these devices.

You cannot talk with us without equations and mathematics. I'm getting into that part of my presentation.

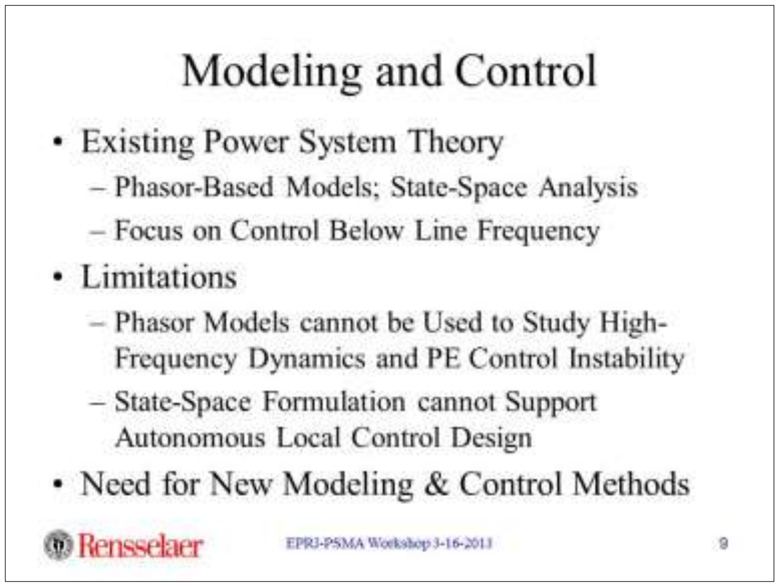

Modeling and control. How do we do that? How do we bring these two communities together? The power electronics people traditionally focus on circuit design and perform the functions locally to optimize the performance. The power system people look at millions of devices.

How do we bring them together to really enable autonomous local control that gives you the desirable, predictable, global behavior? If I look at the more fundamental side, look at the power system theories today; mathematically, are we disqualified? These devices are phase-to-phase models, so we formulate a big matrix to analyze system behavior. We look at algorithms for stability, damping, resonance, and other issues.

Certainly everything we care about is below the fundamental 60 hertz. That's really what the power system theory is about today. Then we take them, and we try to do power electronics control design, that framework. And to achieve our goal of supporting the system, there are some fundamental issues.

The phasor models cannot be used to study high-frequency dynamics and the stability of power electronics impacting with the grid.

Phasor models cannot be used to study high-frequency dynamics and the stability of power electronics impacting with the grid. Phasor models limit you to the fundamental frequency. If you go beyond that, the model is no longer valid.

The state space formulation also has limitations because; to form a state space, you have to have the knowledge of the entire network. So for wind farm development, you try to put a wind turbine here, and you try to analyze how the turbine will impact the grid. You don't have the information about the grid. You cannot form that matrix. That matrix is sort of proprietary to the system owners. So the state space formulation does not allow you to do that. And the local control design will be able to support the system.

That points to the need for new model and control methods. Fundamentally, this is a challenge. This is a gap between the two communities.

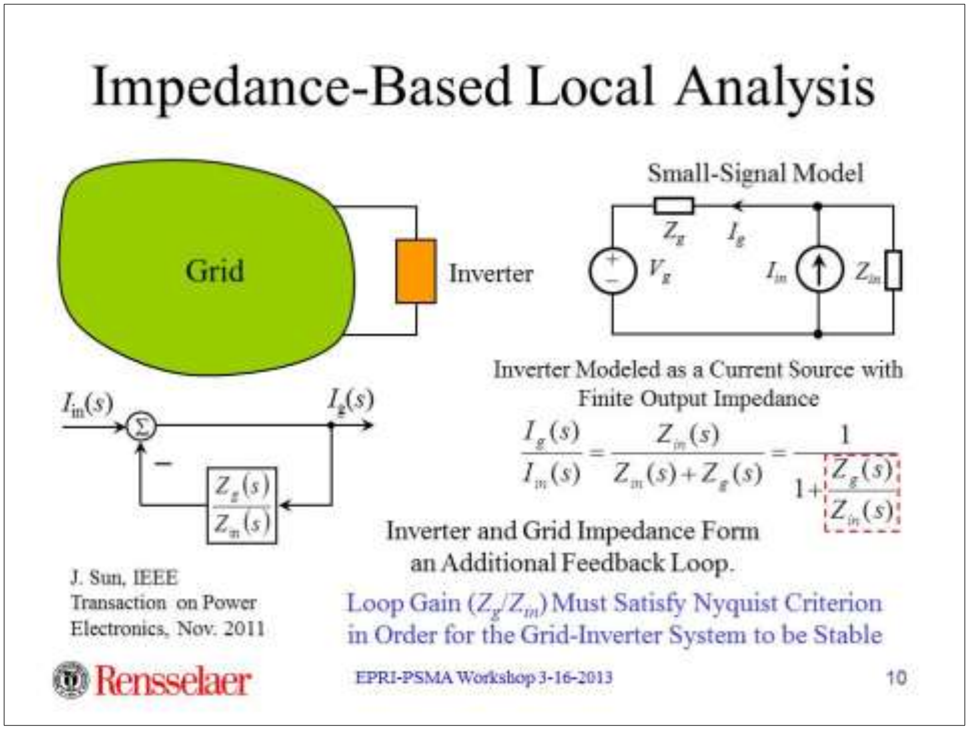

So one of the developments we've worked on in the last few years is to enable these autonomous local designs. We take a different approach. We take the sort of equipment manufacturing viewpoint.

We can characterize the grid by impedance.

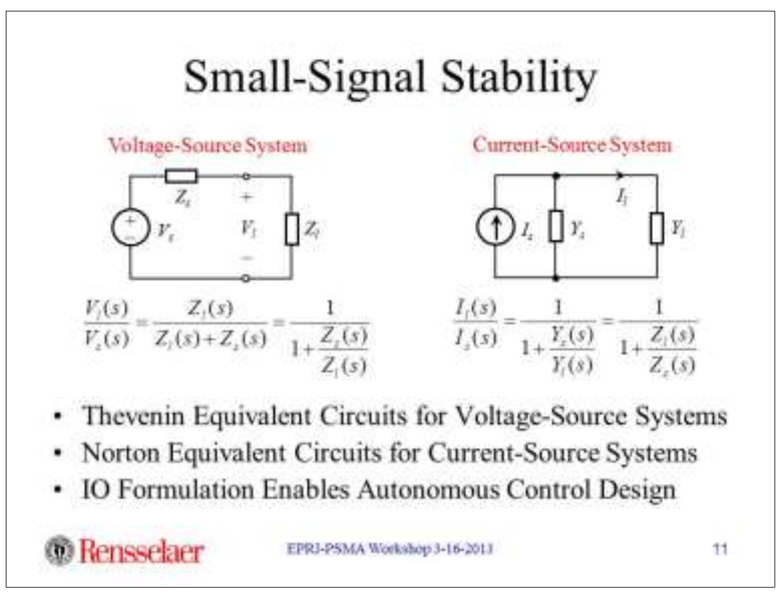

I have a grid network that I don't really know much about. I'm supposed to put something there to integrate wind and solar and store whatever by quick connective devices, a converter, a rectifier. I want to control these in a way that works well with the grid.

So our project is sort of the input/output of that thing. Instead of formulating a big matrix, I look into the grid from my interconnection point and I characterize the grid by impedance. This is not a new concept, but the application is sort of different.

I formulate a small-signal representative of the grid as an impedance or voltage source behind impedance. Most renewable energy sources are controlled as the current source, so we can model it as the current source. Obviously, you never have an ideal current source, so we model that by current source in parallel with an output impedance.

The ideal converter for this application should have infinite output impedance.

I heard a gentleman commenting on the output of impedance of these converters. That's what you do. In fact, I want to point out, the ideal converter for this application should have infinite output impedance. It's not that they have high output impedance. Ideally, you should have infinite output impedance to satisfy the requirement.

Based on this simple small-signal representation, you can formulate a dynamic model for the interconnected inverter grid system. What I want to point out is, by doing that, you actually formulate a local feedback control.

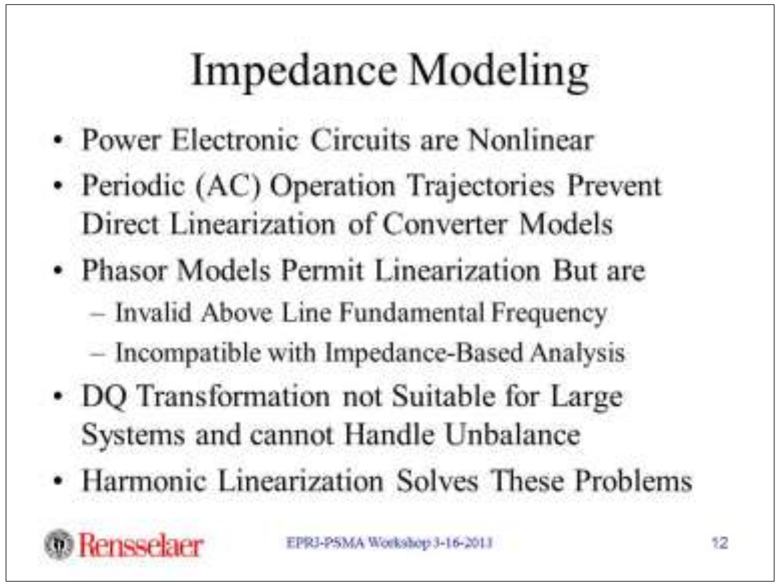

By connecting the converters with the grid, we effectively create a local minor feedback loop.

This is well known in electronics, but the application is quite different here. By connecting the converters with the grid, we effectively create a local minor feedback loop. Where the loop ends is the ratio of the grid impedance or the input/output impedance. You can design a converter to be perfect to be stable with ideal grid and the rest of it can operate without this inverter.

The problem is that when you connect them together, you have those additional minor loops that create a potential for instability. Instability depends on the ratio of the impedance, so that's a rule. So, again, it is zero, and you don't have any instability. The higher the gain, the more likely your interconnected system may become unstable. It comes down to the grid impedance versus the input/output impedance.

Ideally, you want the grid to have zero impedance.

Back to my comment. Ideally, you want the grid to have zero impedance; or you can make your inverter an ideal current source and that impedance would become infinite. Again, the ratio becomes zero. In reality, neither happens... so you have to deal with the instability problem. This is different at every location, so that's the design issue.

In the power electronics community, producing power supplies; we know this voltage source represented here. You have a power supplies, you model that source behind impedance, and you look at how that works with the load. There's a stability issue between the sources and the loads.

The current source representation I showed you for the grid connector is sort of a dual to that. You can look at the impedance ratio versus the admittance ratio. There's a lot of interesting material here, but for the lack of time, I'm going to skip these.

What I want to say here is when you try to apply this methodology to the design of these electronic power circuits, there are some fundamental challenges. This is an ideal sort of topic for a Ph.D. student. The problem is simply that every part of the electronic circuit we work with is non-linear, so impedance is only meaningful in the small-signal sense. That means you must do small-signaling analysis and linearization before you can have impedance models.

Instead of a dc operating point, which allows linearization, we have a sinusoidal trajectory.

The problem here is that, instead of a dc operating point, which allows you to do the linearization; we have a sinusoidal trajectory. How do you linearize a system along a time-varying trajectory instead of a dc operating? Mathematically, that is a challenging issue.

In today's power system theory, we deal with that again by working with phasor models. When you model a system by phasors, you can amplitude the phase angle and these are the quantities you work with. They become dc quantities instead of state, so you can linearize the phasor model. You don't really have the sinusoidal signals, voltage, and current to worry about.

The model you're likely familiar with is the DQ transformation. The idea is you have three balanced phases of quantities you transform into a rotating reference frame and they become dc qualities instead. So you do the linearizing there, do all of the analysis, and impedance output in that frame, then go back to the time domain.

We use harmonic linearization to linearize the design for small-signal analysis.

This works in some cases, but has a lot of limitations for many network applications. The one we've been working on is what I call harmonic linearization. We linearize the system, the model directly around the side of the trajectory, so you develop an impedance model you use to support the system design. That's a lot of the mathematics behind it.

A three-phase system in that framework is decomposed into a positive sequence system, a negative sequence system, and a zero sequence system.

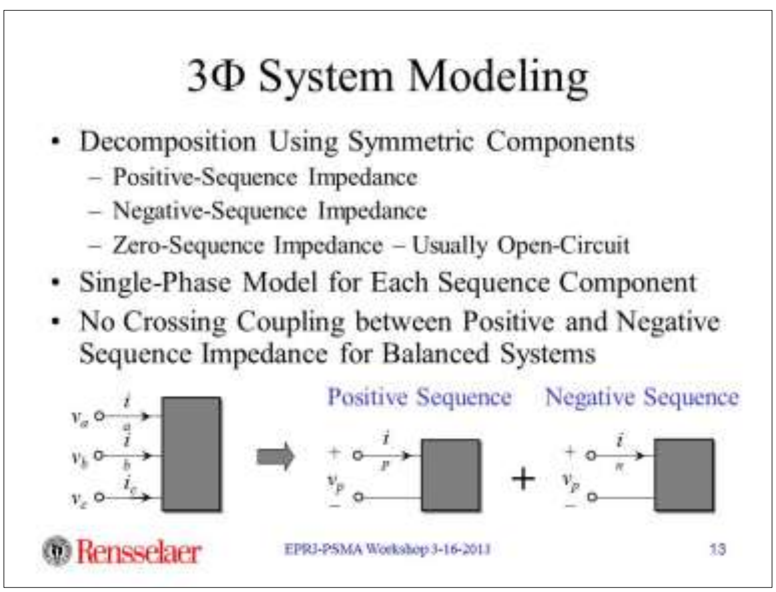

A three-phase system in that framework is decomposed into a positive sequence system, a negative sequence system, and a zero sequence system. This is similar to the symmetric analysis in the existing power system theory, but we generalize the impedance concept into a broadband frequency. Instead of just fundamental frequency impedance, the impedance here is automatic over a wide frequency range, going into the wide frequency range.

It's a mathematic tool and that framework is the three phases, decomposing to a positive sequence, a negative sequence, a subsystem. We can access the stability successfully in each of these subsystems and put them together and determine system stability. With power electronic, normally you ignore the common-mode EMI issues.

In my list of issues, the zero sequence current is zero. That means a zero sequence impedance is normally infinity. Normally you don't need to worry about that.

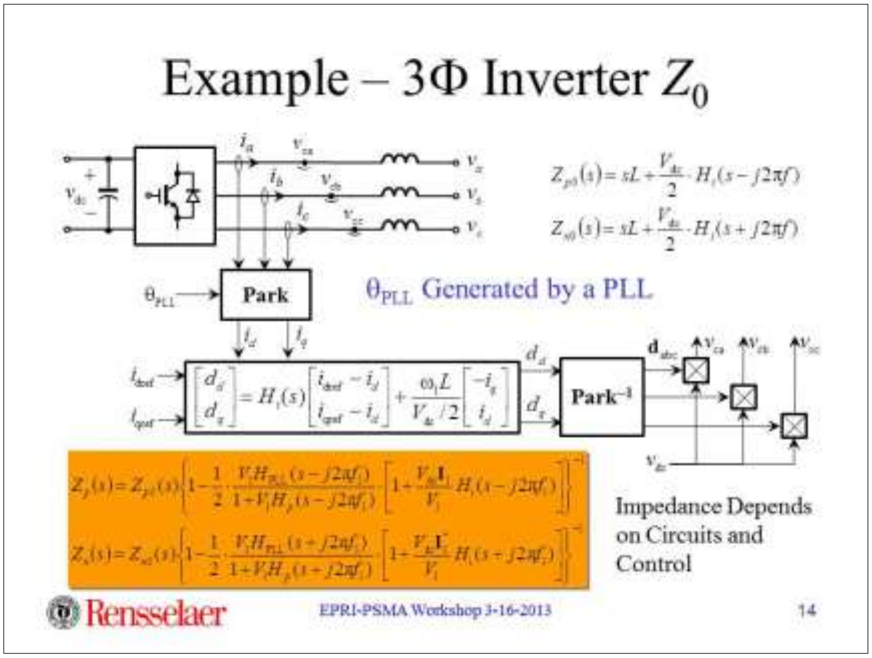

Here are some equations. I'll skip them, but show you what they mean. We can now model a three-phase converter for wind or solar applications. This is a grid and there are different ways to do controls. These are the sorts of impedance we work on.

What I want to point out is this impedance is analytical. Once you have the impedance, you can use that to guide your design. That's the beauty of these analytical models; they are not just quantitative analysis. You can trace back if you have any problems and go back to the impedance model and identify what is affecting the impedance in a particular frequency range, then fine tune the controls so there is no instability.

Harmonics have been considered a static power quality issue. Actually, they're an instability problem related to power electronics.

An issue we have been working on has to do with harmonic resonance. As I pointed out at the beginning, harmonics have been considered a static power quality issue. What we've found in this context is that they represent an instability problem related to power electronics.

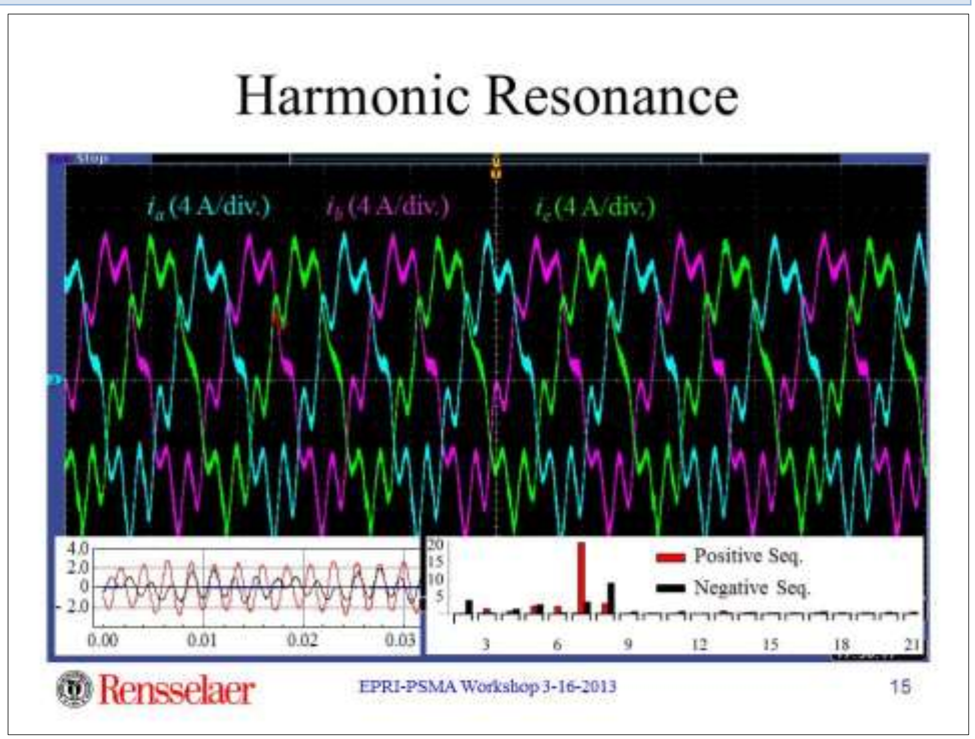

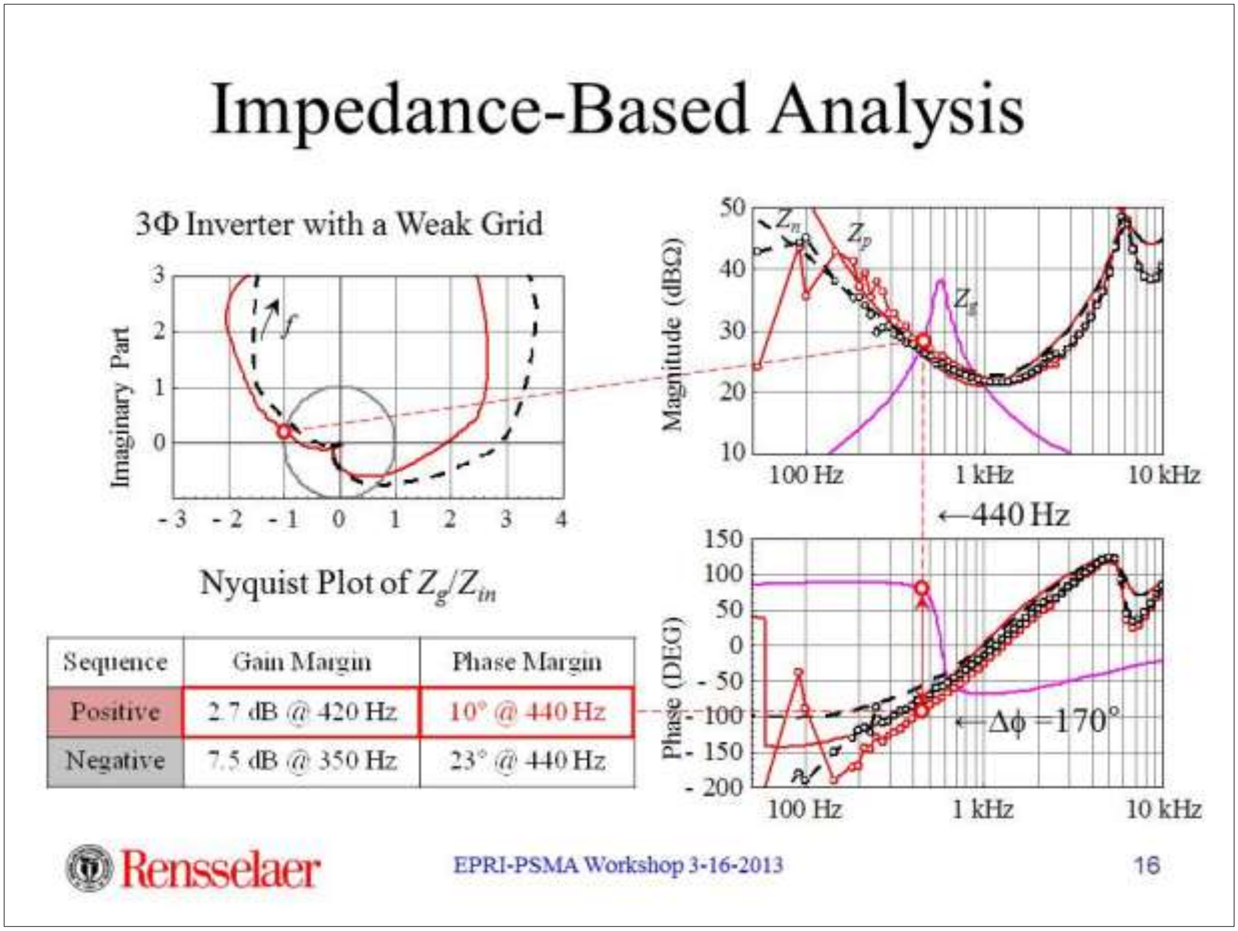

So slide this shows a particular measurement of a three-phase wind inverter in the lab. You can see the harmonic current resonance system resonate in the current. If you do a frequency decomposition, you can see the strong seventh harmonic in the positive sequence. There is also negative sequence impedance. We also found pseudo-harmonic forces. Even for three-wire system in a mechanism, you can see there is strong third harmonic. People wouldn't believe that because, normally, third harmonic is zero sequence. But with this sort of behavior resonance, you can end up with third harmonic in either the positive or the negative sequence domain. They can exist. That is, again, due to the non-linearity.

If we apply the mathematic tools that the impedance ratio implies, you can look at a positive sequence impedance and a negative sequence impedance. The red is the positive sequence impedance; the black is the negative. The purple is the impedance of the grid. You can see they intersect at a certain frequency. When they intersect, the impedance, the Nyquist plot of the impedance ratio, leaves the unit circle.

What's important, then, is the phase angle at that point. You have that to measure again. If the gain exceeds one, the phase angle comes into play. You measure that phase angle at the second point and that's the phase model. The distance to the 180 degree, the phase model.

In this particular case, you can see the positive sequence of the impedance system has a 10 degree phase margin. That, practically, is not sufficient. We are not likely to have resonance at that frequency, and that's actually the harmonic we saw in the seventh harmonic.

So these correlate very well with the theory and the measurement. The power system people may say, “Well, okay... that’s an equipment design issue. You have to design to operate with all this grid for maximum output impedance, so on and so forth.”

I can have a perfectly designed inverter that works in this location, but you install it somewhere else, and a problem appears.

What I want to point out here is that’s an impossible job because the grid changes. It’s different for every location. I can have a perfectly designed inverter that works in this location, but you sell it to someone else, and a problem appears. This is an issue that both sides have to work on together. You can’t solve the problem with one side.

We develop algorithms to measure the grid impedance in real time.

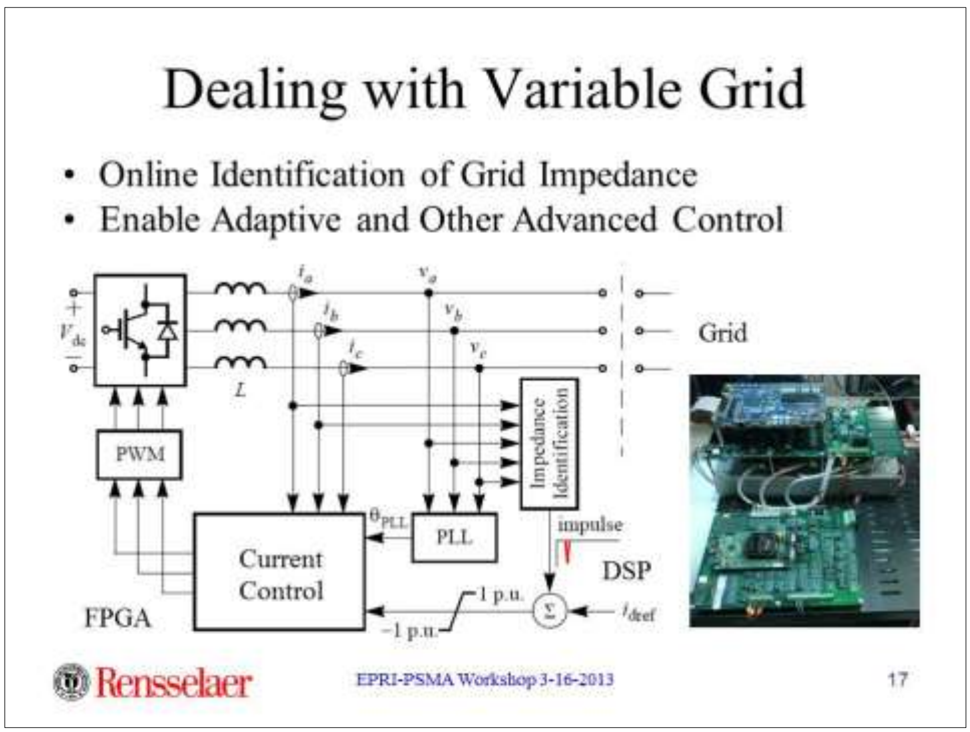

First of all, because the grid impedance is important, you can’t rely on someone to call you and tell you what the impedance is, right? So we develop algorithms to measure the grid impedance in real time. This is the inverter control circuitry and we add the algorithm to periodically measure the impedance of the grid by injecting either an impulse or other forms of perturbation into the current. That responds with the grid voltage and, from there, we can figure out what the impedance is in real time online. We have the ability to do that.

One step further, once you know the grid impedance, you can use that in conjunction with the impedance model inverter. Now you can synch with adaptive control. As impedance changes, you can change it; go back and choose it again. The phase lock load, current, and different functions are trying to stay with the grid without any stability problem.

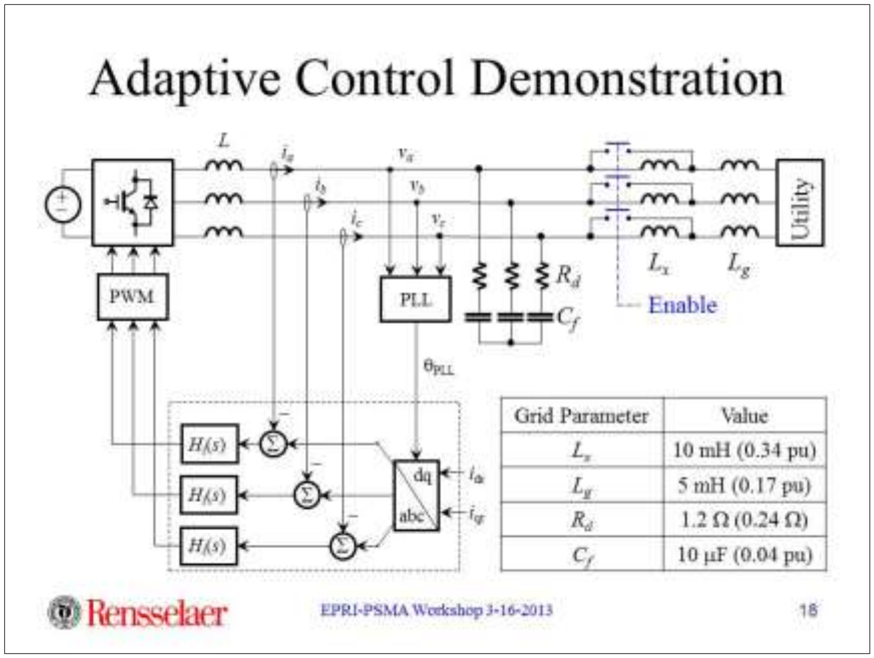

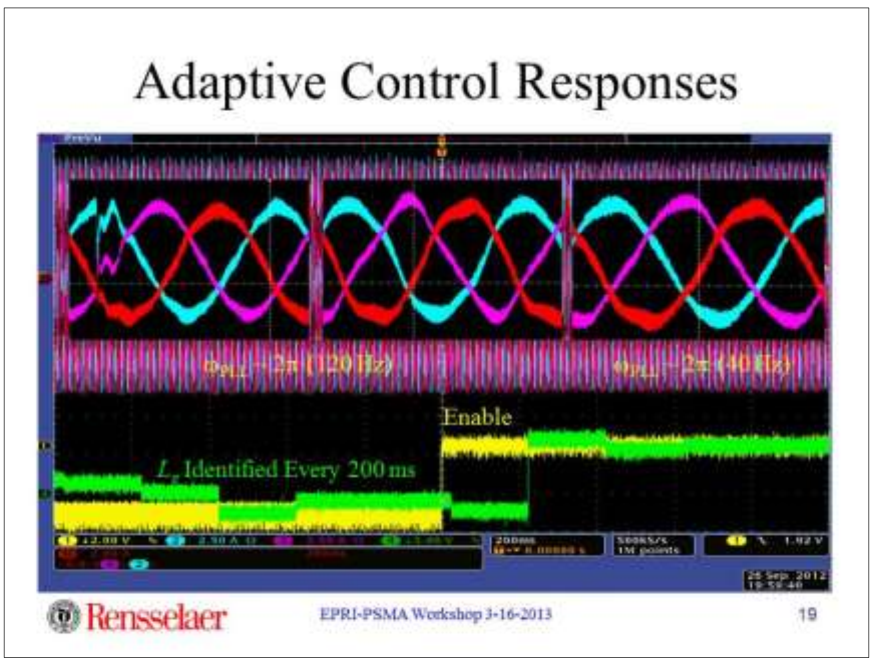

Here's a demonstration. It's complicated, but what it shows is that we have the impedance of the grid. Make it strong or weak and we can use an algorithm to identify the impedance of the grid gradually so that one line here shows the identified impedance value in the algorithm. You can go from no harmonic resonance, where you have a strong grid, to significant harmonic distortion in the current. Then change the impedance in the grid, but not change the control inverter. You'll eventually change to an adaptive converter that can control impedance to resolve the resonance.

In dc power system, constant power loads translate into negative incremental resistance that destabilizes the system.

That's what we've been working on. There's a lot of detail here, but what I want to emphasize this is an issue between equipment design manufacturing and the grid. So, speaking about the gaps, this is one of the areas the two communities need to come together to make sure this doesn't become a problem.

Constant power loads are spoken of very frequently. In a dc power system, which I think is most of what PSMA worries about – power supplies; constant power loads translate into negative incremental resistance. That’s what we all know. That is to destabilize the system and all the power we know about.

Look at the ac power system. First of all, we have to redefine that… because, normally, particularly when you have a single phase load, you can’t have constant power behavior. In most cases, when we talk about ac power system grid, constant power really means constant average power.

Instantaneous constant power is only possible with a perfectly balanced three-phase load.

What matters over fundamental cycle area is power average over just that cycle. Instantaneous constant power is only possible when you have a perfectly balanced three-phase load. If you design your circuit properly, you can achieve that instantaneous constant power, but that’s not really common.

Most of the time when we say constant power, it really means average constant power or fundamental cycle. And that has, although it’s just the wording or terminology, huge implications for the dynamics.

The third category is constant power loads behind the ac-dc converters. You have the constant instantaneous power, but behind something. It's a dc load behind the converter. And it's really common.

I want to just briefly touch on these. Each one of these has a different effect on system stability. You have to model them differently because so-called negative resistance does not apply in any of these cases. It does not apply because it's only meaningful when you have instantaneous power, constant power. If you have average constant power, these terminologies don't work.

Just to make my point, I want to show you the input impedance of some of these so-called "constant power loads" …which are really constant average power loads.

This is a single phase PFC, a single-phase power factor correcting rectifier. This is really common for a constant type of power load. We model input impedance and we measure the impedance of the rectifier. This shows to the measurement, the magnitude of the impedance, and phase angle of impedance.

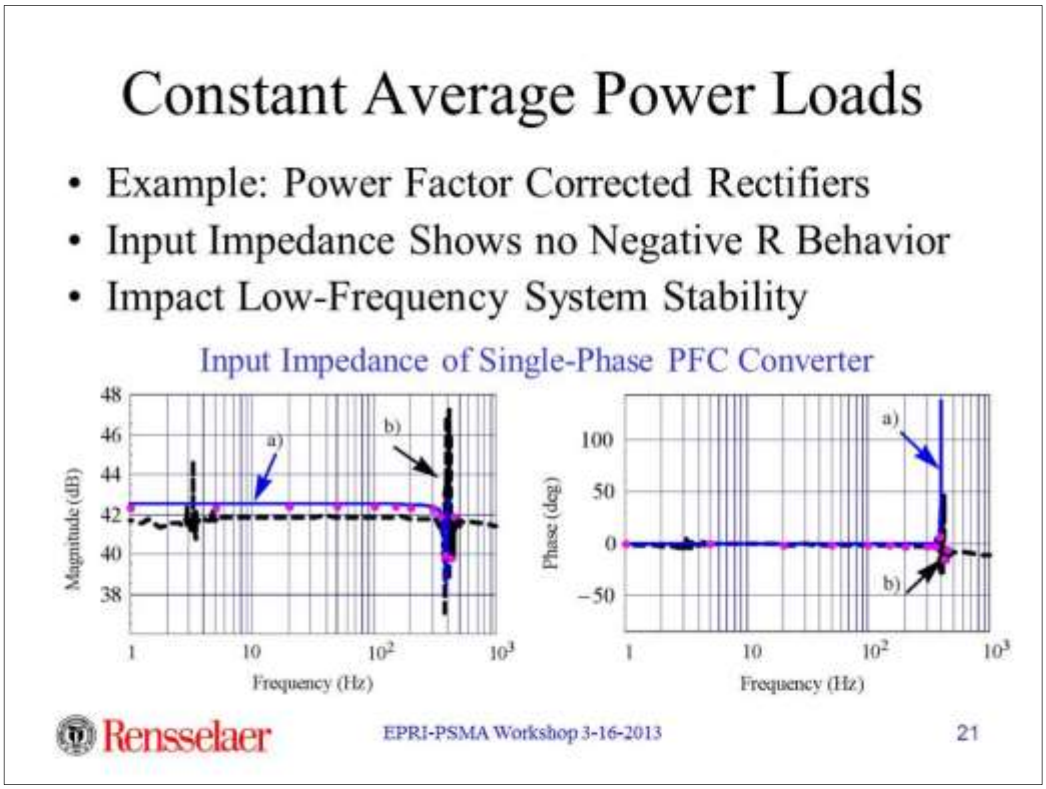

Notice the frequency of this from 1 Hz to a few hundred Hz, about 4 kHz... it's low frequency and high frequency. What you see here is that the input impedance has no phase; virtually no phase here. Instead of -180 degree phase when you have a negative resistance, this is a normal resistance. Perfectly normal positive resistance. There's no constant power of a behavior in that sense. It's not a negative resistance.

In a single-phase power factor correcting rectifier, as we model the input impedance; we measure the impedance of the rectifier as well. There's no constant power behavior. It's not a negative resistance.

It doesn't mean that these types of converters will not affect the stability of a system. They just have very different effects. We have to use different models to assess that, but the natural resistance sort of model does not work here.

The other kind of a circuit, which I show here, is a single-phase power factor corrector actually used as a control circuit. This is about an uncontrollable rapid fire in the front… a three-phase, uncontrolled, rapid fire. Then you have a constant power load, dc load, at the dc output.

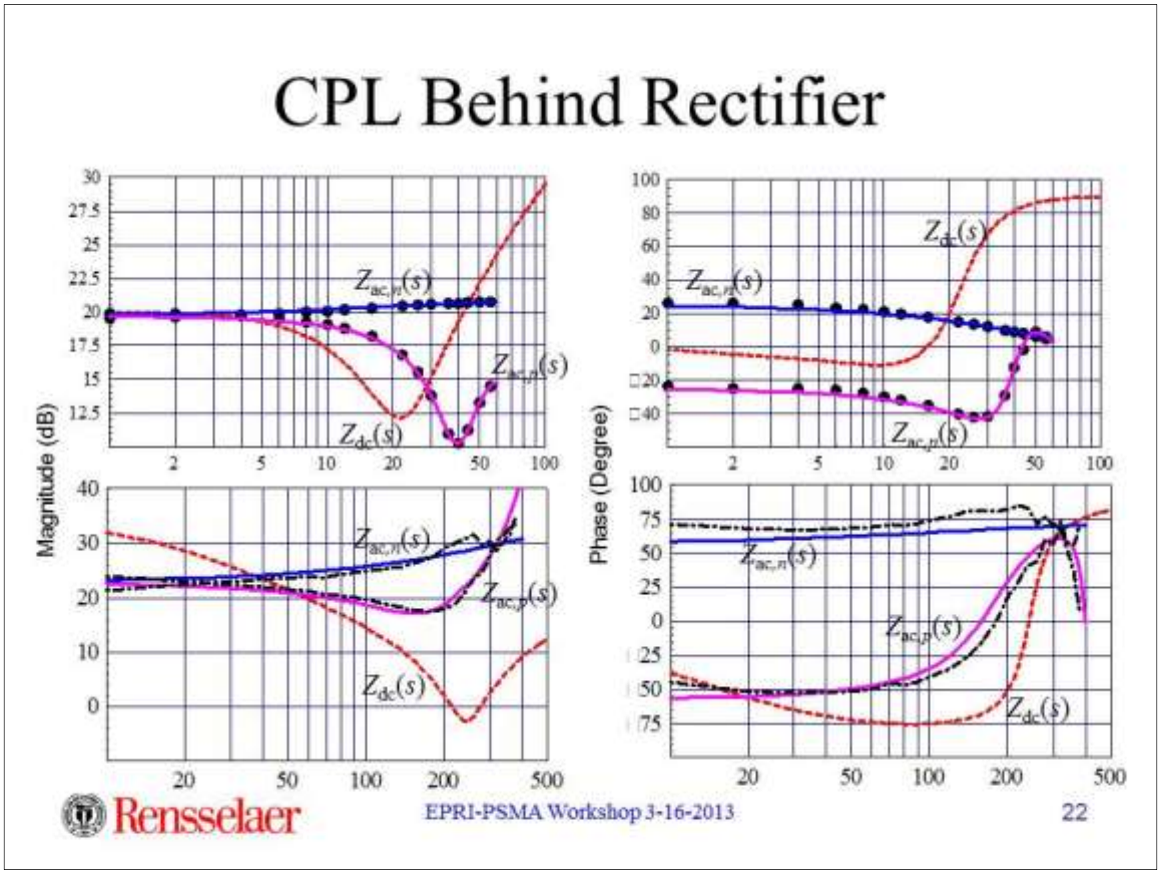

This is the input impedance measurement of that circuitry because it is three phased. We decomposed that into a positive sequence and negative sequence, until you both have positive sequence and negative sequence impedance.

There are two cases. The first case is when we have a normal RLC load, so the dc load is just normal, passive RLC loads. You see the impedance of dc loads, the magnitude and phase angle, then versus the impedance when you max out the ac terminals, the positive sequence/negative sequence impedance. The dots and the lines represent measurements and active model predictions.

The thing I want to point out here is that there is no constant power behavior… it's passive.

The second case is the load constant power load. The red line is representing impedance of the dc load; that's the magnitude, and this is the phase angle. If you look at the phase angle response, this is a minus sign. Somehow it's been messed up. What you see here is: up to a certain frequency range, there's a constant power behavior on the dc load. Once you translate that, transform that, and rectify it… you can look at the ac side, the natural phase angle. Look at the magnitude in response here… that's the phase response. If you look at ac input impedance? There's positive sequence/negative sequence versus positive sequence. There's no constant power -180 degree phase here. This is a combined phase response of the constant power load with theory in front of it. So you don't see the -180 degree phase itself.

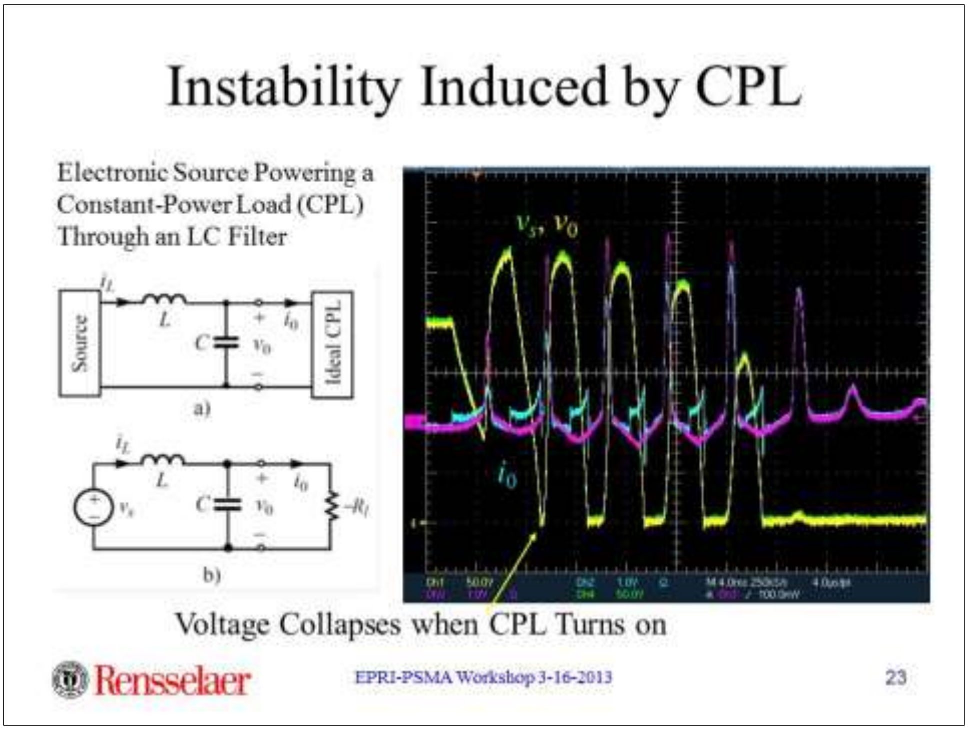

The point I want to make is that you can have that constant power load behavior behind the uncontrolled rectifier; but look at the ac terminal and you don't see it anymore. The whole point is the constant power load needs to be very carefully examined when you access the faculties on a power grid.

The dc side is easy. We have the tools to assess the instability and so on and so forth. The measurement is to show the ideal constant power load. On the dc system, we have the power source and we can stabilize the constant power loads by passive filters. This shows a constant power load with ac filters. You run the system, it's unstable. So the voltage, you try to stop the load. The load, the motor, oscillates for a few cycles then shuts down. It can't operate due to instability caused by the constant power load. Obviously, you look at that and, minus active resistance, the system won't be stable.

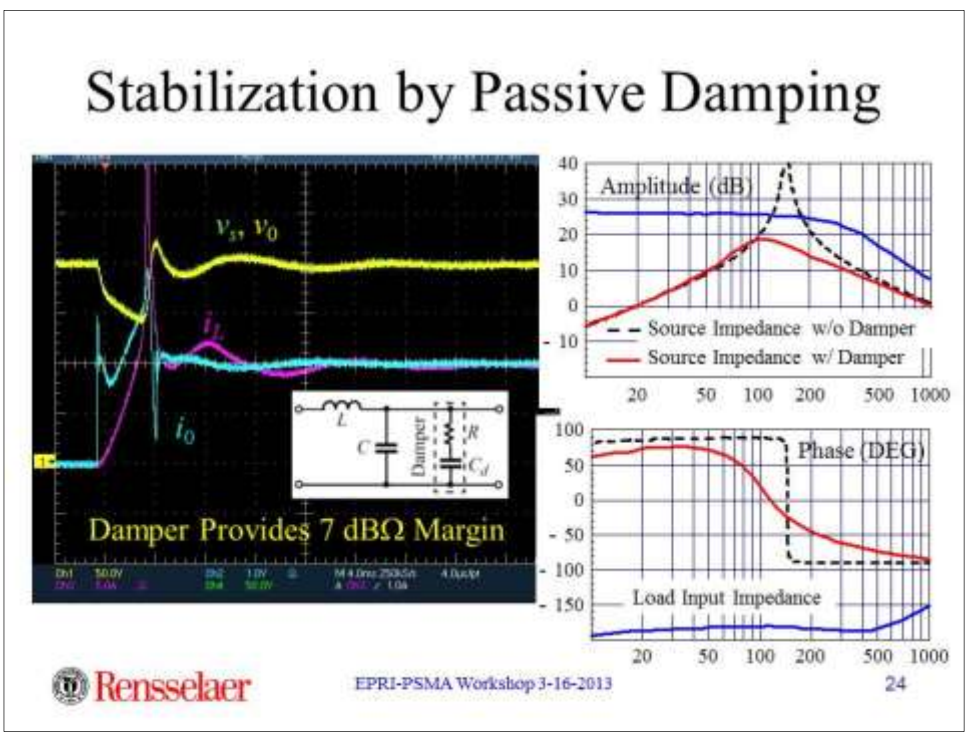

We can stabilize that by many different methods. One simple thing we did in the past was to put a damper there. I'm not touching anything else; just put an RC damper with filter capacity and that takes care of the instability so the voltage can stabilize. You can do all of the analysis by looking at the impedance.

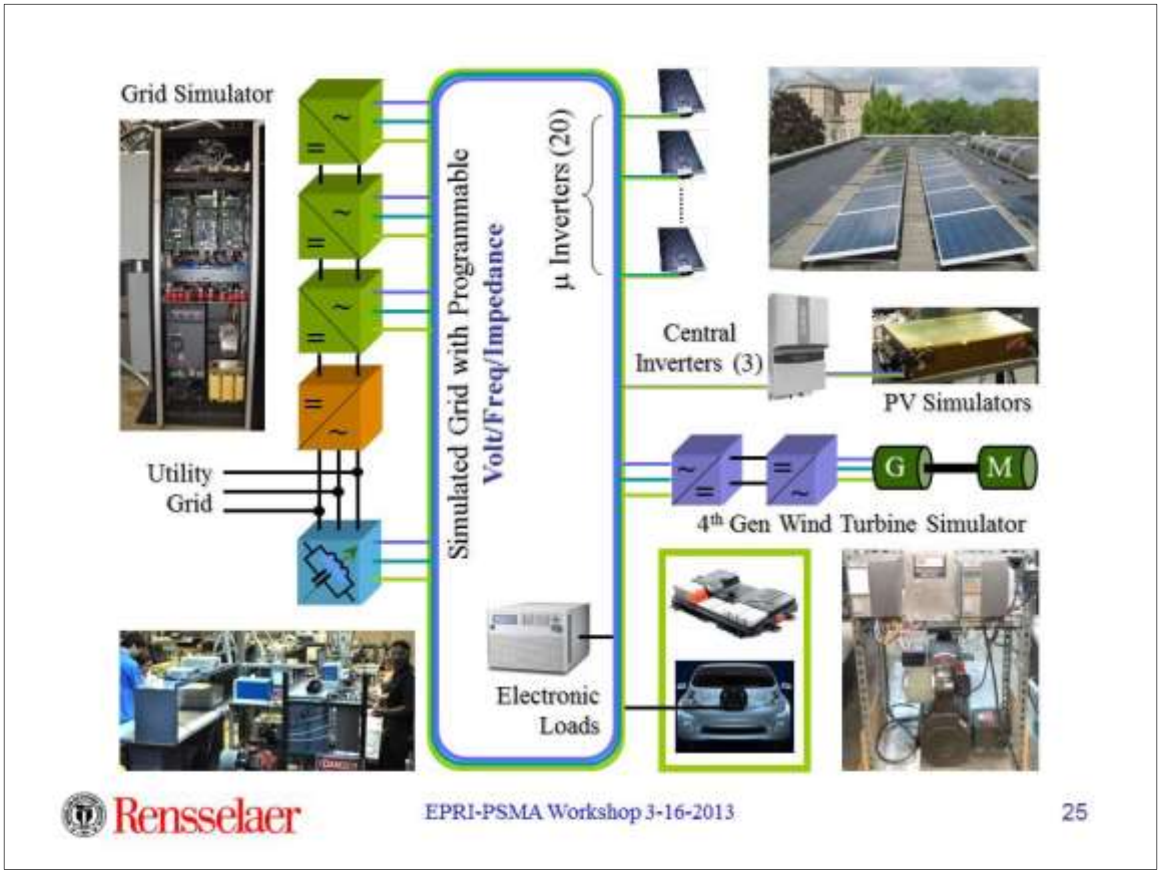

To support a lot of this research and be able to test this in the labs, we designed this programmable grid that we can control. We do that by using passive elements such as back-to-back converters, so we effectively have a simulator and we can program all the parameters we want. That's the platform we use to do the experiments.

To summarize; converter design, the power electronics part, needs to focus more on system impact and performance to prepare for the future. My message for my system colleagues is that system operating and control must consider power electronics. We can't just look at them as black boxes where everything is designed perfectly so we don't need to worry. We have to worry about it.

That's the end of my presentation.

Thank you.

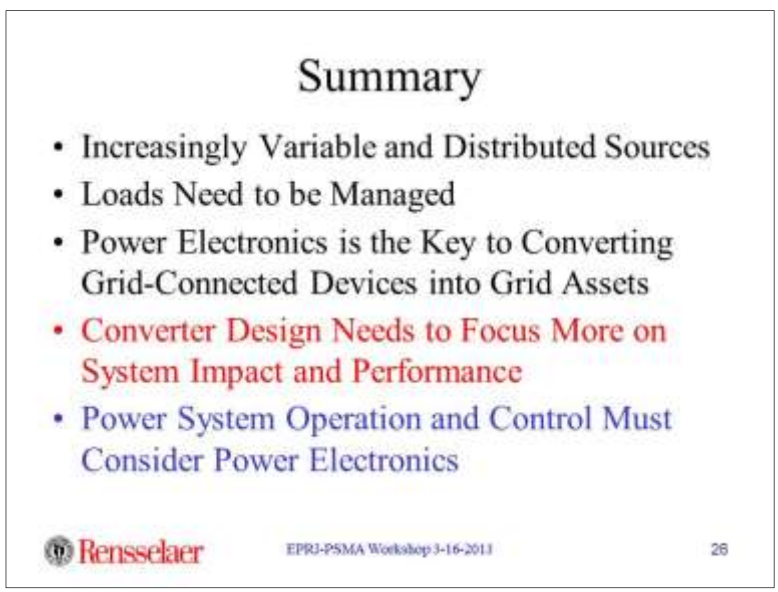

Evolving Communications Agenda

DUSTY BECKER:

Next we have Harley Garrett. He began his career as a Command/Instructor, Pilot/Flight Examiner of multiple aircraft at Squadron, Wing and Major Command levels. While there, he developed Command operational requirements and aircrew training standards.

He has worked with SCI Systems, Inc. In 2003, he joined Global Technical Systems and developed strategies to leverage the company's open architecture advanced storage area network technology with Navy Surface Warfare program offices.

Currently, Harley is the Vice President and set up the Mississippi-based software engineering team to develop integrated systems solutions for Command and Control by using advanced processors with emerging communications media, open architecture standards, and open-source software.

Harley has supported GTS's common processing system and published multiple papers on Internet and electric grid vulnerabilities and how future control systems could be made more secure. He has developed a robust cyber secure, scalable EMS/multi-grid architecture to control alternative energy generation power sources.

He has a BS in Mechanical Engineering and a Masters in Systems Management.

You're on, Harley!

HARLEY GARRETT:

While we're waiting, this is a large subject, like all the subjects, and I always try to get more than I can handle, then downsize it. Sometimes I make it, sometimes I don't. It's like a poem my mom told me:

There was a young man from Japan

Whose poetry would never scan

When told it was so

He replied, yes, I know

But I always try to get as much in the last line as I possibly can.

Greetings from Mississippi!

If there's anything that we like to do better than power electronics down there; it's fishing, going to football games, and deer hunting.

Evolving communications. I thought about how to handle this and I decided to try to give you the perspective of two evolutions. One is really a revolution and how these things have started to come together. That's the digital revolution. Then the evolution of power electronics... the impact of digital networks has been absolutely spectacular. The pace of innovation is now quickening over the past ten years, and it's moving toward interoperability and convergence, wired and wireless. I'll wrap it up with some future trajectories on what I think is happening.

The impact of digital networks has been spectacular. The pace of innovation is now quickening, moving toward interoperability and convergence, wired and wireless.

This is a quick and dirty overview of the last 40 to 50 years of history. The bipolar transistor was a game changer that created an absolute revolution in communications. It went from analog to digital. There was a first transistorized computer that came out in '54. Computers back in those days were gigantic things. Everybody remembers what they looked like. By '57, you had the first silicon controlled rectifier, which basically was where the old mercury arc rectifiers have now gone and everything now is transistorized.

So the purpose is to show you how electronics and digital information have sort of intersected in the past 50 years. CDC came out with the first commercial computer. That was about the same time they invented COBOL, which was a business-oriented software development language. That computer set the stage for the IBM mainframe, which came out later. The DEC, PDP-8 was probably the first smaller computer that was affordable. They called it a minicomputer. It also was a predecessor of linking computers together.

So off we went. 1968, Dick Morley created what he called a "modicon," which is a modular, digital, controller. He avoided calling it a computer like the plague, but that was really the evolution of the Programmable Logic Controllers. We have millions of PLCs today, as you well know.

Intel 4004 was the first solid-state logic processor. A processor is exactly that: It processes information, adds, subtracts, multiplies, divides, etc. It was manufactured on a 10 micron size, which means the transistors were no wider than 10 microns. A micron is a human hair, a millionth of a meter.

That particular device had 2250 transistors and each transistor, as you know, is bipolar. A lot like some people we might know. But it's either zero or one. You pass a current this way, it's a zero. You pass it that way, it's a one. So you're into the one and zeroes game.

In 1979, they made the first static VAR compensator.

It wasn't very long after that they came along with the first static VAR compensator. Micro-computers showed up in 1980 and we're using those today: the PC, the Apple, and so forth. The Internet came along about a decade later… and guess what? All of the old analog-type devices that used VARs have now gone to a whole family of thyristor-controlled systems, which is a hybrid word that is a combination of resistor and transistor.

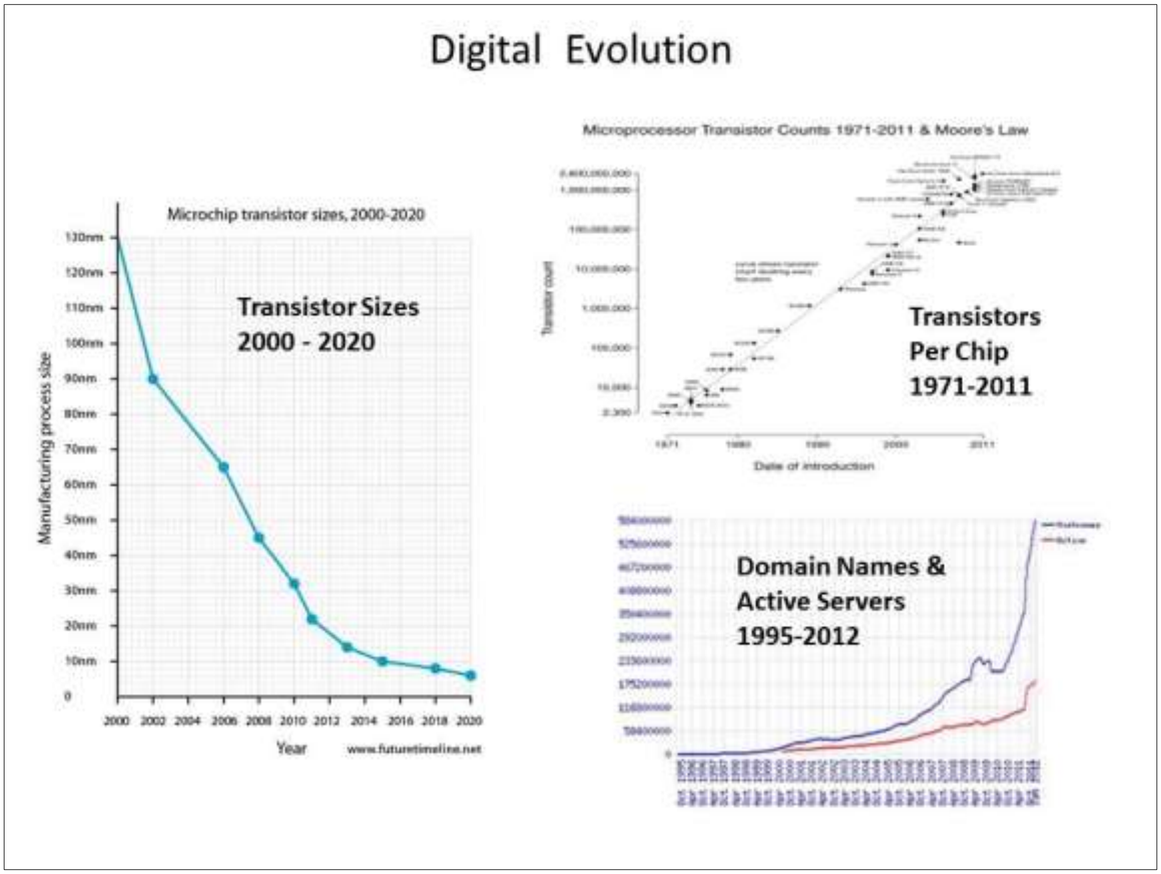

So these charts put it together for you a little history of how fast this technology has developed. On the left side is the manufacturing process used in developing chips. I wish we could go back to 1971 to that Intel 4004 because you really see how it's advanced. I remember, in the Pentagon, we were blown away by the fact that DARPA was putting money into an advanced chip that would be like two microns.

Well, they went well below 2 microns. Now we're down in the submicron area, 130 nanometers. They were all the way down here.

Intel is working a chip that will be a system on a system on a chip, which puts it at a 22 nanometer layer. Some people think Moore's law, which is up here on the upper right, may run into a brick wall one of these days, but so far it hasn't.

We're up to 2.6 billion transistors on a single chip…. you can do a lot of computing a lot faster.

We're now up to about 2.6 billion transistors on the single chip. What does all this mean? It means you can do a lot of computing a lot faster. That's what's been going on.

This chart is Netcraft, over in England. It's basically tracking the growth of the Internet. They're doing that in two ways: A domain name is ".com" or ".net." Below those, you have all kinds of sub-domain names. Because this is all controlled by the ICANN (Internet Corporation for Assigned Names and Numbers), only one person can have a domain name. When I tried to get AOL to give me GHarley@aol.com, they didn't want to do it. Somebody else had GHarley@aol.com.

These are the domain names: basically in 1990, up through today. The domain names are not all necessarily being used, but they've been assigned. Those get implemented on servers. So how many servers are out there in the world?

Well right about there, you can see what happened when the Internet started. It started to go asymptotic. This red line is actually much higher than it looks on the chart.

There are 250 million servers — that gives you an idea of how big the Internet really is and what Moore's law is doing.

Then we hit the dot-com bust. Everything hiccupped up for a minute, then it kept right on going. Today, I think it's in the whitepaper… but if I remember correctly, this is about 690 million, this is 250 million servers, something like that, and growing. So that gives you a little bit of idea of how big the Internet really is and what Moore's law is doing.

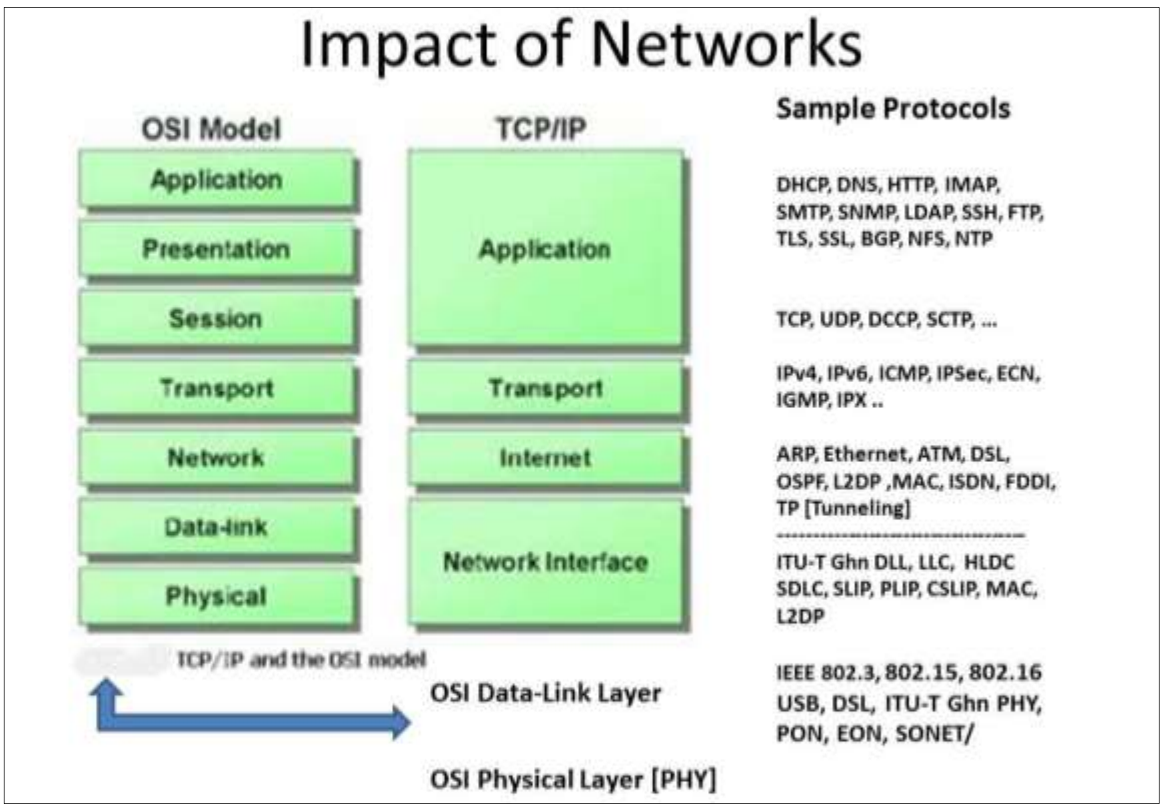

This is key. It's key because in a capitalistic society, we want people to be able to design and manufacture widgets, smart phones, cell phones, whatever, and do it in a proprietary way so to keep their intellectual property. At the same time, that device will not work over a network unless it's compliant with the protocols involved in the network.

A device will not work over a network unless it's compliant with the protocols of the network.

So this particular model (on the left) has been around since about 1980. It looks at networks in a layered fashion. You start down here at the physical layer, which is where your hardware is... literally. The server is a piece of hardware, your cell phone is a piece of hardware. If that has been developed against a certain set of protocols, then those protocols calls are compatible with the ones that are in the next layer above.

The data link layer is where a couple things occur. One is each device on this model has its own media access control number, which is unique in the world, implanted on the device itself. It's burned into it. So the upper part of the model knows that that's a Cisco server that operates according to these protocols. It knows that. The data link layer itself has a logical link control that goes right up the chain of command.

The people who developed the Internet didn't want to create a stack that was different from the OSI model. They didn't want to mess with the data link in the physical areas right here. So they came up with protocols, where the network interface on this side of the model, the TCP/IP, which is Transport Control Protocol / Internet Protocol. TCP/IP actually does the same thing as a data link in the physical. It has to, to some extent, but it sits on top of an Internet network. Transporters are the same. On the application level of the Internet TCP/IP, you actually do the presentation in the system protocols in there, although they're not shown.

If the hardware is designed to be compliant with that protocol, then it can be designed by anybody using whatever intellectual property they want.

Do I know what all those protocols are? Yep. I got a whole page of them if you want me to read them off to you, but I won't do it now. The point of a protocol is this: If the hardware is designed to be compliant with that protocol, then it can be designed by anybody using whatever intellectual property they want.

Let's look at it this way: In the cellular area, we have basically two technologies competing, CDMA and GSM. I'll get into that in a minute, but if you had a GSM cellphone, it would not run on a CDMA network and vice versa. If you have an Apple commuter and you go down and buy a Microsoft application, it will not run. Why? Because they're not compliant with the same protocols.

Networks are expensive, but that's where you're going to get commonality and standardization. Not necessarily standardization, but convergence. The Navy has got Grumman radars on their ships. They've got Raytheon radars on their ships, they've got General Electric radars. Great radars, they are stove piping. They don't talk to each other. If you're looking at an airplane on this radar and it says that the airplane is there, this radar says no, it's not. It's over there. They're stove-piped. How many databases do you know that are stove piped? This database doesn't talk to that database.

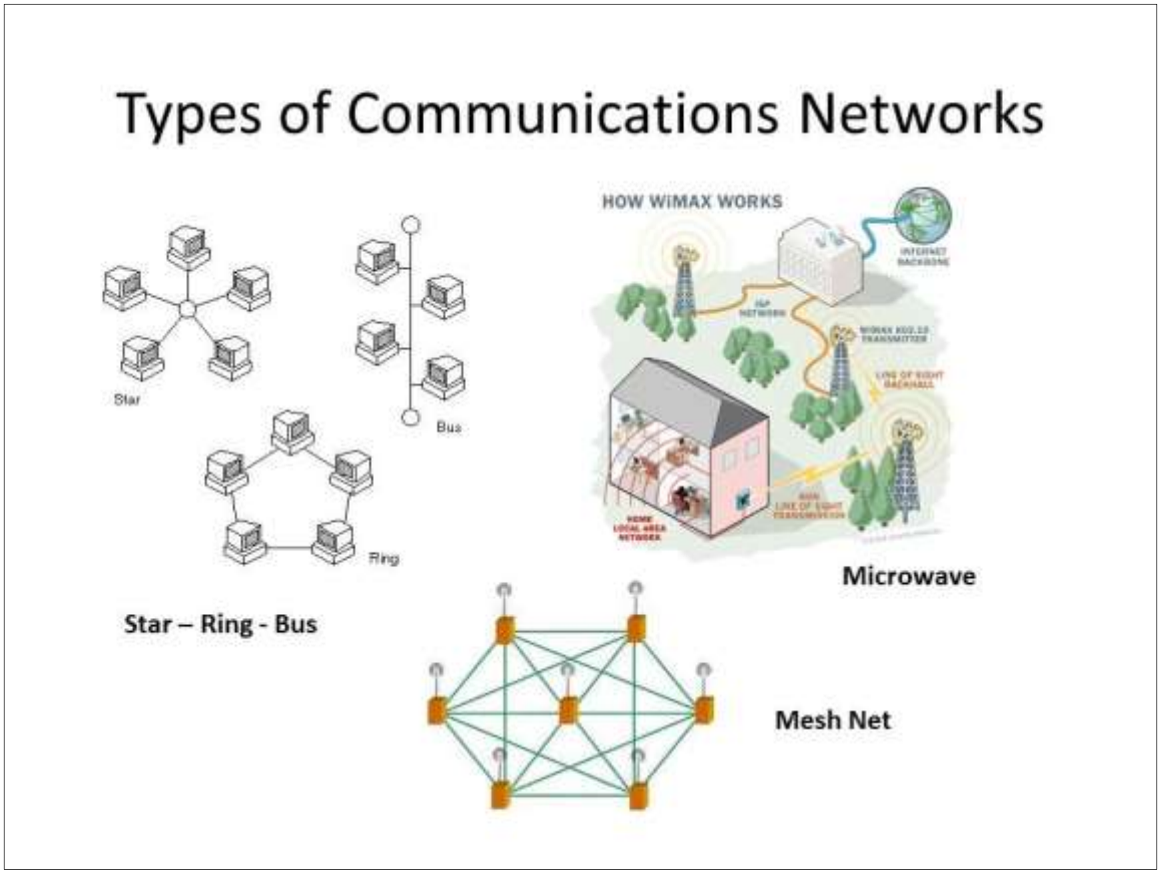

There are different types of networks. We can go to a star network, where you have a master controller in the middle and the units all around the outside don't talk to each other. They can, but only if they go through the middle.

Meshnet is an interesting network because it's very redundant.

A ring bus is the same thing except no center. The standard bus that you see here is a lot like the Mil-Standard 1553 bus that tied all the avionics together on an aircraft about 25, 30 years ago. Meshnet is an interesting network because it's very redundant: any node can talk to any other node three or four different ways. If you have communication breakdown in there, there are a bunch of redundancies that pick that up.

I put microwave in because it's a little bit different. Microwave is a communications-oriented communications medium. That means that, for a microwave receiver to transmit information back upstream, it can't until the base station assigns it a channel. That channel becomes that receiver's channel. A little bit different from the star and the ring bus, where they're all running peer to peer, point to point, or point to multi-point.

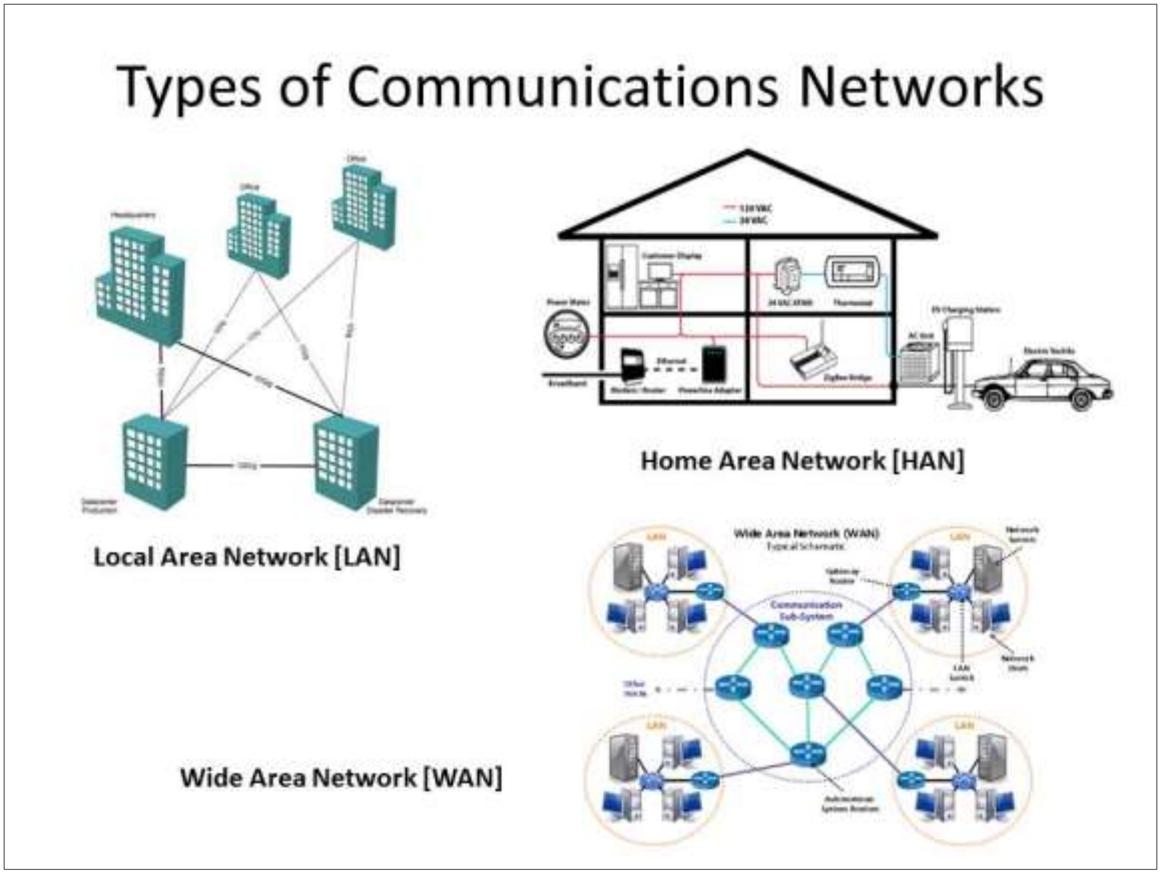

Now let's look at networks themselves. If I drive down a road and I've got my ear piece in -- which I don't, by the way -- and I'm talking on the phone, I'm using Bluetooth. Bluetooth is great, but it only goes about 4 feet, maybe 8. That's a personal area network, that's what that is.

Now add a bit more; let's hook rooms together, maybe a couple of buildings together, and we have a local area network. If you have a network that oversees New York City, that's a metropolitan area network.

You've heard about the home area networks. That's because we're coming of age, now, into the energy demand/response... also communications, but I'll get to that in a minute.

It doesn't matter what type of network you've got, you can interconnect them.

The important part about that chart is: it doesn't matter what type of network you have, you can interconnect them. That's precisely how the Internet works. They're all interconnected. Once you have an entry point into the Internet, you can go worldwide. I will leave the security of that up to the next speaker.

So what is happening here? Let me tell you what's happening. I'm a firm believer that if you want to know what is going on in a particular area, you have to try and understand how you got where you're at. I tend to go back in history... so I like very much some of the presentations I've seen. If you look in my whitepaper, you'll find that I started this about 1830, somewhere back there, and walked my way up to see how we got where we are. All the way through Edison's current wars, Westinghouse, and all the rest.

Everything was percolating along, the digital revolution was going along like that, power electronics were going on like this, and they cross a little bit here and there, but they weren't being driven by the same thing. The digital side, though, was being driven by Moore's law. And the consumers of information want more and more information.

On the energy side, it was being driven by modernizing FACTS, better devices, and so forth. That all changed in 2001.

> *In 2001, President Bush signed NSPD (Directive 7) which established DHS. Ten years later, the Congress passed the Energy Policy Act, EPACT.*

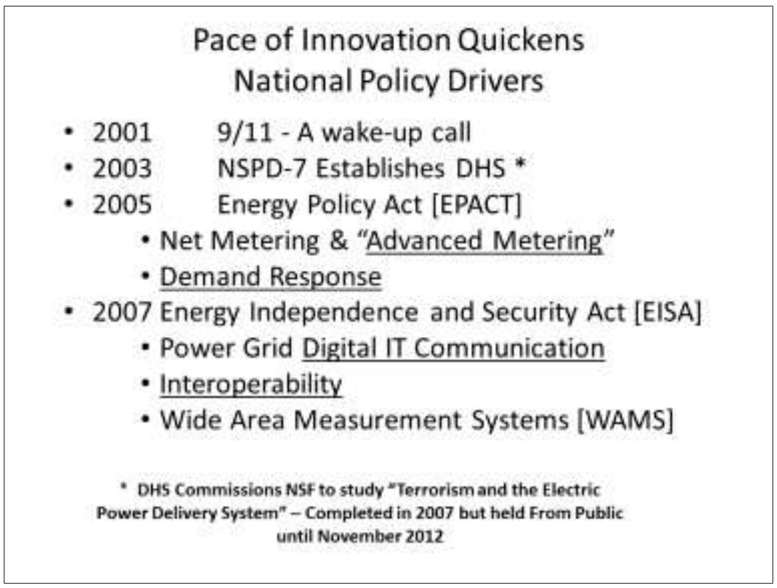

President Bush signed NSPD (Directive 7), which established DHS. Two years later, the Congress passed the Energy Policy Act, EPACT, which, by the way, was heavily focused on renewable energies.

I pulled out a couple of key things in that act that are very important. One of them is net metering — if you have a solar panel on your house, you are entitled by law to only pay for the electricity that you use if you don't generate. If you're generating what you use, you don't have to pay any utility for that electricity.

So you're only paying for the electricity that you use and if you don't use what you're generating, you can sell the excess back to the grid. By the way, if that actually happens and your whole house is running off of your own meter, it will run backwards. It really will.

> *Net metering: if you don't use what you're generating, you can sell it back to the grid.*

Advance metering was a title that they gave the meter that turns backwards because they realized that to do net metering, they had to have a little bit smarter meter. They didn't know what to call it, so they called it "Advanced." That very quickly turned into what we used to know as AMR, but now know as AMI. We're basically talking about a meter that used to be one way, now it's two. You cannot do demand/response without a two-way meter.

Demand/response. There you are. In 2007, EISA came along and it came down very, very hard on digital communications. We need to use more of it in our systems and we really need interoperability, going back to the radar stove-pipe thing. They recognized the utility of phasors. Phasor Measurement Units could be used to measure the units, PMUs.

So now there's a heavy emphasis on evolving what's going to wind up being a coast to coast, Mexico to Canada, WAMS: Wide Area Measuring. Because if we really had that now, and we have the technology to do it, I think, Con Ed could know what's happening in Bonneville… just like that… in real time.

There's a heavy emphasis on a coast to coast, Mexico to Canada, WAMS: Wide Area Measuring.

I put something on the bottom here that I think is important. When 2007 came along, DHS commissioned the National Science Foundation to do a study to figure out how to evolve the smart grid. They did that, but when they delivered the results of the study (they started this in 2004 and delivered the results to DHS in 2007), they immediately classified it. They wouldn't release it to the public.

Last November, the National Science Foundation finally got them to release it and it's available to the public. I strongly recommend you read it. Find some time to download that thing. It is excellent. Maybe a little out of date, but it is excellent.

JONATHAN POLLET:

Really only the cyber portion of it is out of date. I was on that panel.

HARLEY GARRET:

Yeah, exactly. Thanks, Jon.

With cellular, the pace of innovation is now quickening.

HARLEY GARRET:

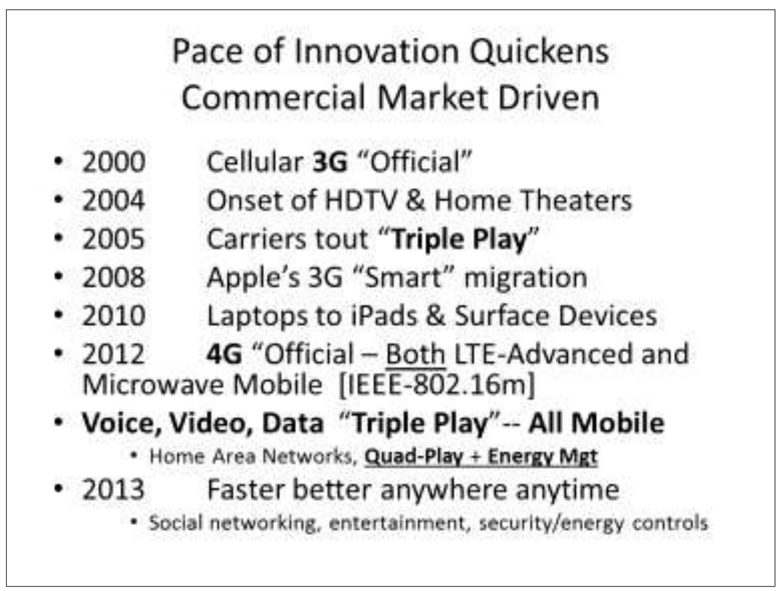

In cellular, the pace of innovation is now quickening. 3G came out in 2000; that was the ITU (International Telecommunication Union) and pretty much the IEEE and convergence on what 3G really is. In 2004, we got HDTV in home theaters, triple play. Triple play is voice, data, and streaming video over one connection. That's important. Now you can have quad play, which is the same thing except you can go mobile with it.

Why do I make a point out of that? Because the consumer is driving this. They want movies. They want football games. They want everything on Internet right there on their phone, NOW. Immediately. That includes their home.

The providers of that service are lining up because they see it as a great market. So this is what's pushing it right now.

Now, that doesn't mean that if this stuff goes that way -- and it's going that way -- that you can't use it for other reasons. This is where home area networks and energy and demand/response come in.

Broadband over the power line came out a few years ago.

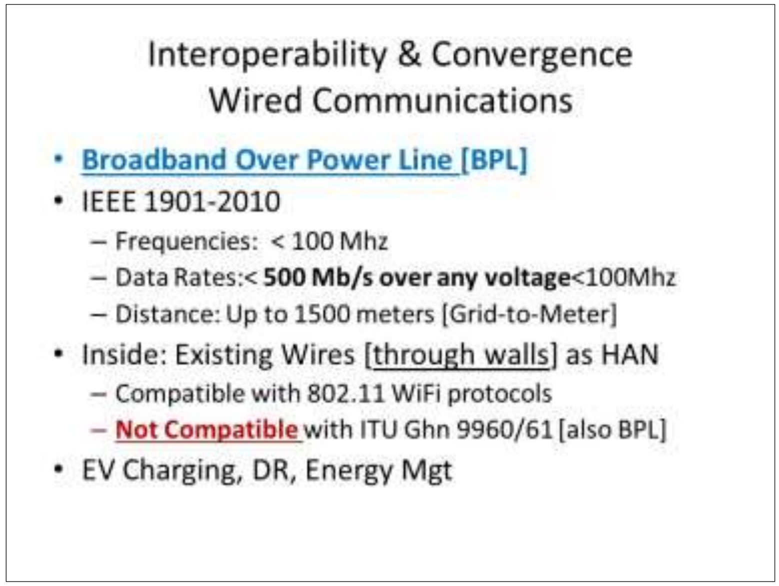

Broadband over the power line came out a few years ago. That's the IEEE standard for broadband over a power line. Data rates about 500 megabits per second over any voltage. This is for about 1500 meters. Not bad, but it doesn't work very well beyond that; as some large industries that put money into this found out.

Inside, it uses the existing wires. That's a neat feature because nobody wants to rewire their house. It's compatible with all the Wi-Fi protocols, 802.11. It is not compatible completely with the ITU standard for home area networks. It has to do with how Forward Error Correction is accomplished with different approaches to how signals are multiplexed. IEEE decided to accommodate both a Fast Fournier Transform and a wavelette approach using turbocodes to achieve forward error correction and, on the European side, Low-Density-Parity Codes. The European LDPCs are not compatible with the two types of IEEE approaches.

The latest is they've decided to become "coexistent" though I'm not sure exactly what that is.

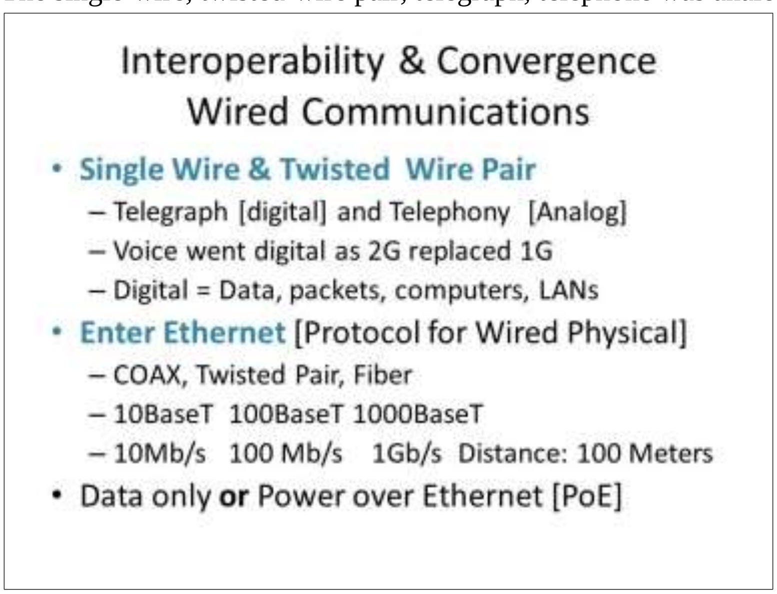

The single-wire, twisted-wire pair, telegraph, telephone was analog voice and went digital as 2G replaced 1G where data and voice was digitized and flowed over LANS in packets instead of legacy waveforms. This all started as part of the digital evolution and gave birth to all the protocols on that other chart.

Ethernet has been a major player. It's sort of like it went away for a while, but now it's back. The Ethernet is a protocol, but it's interesting because, if you comply with it, you can run your data over a number of different physical lines: co-ax, twisted pair, even fiber. With 1000 base T, we were up to 1.1 gigabit per second up to 100 meters.

Yes… you can do a little bit of power over the same Ethernet, but not very much… about 22 watts, something like that. That's actually a niche market for putting advanced sensors into substations. You don't want to necessarily run a copper wire to that particular sensor. You can run a fiber now and still power the sensor.

Fiber optics… the Holy Grail. They've demonstrated 54 kilometers using a 14-core pipe at a little better than one petabite.

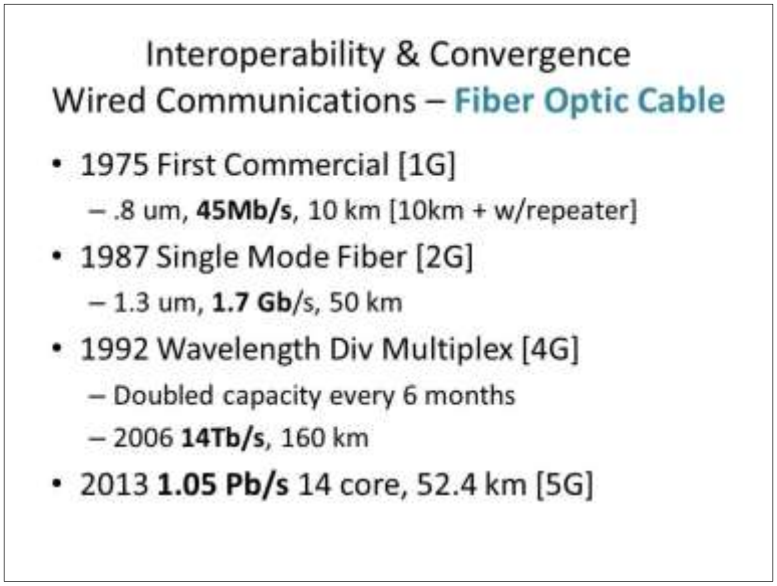

Fiber optic cables are the Holy Grail, as most people know. Eight-tenths of a micrometer and, back in '75, they were able to pass 45 megabits per second over it. It went to single-mode fiber in 1987. By '92, they actually took a jump because there are several multiplexing ways to achieve more data throughput over a single pipe. This was along the wavelength of that division multiplex technology that doubled the capacity every six months, up to 2006, before they demonstrated unequivocally down to 14 terabits per second.

Now, just a few weeks ago in the laboratory, they demonstrated 54 kilometers using a 14-core pipe at a little bit better than one petabite. A petabite is a thousand terabits. A terabit is a thousand gigabits. A gigabit is a thousand... And, you know, so ten-cube, ten-cube, ten-cube, ten-cube, ten-cube. Let me run through this as fast as I can.

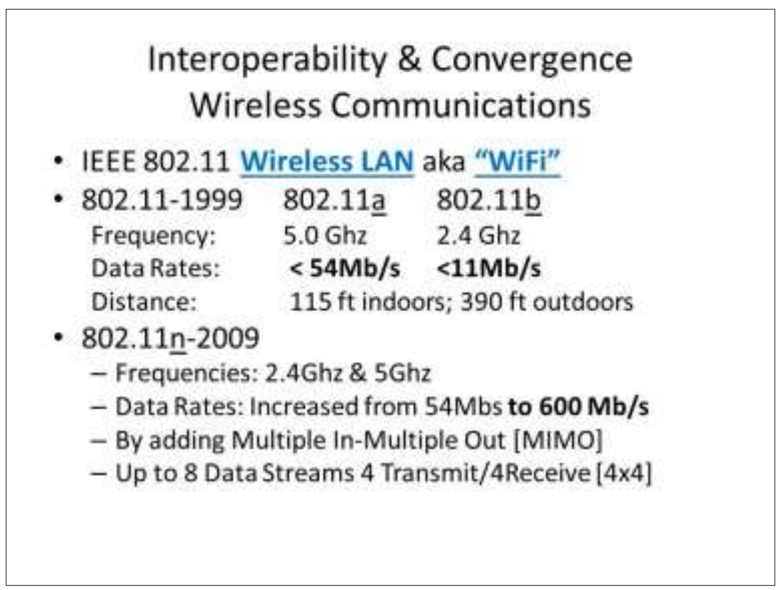

This is the wireless side of the house. You can see it started out with 802.11a, which is 5 gigahertz frequency range. Then 802.11b, which is 2.4 GHz. Back around 1999, is when that happened. 150 feet indoors, about 400 feet outdoors, and that became known as the wireless LAN standard, or Wi-Fi, as we know it today. It's up to 802.11n as of 2009. Per IEEE; those of you who belong know this: A has to be compatible with B, B has to be compatible, all of these have to be reverse compatible unless they get deleted. So N put everything together in 2009 and, at that point, they were using 2.4GHz and 5GHz modes. Data rates went from about 54 Mb/s up to 600 Mb/s.

So what's going on here? By adding MIMO, they basically put eight antennas on the transmitter device and eight antennas on the receiver device… and they could get up to a 4x4 throughput, eight mega streams, without increasing the size of the pipe.

Data rates. Data rates keep going up, up, up. Why? Because everyone wants more, more, more.

So what's going on here? Data rates keep going up. Why? Because everyone wants more, more, more. We're now up to 802.11ac, which, by the way, hasn't been published yet, but should be published later on in 2013. 802.11ac operates at the 5 gigahertz. They said, "Well, the 2.4 is getting crowded, so we're going to go to 5," and they did. Using MIMO, they didn't give up any distance. It's now competing with the cell phone and cable. It's not compatible with ITU for the same reason broadband over power line was also not compatible.

3G cell plus WIFI gives you triple play. That means you can have a cellular bridge to this particular standard here and connect devices inside the HAN to destinations outside the HAN at super-fast speeds. Manufacturers are shipping product now. Not to be outdone, wireless gigabit, now called WiGIG, is working on the 60 GHz band; which is seamlessly backward compatible with 2.4 and 5 GHz.

Now those data rates are probably going to exceed 7 gigabits per second. The only problem is it doesn't go very far. But, in fact, it doesn't have to go very far.

All of these wireless technologies bridge to other technologies.

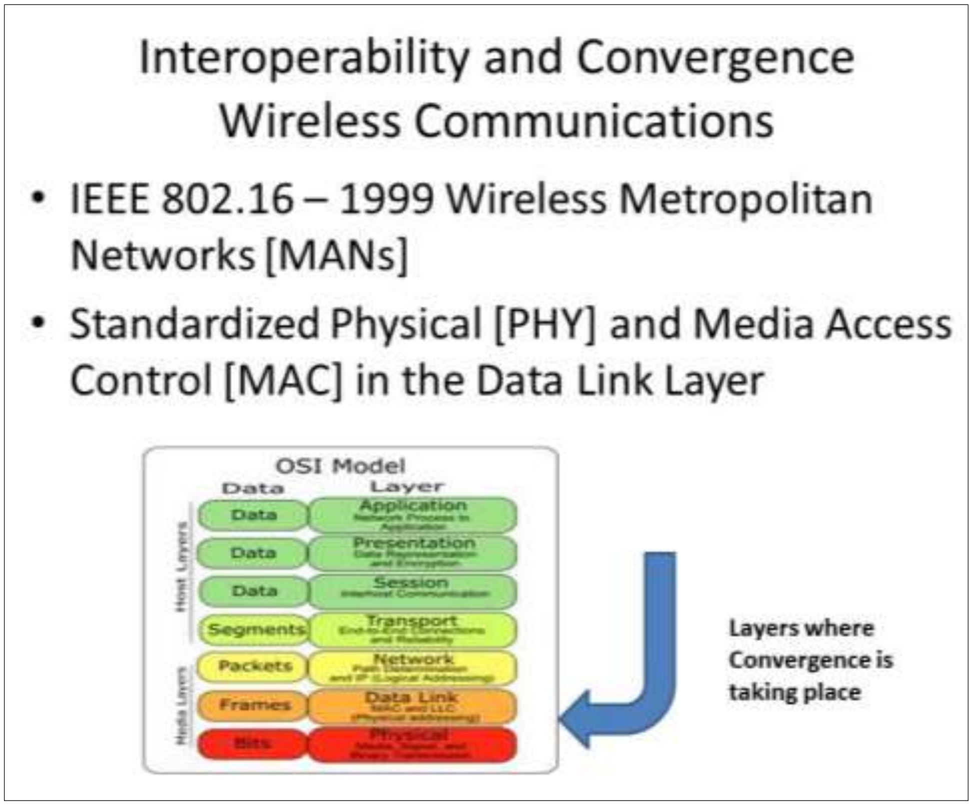

All of these wireless technologies bridge to other technologies, such as a wired technology. They all bridge. The manufacturers involved are going to make sure that happens.

We're back to the interoperability and convergence. I want to talk a second about metropolitan area networks, which traditionally were microwave. Microwave was always a back-haul capability for the utilities to connect substations to far-off generation or control sites. If you're going to converge, you have to do it down here at the physical and the data-link layer. That's what this seeks to do: IEEE 802.16 is the wireless metropolitan network standard.

Right now, they're up to 802.16m, which is mobile; that's wireless interoperability for everybody. This is what WiMax means. This is the mobile version of it, which some call Wi-Fi on steroids. That's about what it is. It's connection based, like I said earlier.

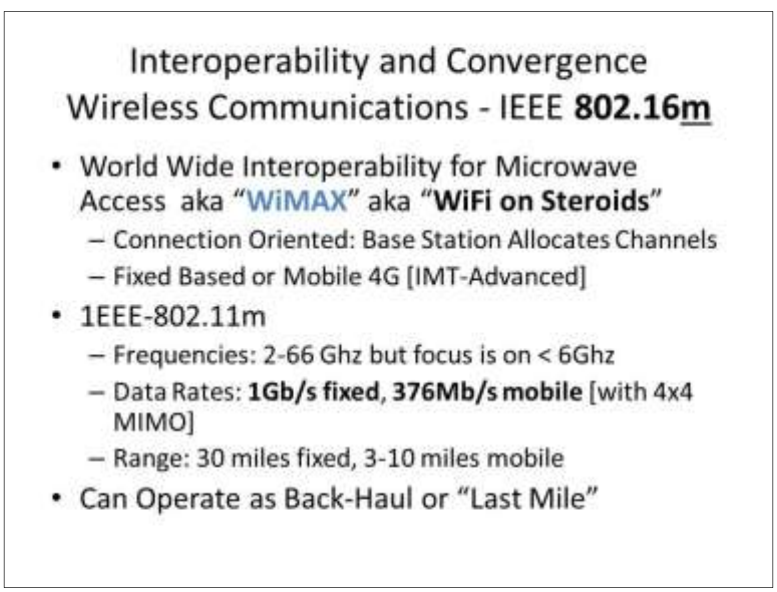

The important part here is microwave is in no longer fixed, but it's now going mobile, which means you can extend the range about as far as you want through two-way repeaters. By definition, they can do anything within 2 to 66 gigahertz. They're focusing on less than six gigahertz frequencies, which is smart. Data rates are 1 Gb/s fixed and 371 Mb/s mobile with 4x4 MIMO. That's not bad, but remember; if you can extend that with mobile repeaters.

What is really interesting is that they've been selected as one of the 4G technologies.

FHSS, frequency hopping spread spectrum, is a technology that was developed by the military.

Frequency hopping spread spectrum is a technology that was developed by the military for a couple of programs back in the ‘80s. We’re talking low probability of interception, low probability of detection, because it was a military technology. They didn’t want it to get jammed. They didn’t want anybody to know they were operating. Basically, take a little frequency sector and chop it up into little bits, then spread your signal over those little channels in the grid within that frequency range.

By the way, they’ve come up with all kinds of ways to hop around so you don’t know what they’re doing. The only way to jam it… is really to jam the whole thing.

There are commercial companies out there doing this now and they’ve been very successful selling this to the utilities to reach remote substations, in particular, Mexico and some other countries. The range is actually frequency dependent, but works quite well. It can go about 60 miles.

IEEE1901.2 is the broadband power line communication. They call it Netricity.

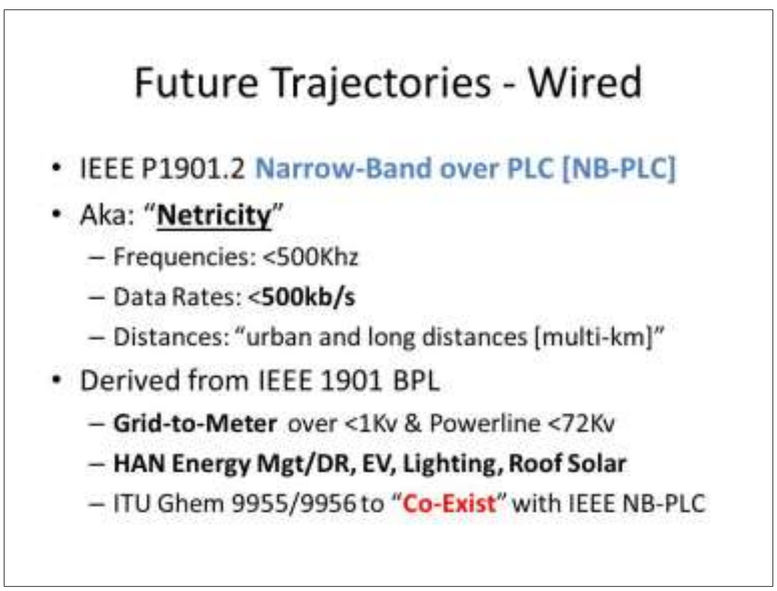

Future trajectories. The wired future, I talked about in IEEE1901.2, which is the broadband power line communication. This is the narrow band. They call it Netricity. Data rates about 500 kilobits per second. They haven’t really said what the distance is going to be, but they hint that it’s going to be reasonable, urban and long distances. It’s derived from IEEE 1901, so is therefore compatible with it.

So appears the wireless, another PLC standard that is compliant with the other. It's great for grid to meter. It will work over any voltage, and it's really going to be good for HAN energy management, home, like we talked earlier.

Demand/response, EV lighting, roof solar, and so forth… are going to coexist right now because of the legacy with BPL.

IEEE P1905.1 is called convergent digital home networking. It is an interesting one. They're going to take four different protocols; Wi-Fi, BPL, Ethernet, and multimedia over co-ax; and merge them at the physical level, at the data-link level. When the Internet or the stack looks down, it will only see one interface, not four. That's going to be interesting.

You're seeing convergence and interoperability.

So what you're seeing here is convergence and interoperability, but the manufacturers can produce their own designs and keep their own proprietary products to make money from their development. At the same time, if they comply with these protocols, we'll have networks where almost anything will run over it.

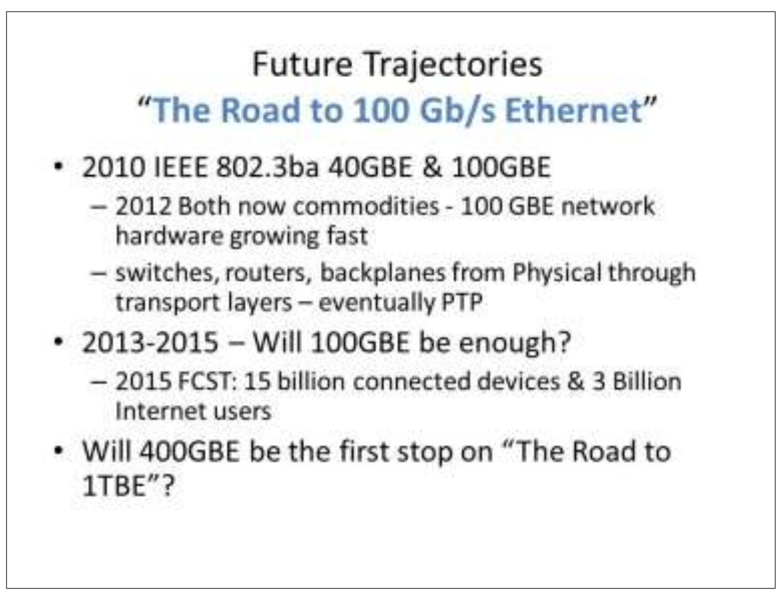

The road to 100 Gb/s is already being implemented with Ethernet running 100 Gb/s over fiber right now. The hardware is growing fast. They're already putting in switches, routers, and back-planes... all the way through the entire protocol stack and by2015 will move beyond that.

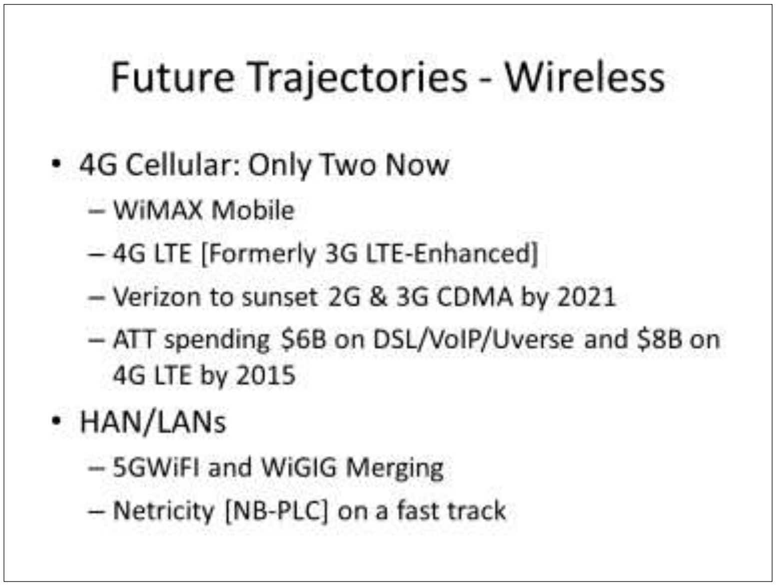

They're already looking at 400 Gb/s and they're going to try to get to a Tb/s. So the forecast right now, and I think one of the presentations said this earlier: we got 15 billion connected devices and 3 billion Internet users. There you go. What did we hear about the smart phones? There are twice as many smart phones as there are people on the planet, now, or we're going to get there soon.

4G cellular. There were seven candidates; they eliminated five of them and they just recently announced that there are two guys now officially 4G. One of them is LTE, which is a GSM standard. The other one is WiMax Mobile. They both evolved from GSM.

Verizon is going to sunset all of their 2G and 3G CDMA networks by 2021.

So what's being left out? CDMA. Verizon just announced a few weeks ago that they are going to sunset all of their 2G and 3G CDMA networks by 2021. They're doing that because they see CDMA as being a network technology that's no longer going to be in the forefront. They actually made that decision back in 2008, by the way.

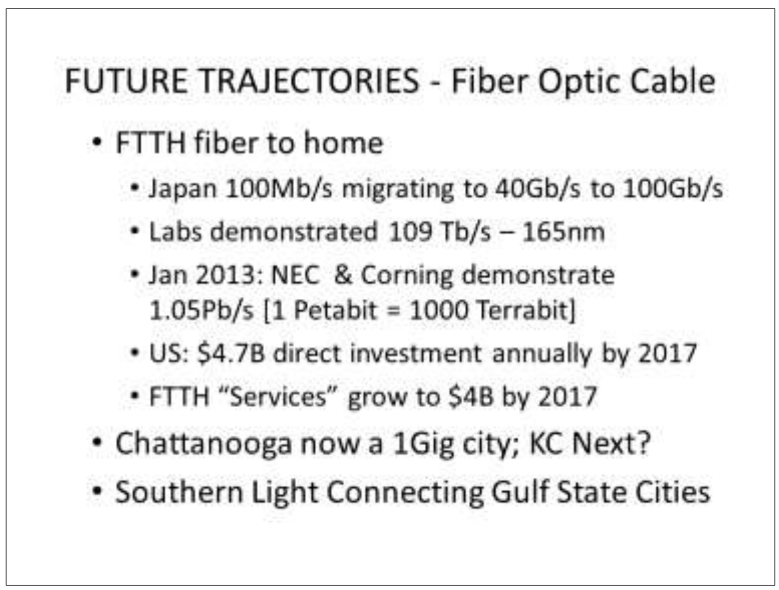

ATT is spending about 6 billion on U-verse and voice over IP VOIP/DSL. They bet on the 4G by 2015. ATT has also indicated that they're on a path never to put a twisted wire pair in the ground again.

Why? Fiber. HANs and LANs 5G Wi-Fi and WiG are merging. That was sort of ordained because they're both wireless technologies, 2.4 and 5, now 60 GHz. Each of these technologies has its own industry alliance pushing it… wanting it to go that way. That's the nature of the game.

These two alliances, the WiGig and the WiMax, have decided, "Hey. We ought to just merge." And they are. That's one reason that some of the products are not going to be out for another year.

Netricity, by the way, is on a very fast track.

Netricity is on a very fast track. Fiber to home? Japan's been doing it for years. We've got labs that have demonstrated now 165 nautical miles and 109 terabits per second. *That* is whistling Dixie.

January 13, NEC and Corning have demonstrated a petabite. A petabite!

The U.S. says that, by 2017, there will be 4.7 billion invested annually in fiber and 4 billion being spent on fiber services. We're talking triple play and quad play, that's what we're talking about.

Chattanooga is now a 1 gig city. Boy, do they make a big deal out of it. So move to Chattanooga and you move to your house, wherever your house is, you instantly got fiber and 1 gig.

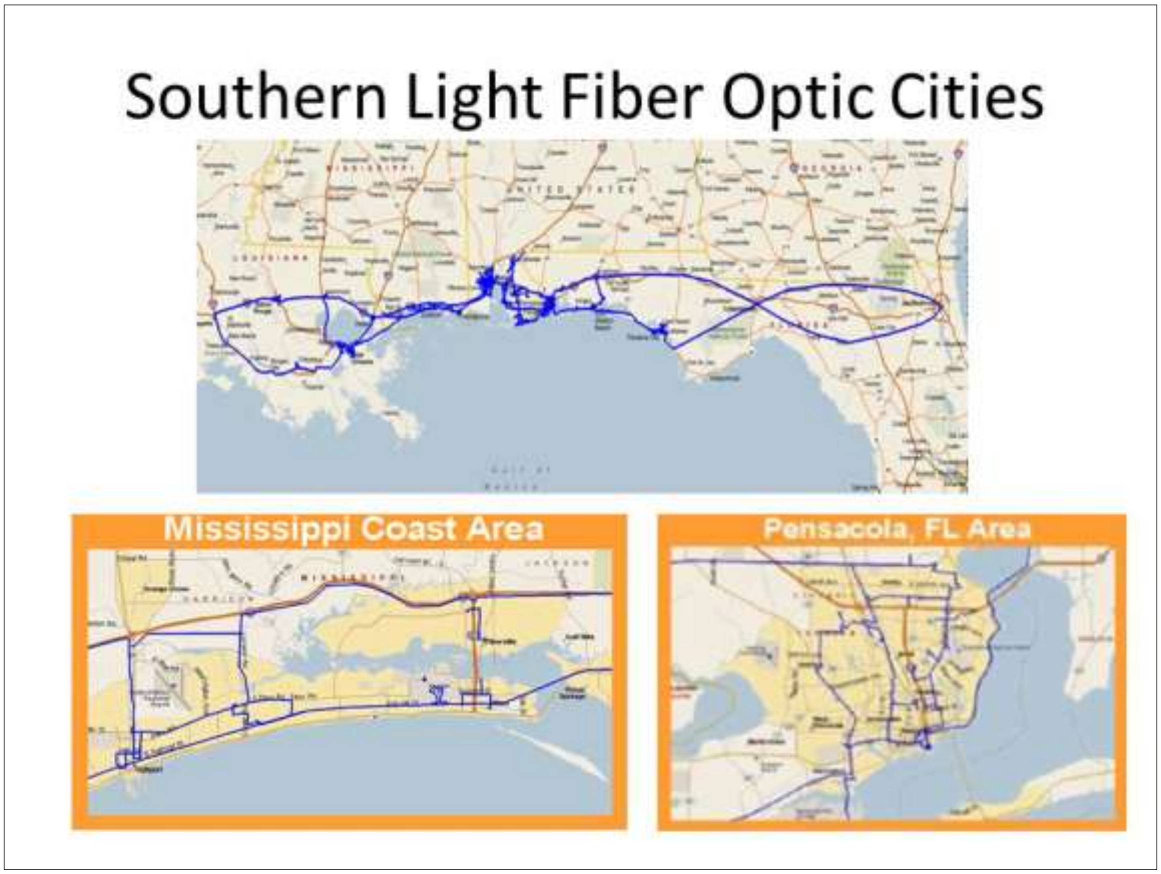

Southern Light is connecting fiber all along the coast, and this is going on all over the country right now. Looking at Louisiana to Florida, with Mississippi right in the middle, fancy that. As these fiber backbones are going in -- and they're going in, believe me – they're branching off and going into the cities. Pensacola, Mobile, and so forth. And they're going to become MANs.

So the population centers are now being connected more and more by fiber. It's just a matter of a time before it comes to your house.

Now the utilities have all kinds of options.

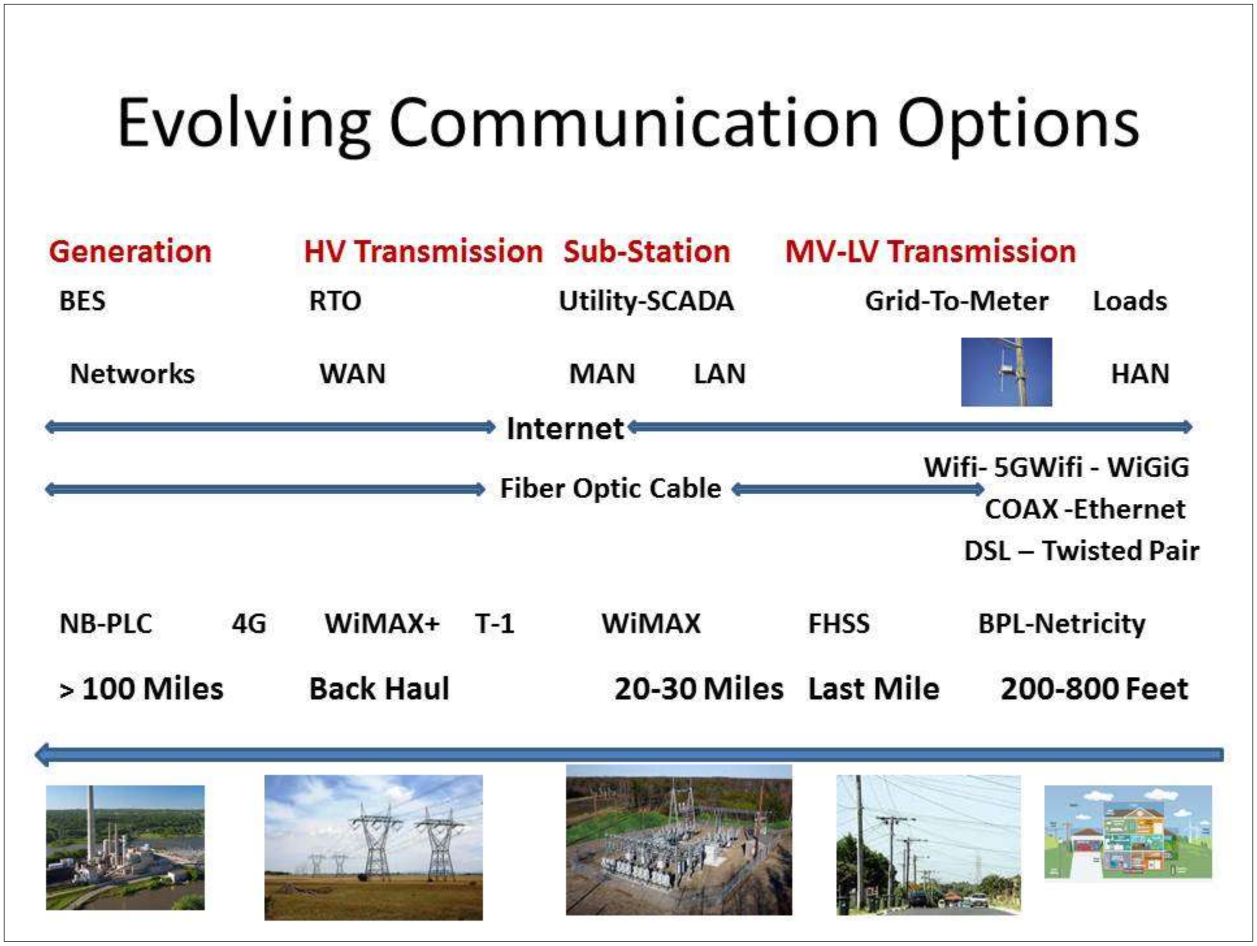

What does this mean?

Well it means we've really gained momentum for three or four reasons in the last twelve years. Policy and consumer demand. So now the utilities have got all kinds of options. A lot more options than they used to have. So if you want to connect to a house, you have a number of different ways you can do that. If this one is more expensive, choose another one. If you go beyond the house and you're out in the last mile, if BPL will go that far, if not, FHSS that I mentioned, and WiMax are both out there, along with some long haul ones. WiMax is one of them: mobile, narrow band PLC, 4G, these are all great. Of course you have fiber optic, you can go about as far as you want to go.

Summary
Evolving Communications [EC]

- Since 9/11, EC has been driven by at least three convergent factors:
  - Federal Policy and Awareness
  - Energy Management and Renewable Power sources
  - Commercial demand for "Triple/Quad" Services
- Technology has been driven by advances in IT
  - Higher Bandwidths, Faster data rates, longer distances
  - Competition among technology loyalists
- The Power Industry has:
  - More communications options available
  - To network their assets within their cost & performance constraints

On the top, you see the Internet. I think that's a great thing to do too. But not for your control system. Jonathan is going to talk about that.

That's kind of the summation. Since 9/11, evolving communication has really been driven by these three factors: federal policy and awareness, energy management and renewable power scores, and the commercial demand for triple/quad services.

Competition gives you more choices, and it gives you lower costs.

Technology has been driven by advances in IT: hyper bandwidths, faster data rates, longer distances. This means competition among technology loyalists and that's great. I love competition. Competition gives you more choices and it gives you lower costs. Fact of life. The power industry now has a growing number of options to network assets within cost and performance constraints. As you well know and probably more than I do, that for the utilities in the United States, one shoe does not fit all. They have their domains and cost structures they have to operate within.

That's it. If you ever come to Mississippi, give me a call.

Hacking SCADA — Are You Ready for Stunt 2.0?

DUSTY BECKER:

Next we have Jonathan Pollet, Founder and Principal Consultant for Red Tiger Security, USA with over 15 years of experience in both industrial process control systems and network security. After graduating from the University of New Orleans with honors and receiving a B.S. degree in Electrical Engineering, he was hired by Chevron where he designed and implemented PLC and SCADA systems for onshore and offshore facilities.

In 2001, Pullet was one of the first to publish several whitepapers that exposed the need for security for Industrial Control Systems (ICDS) and is still active in the research of vulnerabilities within critical infrastructure systems. He has led security teams on over 200 assessments, penetration tests, and red team physical breaches involving SCADA and Industrial Control Systems. He is also the developer of the 5-day SCADA Security advanced training course and has trained over 1,300 professionals since 2009.

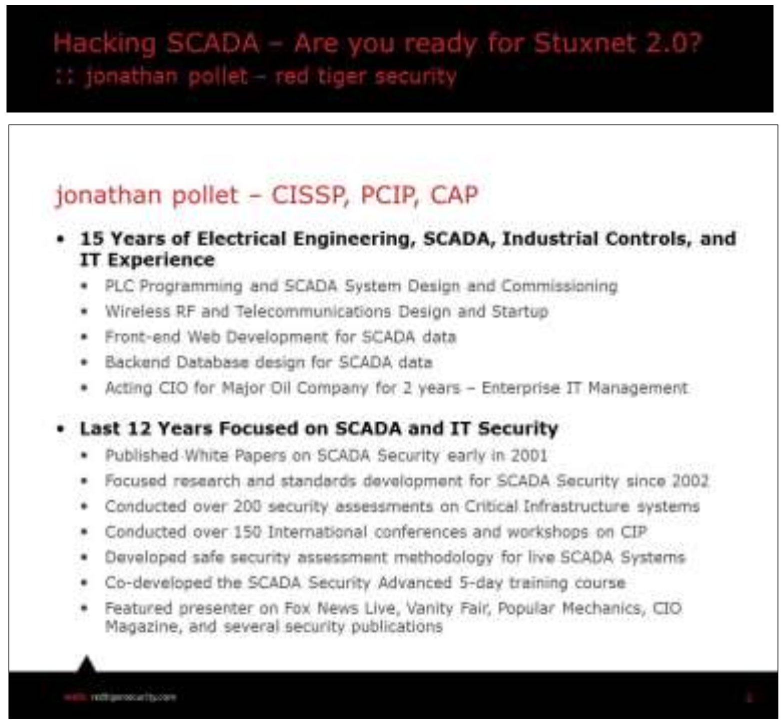

Throughout his career, he has been involved with SANS, IEEE, ISA, ISSA, UTC, CSIA, SPE, and other professional societies. Pollet has developed and presented workshops on SCADA security to the FBI, Department of Homeland Security, and the Utility Telecom Council, and has spoken at over 125 conferences and workshops around the world.

Pollet has been quoted in many periodicals, trade journals, and magazine articles as an expert in SCADA and critical infrastructure protection. He was recently interviews in the September issue of *Vanity Fair* and appeared on *Fox News Live* for an interview concerning hackers and critical infrastructure.

JONATHAN POLLET:

So my background is in electrical engineering. In the 1990s, I worked for Chevron; programmed some PLCs, SCADA systems, telemetry. Then we had a rash of things happen to some of our plants in late '90s and I became instantly aware of cyber security and what packets on networks do to control systems.

So in 2001, I shifted from control systems and engineering into cyber security and I've been doing that since then.

I run a small company. We specialize in consulting, assessment training, and research. The training side of our business has really taken off in the past two years. People want to take their employees' knowledge of cyber security for SCADA and for smart grid and raise it up. We're going to be at a couple of events this year, if you're interested in obtaining more hands-on training in cyber security. We can hook that up on the website.

A lot of the threats and hacking techniques have been shifting from IT systems over to SCADA and process-control systems.

I want to just touch on why cyber security matters today, then shift into some things that are happening in the industry. You've probably seen things in the news, as well, that show a lot of the threats and hacking techniques have been shifting from IT systems over to SCADA and process-control systems. Smart grid meters, as well, have been under the spotlight. Then we'll kind of build toward why this is an issue now and where it's going.

In 2004, the Sequel Slammer worm was taking out 75,000 machines within ten minutes.

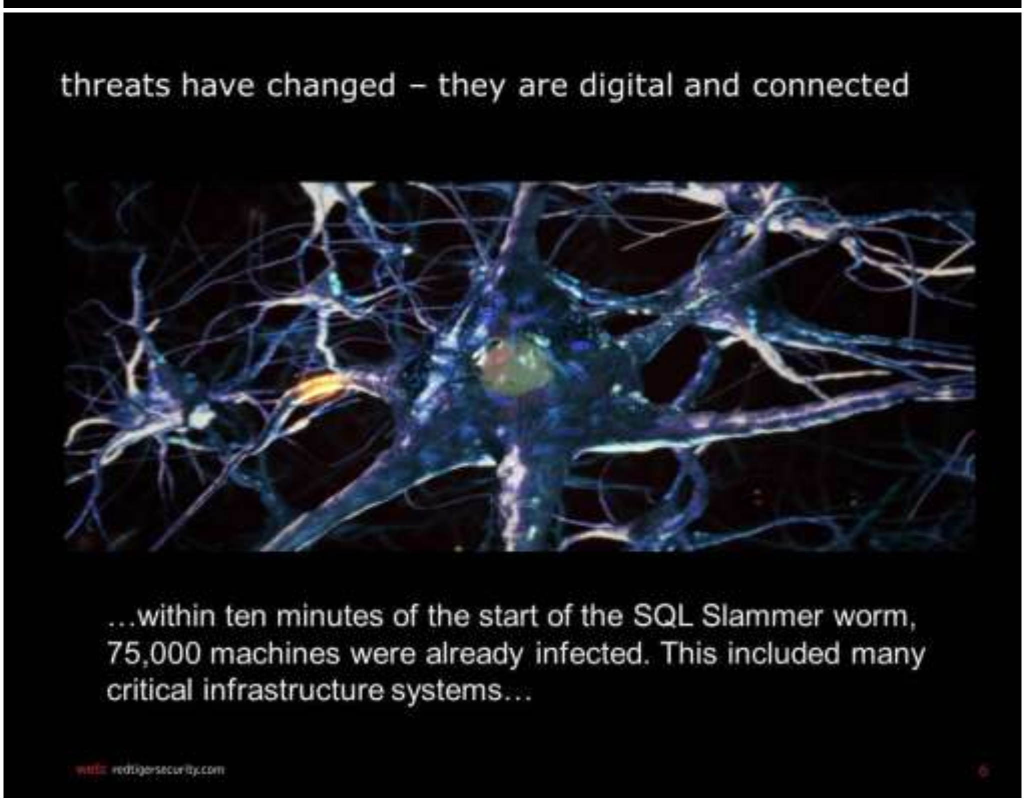

As the last presenter said, it is a digitally connected system we live in today. Malware spreads extremely fast. In fact, in 2004, the Sequel Slammer worm was taking out 75,000 machines within ten minutes. This stuff is just opportunistically spreading through open firewall ports. It doesn't really care what it is. If it has vulnerability, it will blue-screen it.

I was flown offshore to several offshore platforms. I went to a couple of plants, onshore plants as well, where the control systems where crashed because of the Sequel Slammer worm. The guy who wrote Sequel Slammer didn't even know what a control system was. He was a 17-year-old kid who didn't like Bill Gates and Microsoft.

Our control systems, SCADA systems, and smart grid platforms are being built on top of Microsoft and on top of the same TCP/IP protocols that we use in the corporate IT infrastructure, so we have similar problems to deal with.

> *Hackers use a "botnet" army of computers. If you're going to do something nefarious, why spend money to buy computers and bandwidth when you can steal someone else's computers and bandwidth?*

The type of attacks that we're seeing is shifting. About five or six years ago, you probably saw a lot about "Distributed Denial of Service" attacks. Basically, hackers would take a botnet army of computers -- your home computers, if they've been infected — and point them all at someone's web presence and take them down. Distributed denial of service is a way you take a thousand different computers and harness that power to take over someone's website. Because… if you're going to do something nefarious, why spend money to buy computers and bandwidth when you can steal someone else's computers and bandwidth?

How many of you guys have Xbox 360s at home? Interestingly enough, when they were doing the contract negotiations between Microsoft and either NVidia or Matrox, Matrox's website was not available for three days. NVidia got the contract because Microsoft said, "Well, if you can't even keep your own website up, then we probably shouldn't be doing business with you."

A lot of us in the cyber security were like "Hmmmm…. Kind of questionable that Matrox's website went down just as the contract negotiations were underway."

Now, a lot of the threats are very stealthy: Get onto someone's network, maintain a presence, and extract as much information as possible.

There's usually a reason why things happen in the cyber security world. Nowadays, the threats aren't distributed denial of service because, for one, you can tell when your website is down very quickly. You can see this barrage of attacks coming at you. Now, with a DGP router, you can divert the attack. You can get around it. Now, a lot of the types of threats I'm going to talk about have to do with being very stealthy about how you get onto someone's network, maintain a presence, and then extract as much as information as you can.

So it's like looking in, then hooking on like a leech or a parasite and just extracting as much information as possible. A lot of times, it's intelligence based.

This is like a trick question. Do you guys know who this girl is? She's not a real girl. Well, I guess the girl was who's in the picture, but this guy put up a profile for someone called Robin Sage. He put a nice, pretty picture of a girl up and started contacting government agencies saying, "Hey, I'm new in this agency. I need help. Can you let me know how to move up and who I should talk to?"

Within two months, this fake profile had over 500 contacts sending her sensitive information!

Within two months, this profile had over 500 contacts sending her very sensitive information. It wasn't even a real girl.

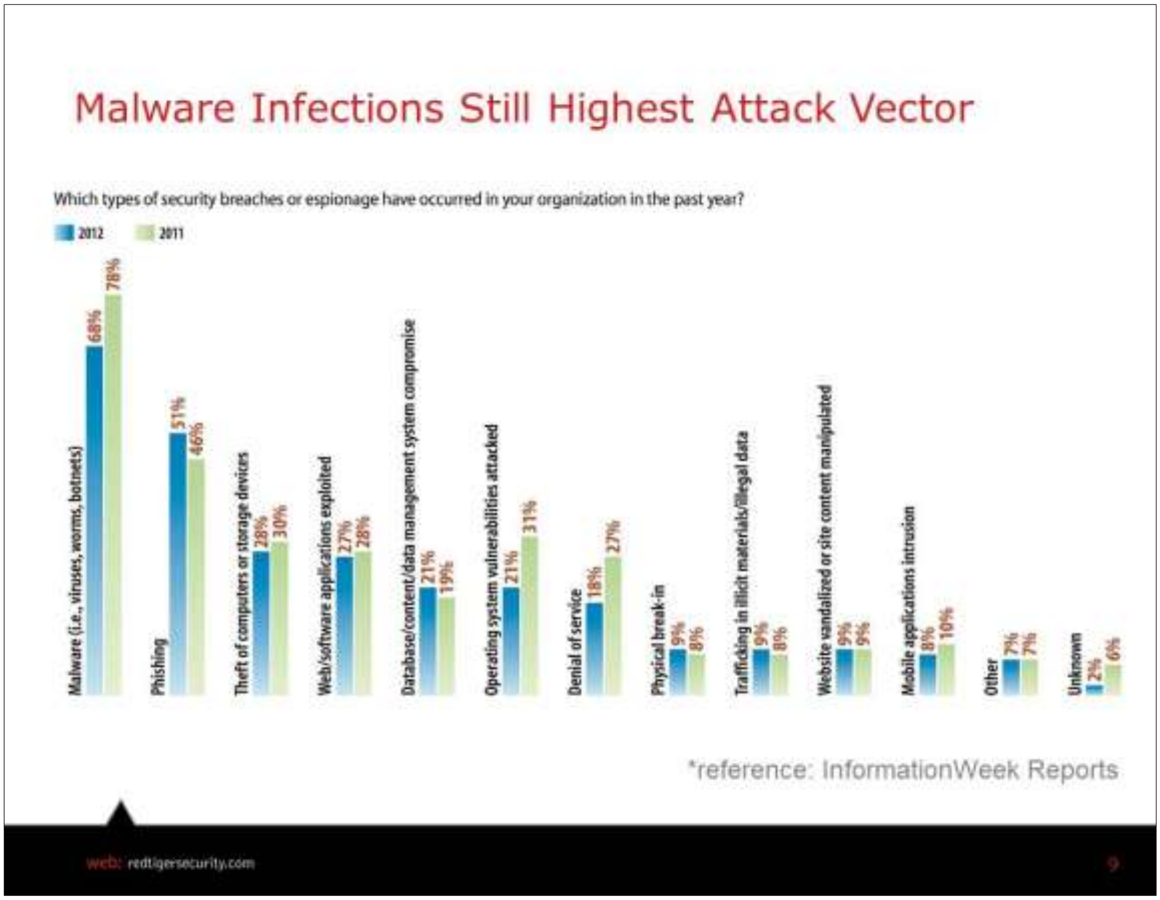

What we're seeing as the biggest problem is not really physical break-ins; it's not really these other things down here; even denial of service dropped down. What we're dealing with now is malware and phishing techniques. That's the biggest problem because in the typical bell curve of workers, there's always going to be someone who is going to click on the PDF file or on that little link… and there's really nothing you can do about it. Someone is always going to click on something they shouldn't.

You can't really just stop it at the border because you can't block .exes. You can try to block as many extensions as you want, but they're coming in on PDF files, .dot files, and Excel files.

Someone is always going to click on something they shouldn't have.

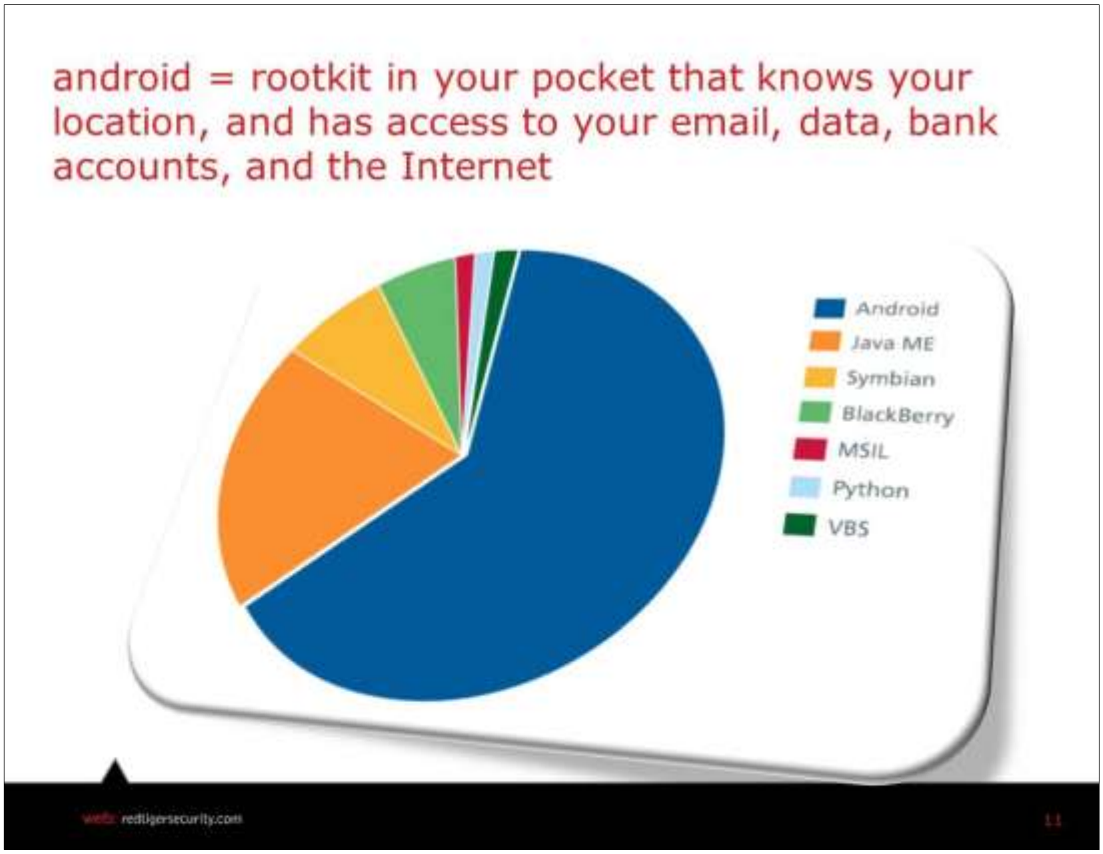

Then we walk around with cell phones in our hands. If you would have asked me back when I was in college -- this is before cell phones and I was lucky enough to have a pager back then -- if you were to tell me that there would come a time where everyone would have their own phone number and carry it on this little device… and the device would know where you are in the world and have access to all your bank accounts… and know everything about all of your business -- your personal business and your computer, your professional business — and it's always connected to the Internet, which is a network that everyone else is connected to; I would have said "That's a great device to have control over if I was a hacker because this device could give me access to anything that person has."

A problem for a lot of the local phone companies is that it's really difficult for them to keep malware off their phones.

For a lot of the local phone companies, it's really difficult to keep malware off their phones. Especially Android, because it's an open source and it's really difficult to stay on top of 600,000 apps that are constantly being updated and changing their play environment.

This is also a concern in the corporate world because all of these devices can become hotspots. It's really difficult to force everyone through web sense, you know, do proper filtering when everyone is jumping on their own little hotspots connected into their corporate laptop.

We go out to power plants and check control centers and we see operators with their phones plugged into their operator console. They're not aware that they're bypassing all of their cyber security controls and connecting the Internet directly into the plant.

When we go out to power plants to check control centers and we see operators with their phones plugged into their operator console, we ask: "Do you realize that this is a 4G LTE device with high-speed Internet... and it's plugged right into the plant's control system?" They tell us that they're just trying to get juice on their phone, to charge it. They're not aware that they're bypassing all of their cyber security controls and connecting the whole Internet directly into the plant.

We have to worry about our phones, now, because they're not just telephones; they are complete computing systems.

So now we have to worry more about our phones, because they're not just phones; they're complete computing systems. For instance, this was the one of many applications pruned from the Android marketplace because, while you're playing the game and bowling with your phone, it actually is sending out all of your information outbound. You can read more about this; it's call the Droid Dream.

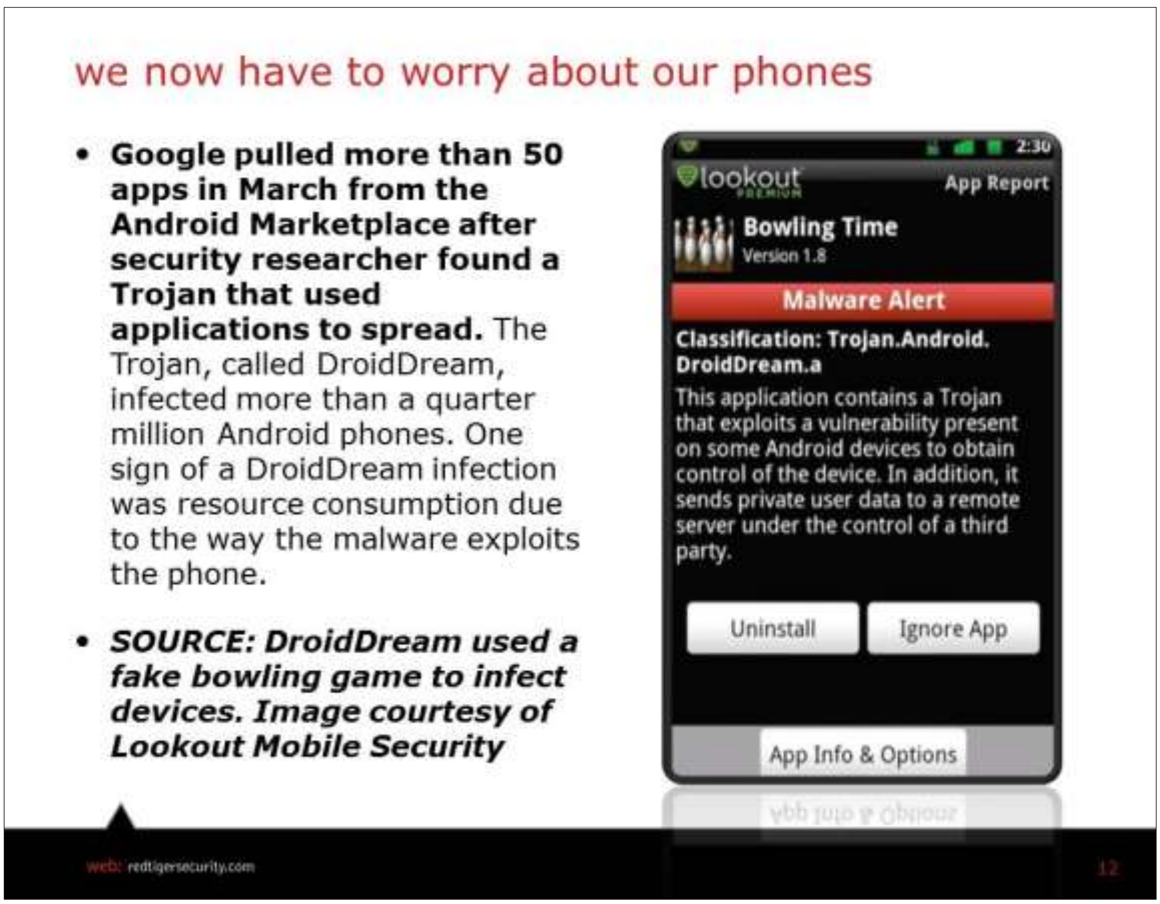

This was what a control system looked like in the 1950s. All you had to worry about was physical access. If I could be in the room, I could push buttons and operate a substation. *IF* I could be in that room. All we had to worry about was physical security.

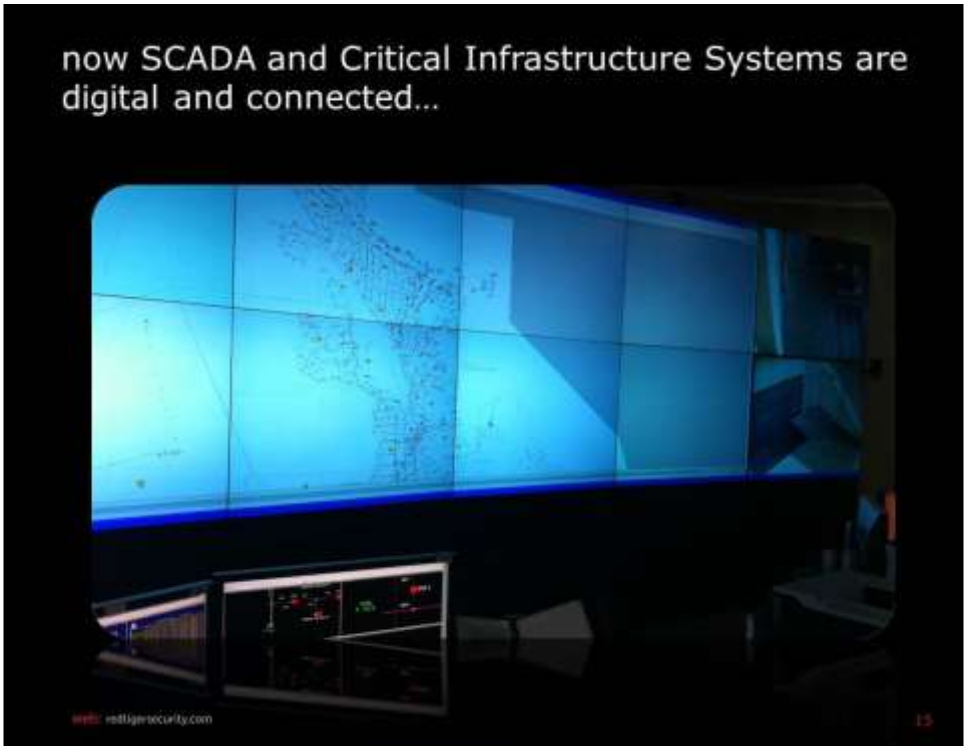

This is what a control room looks like now. Computers running Microsoft® Windows, connected to everything, connected to a sequel server, SAP, GIS, management systems, market trading systems… all connected to generation systems and connected to distribution systems. It's all connected and integrated because it makes sense on the business side.

It's running the same software as our corporate PC, but it doesn't afford the same level of security as our corporate PC. It progressed with the IT industry; as it developed the technology, the security progressed with it. For instance, when you log into your corporate desktop, you know you need a user name and password that's unique to you. You know that every 30 days you have to change that password. You know that if you want to access a certain file on a server, you have to put in your user name and password to access it. Sometimes you have to have two factual identifications if you're working remotely. You know that if you're going to connect to something that sensitive, there are probably going to be certificates involved, like SSO or HTTPS.

SCADA process control and smart grid applications have grabbed onto the IT infrastructure, but the security is not yet built into the protocols. It's a ten-year-old protocol running on today's technology with no improved security.

We take all of these things for granted because the security has evolved with the IT infrastructure. Unfortunately, SCADA process control and smart grid applications have grabbed onto the IT infrastructure because it's low cost, has good standards, and it's good equipment; but the security is not built into these protocols.

So we have all of these nice, shiny Microsoft-based systems; but the underlying protocols do not have authentication or any encryption. It's really like a ten-year-old protocol running on today's technology.

Making a dinner reservation has more security than to send a command to an RTU.

If you ever sniff SCADA protocols and capture the traffic or if you watch DNP or Modbus; you watch these protocols and all of the characters are in clear text. You can actually just read KW this, KVAR this, power factor this. It's clear. It's not encrypted. It's not encoded. You can just read it. You can't do that with your banking information. In fact, if you make a reservation at a hotel tonight to go eat somewhere at a restaurant, your reservation process and making a dinner reservation is going to have more security than to send a command out to an RTU. If you want to make a reservation, they're going to want to know something about you: your name, maybe your room number. When you connect to that website to set up a reservation, that website is going to have SSO and HTTPS certificates. It's going to be encrypted. There's more security wrapped around the normal things we do day to day in our lives than there is in sending out a command to an RTU, which is in the clear text, with no certs, no authentication.

Now, if I were to describe to you that you're on a network that had Cisco, firewalls, Cisco switches, servers running Windows, active directory and file print servers… you'd say that's a business network. That's typically what this looks like.

When you're buying a SCADA system today, a DCS system, today, an EMS system today, smart grid system today; this is what you're getting. You're getting Cisco here, you're getting servers running the same platforms as IT systems on the same active directory infrastructure, print and file servers. It looks the same.

While this critical infrastructure is on a parallel path with IT, it's not yet ready. It doesn't have all of the security features yet.

We are not the only ones who know about this. Researchers know this too. When I was back programming PLCs in the 1990s, I could say, "Look. Nobody really knows about Schneider PLCs or GE RTUs. No one is really going to know these weird protocols like Contel or DNP or Modbus." You know, I didn't really worry about somebody hacking into plants.

Nowadays, if you just type in "Modbus" or "SCADA exploits" or "SCADA vulnerabilities" or "SCADA security" …you get pages, tons and tons of people researching this, exposing this, and writing about it.

So we wonder, "why?" Why is this the case? It's because Microsoft puts their equipment and their software and hardware through thousands and thousands of regression tests. They fuzz the crap out of their software. They go through secure coding techniques to make sure that when you buy the product, it's already been tested. Yet we still have to patch Microsoft products every week.

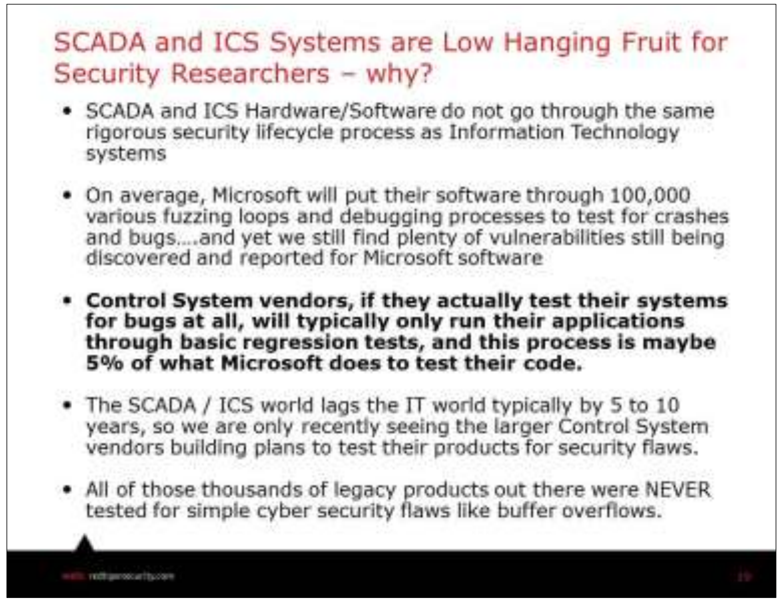

So if a big company like Microsoft, spending a lot of money to make secure products, still can't get it right; what do you think smaller mom-and-pop company control systems, vendors, and smart-grid vendors are doing about their product? Do you think they're actually putting it through the rigorous security testing Microsoft is doing? Definitely not.

There's a huge install base of SCADA and control system products that were never tested for simple flaws.

Even if they started doing this today, deployments are being launched for 10-, 20-, and 30-year deployments. The real-time world lags IT by 5 to 10 years. There's a huge installed base of SCADA and control system products that were never tested for simple flaws. We all know that in this industry.

Two guys, McCorkle and Rios, set out to prove it. One worked at Google and one worked at Boeing. Interestingly enough, they'd never even seen a SCADA system before. These are IT guys who said, "Well, if everybody says this stuff is vulnerable, let's see if we can find bugs in it."

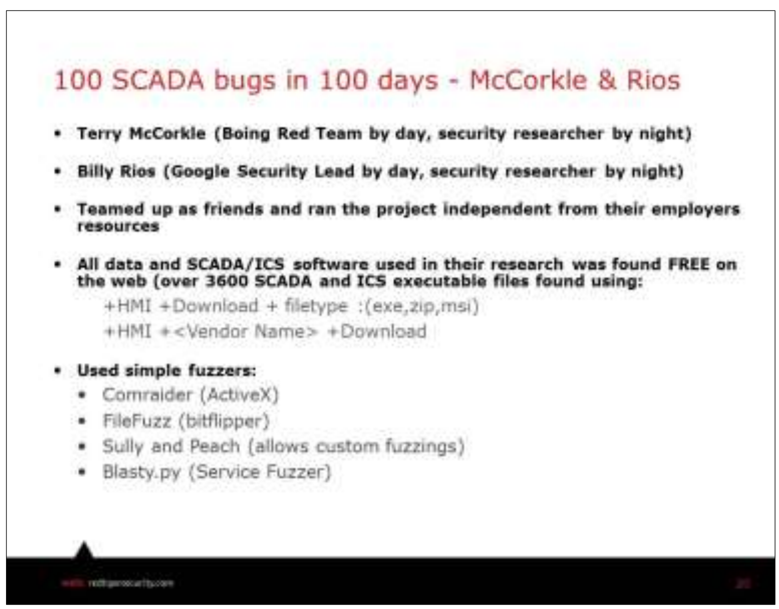

So they did a Google search to try to download SCADA software because they didn't have any money to do this research project. They found over 3600 executable SCADA and ICS files that they could easily, freely download off of the Internet just by typing, "HMI" or "Workbench" and downloading files.

McCorkle and Rios were able to download 380 packages. They only tested the first 76 and they found 665 bugs.

So they started downloading all these files, like Workbench files. They used the program PLCs to create HMI graphics. They ran them through these simple little fuzzers, where you fuzz the software, looking for bugs and crashes.

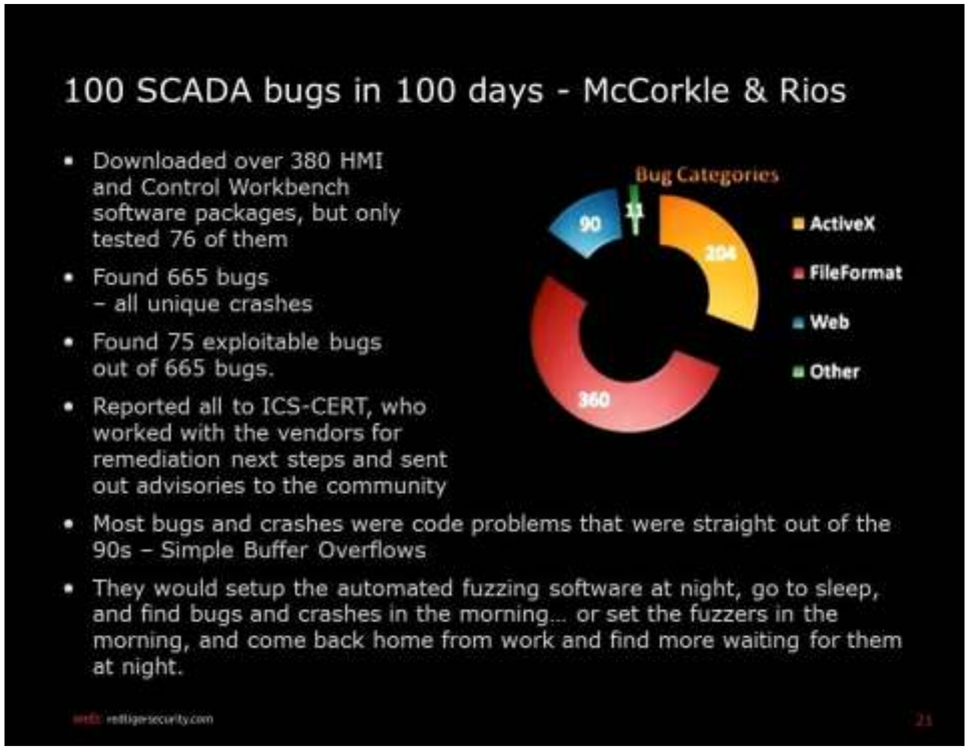

They were able to download 380 packages. They only tested the first 76. They found 665 bugs, far more than the 100 they thought there were. They were doing this at night. These guys had families and jobs, so they were just doing this at night.

So they sent all these bugs over to ICS certification. The certification agency got in touch with all of the vendors involved and said. “Look, you really need to fix these bugs.” The vendors wouldn’t have known about it unless these guys had done this project.

There’s a lot of other software and hardware out there with vulnerabilities.

That tells us that there’s a lot of other software and hardware out there with vulnerabilities. Most of these are straight out of 1990s: simple buffer overflows, sending a field with more characters than it can support, doing a sequel injection. This is the stuff we had to deal with on the IT side a long time ago. These two would just set up the fuzzing software, go to sleep at night, then wake up in the morning to a pile of bug crashes.

Over the past year, there’s been an increase of 753% in the number of vulnerabilities disclosed to the ICS cert.

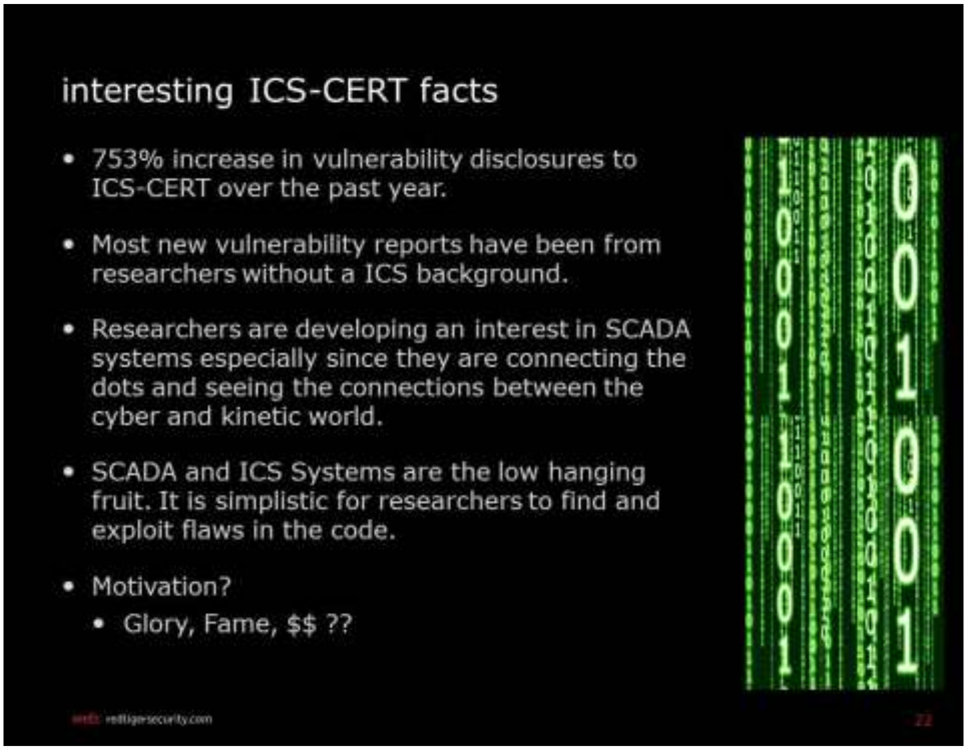

Over the past year, there's been an increase of 753% in the number of vulnerabilities disclosed to the ICS cert, the Industrial Control System cert. A cert is an agency within the U.S. government, underneath DHS, that collects all of these threats, advisories, and alerts. The cert puts them through a triage, finds out if they're real or not, contacts the vendor, and tells the vendor to fix it. That's basically what they do… and they send out advisories to the community.

Think about your own salary for a moment. Raise it 753%. This is a big tell-tale sign. It's not like 10% more vulnerabilities were disclosed this year than last year. 753%! That means there's a lot of more focused activity around trying to find security flaws in SCADA software and hardware.

If I can crack into this SCADA system, I can actually have a kinetic effect in the real world, like turning big motors on and off.

These researchers are really interested in SCADA because they understand that I'm not just stealing someone's identity if I crack this database. If I can crack into this SCADA system, I can actually have a kinetic effect in the real world, like turning big motors on and off or turning out the lights in a stadium. That kind of gets them excited.

So why are they doing this? We have a slang saying in the U.S. that says: "It's all about the Benjamins" meaning it's all about the money. 5 or 10 years ago, hackers would just do this stuff in their basements. If they were living with their mom, they could impress their friends because it was cool.

But this actually takes a lot of work. These guys want to get paid. They want to drive nice cars. They want to have their own apartments. So this guy Luigi, based in Italy, spent a lot of time discovering 34 different SCADA specific vulnerabilities. He went to GE and said, "I have a pile of vulnerabilities on your hardware and software. I'll give them to you for a price."

GE was like, "We don't negotiate with hackers." Just like the U.S.: "We don't negotiate with terrorists." So they refused to pay for it because they said, "It's a slippery slope. If we pay for these vulnerabilities today, you'll go off and find 50 more, then I have to pay you again and again." So GE said "We're not going to pay you for it."

He eventually sold them to ZDI. After that, he found more vulnerabilities. ZDI wouldn't pay for them, GE wouldn't pay for them, and he couldn't find a broker who would pay for them. He just got pissed off, then he basically threw them all up on his website all at once.

"The source code for exploiting all of my RTUs is on the Internet now. What do I do about this?"

I had clients, utility clients, who came to me the next morning and said "The source code for exploiting all of my RTUs is on the Internet right now. What do I do about this?"

That's the world we live in today. We can't hide behind the fact that our stuff is proprietary and different. The security thing doesn't really help us anymore. We have to be able to secure our systems even if the source code is out there; even if someone knows about them.

There is a growing number of researchers and buyers out there in the market. You ever use a website called Stubhub? Like to buy concert tickets or something like that? Check out Exploithub. The basic thing is they keep a closed mind to this, saying: "I know there are people out there developing exploits over here. I know there's people over here in nation states buying the exploits, and we're just going to take a 5% cut off of the transactions."

You don't even need to create your own exploit. Just put in your credit card number and you can get the exploit of whatever system you want.

ReVuln is pretty brazen. They have SCADA system vulnerabilities for sale to the highest bidder.

This is a company called ReVuln. They just came out with this website about six months ago. They were pretty brazen and said that they have SCADA system vulnerabilities for sale to the highest bidder.

Go to this website as well, ReVuln.com (http://revuln.com/), and this little wheel shifts is interesting. It shifts around. [5/1/13: The web site has changed]. You can click onto these different categories: Enterprise, mobile, these are different platforms that they believe have vulnerabilities and exploits. Look at the biggest section of the pie. It's SCADA.

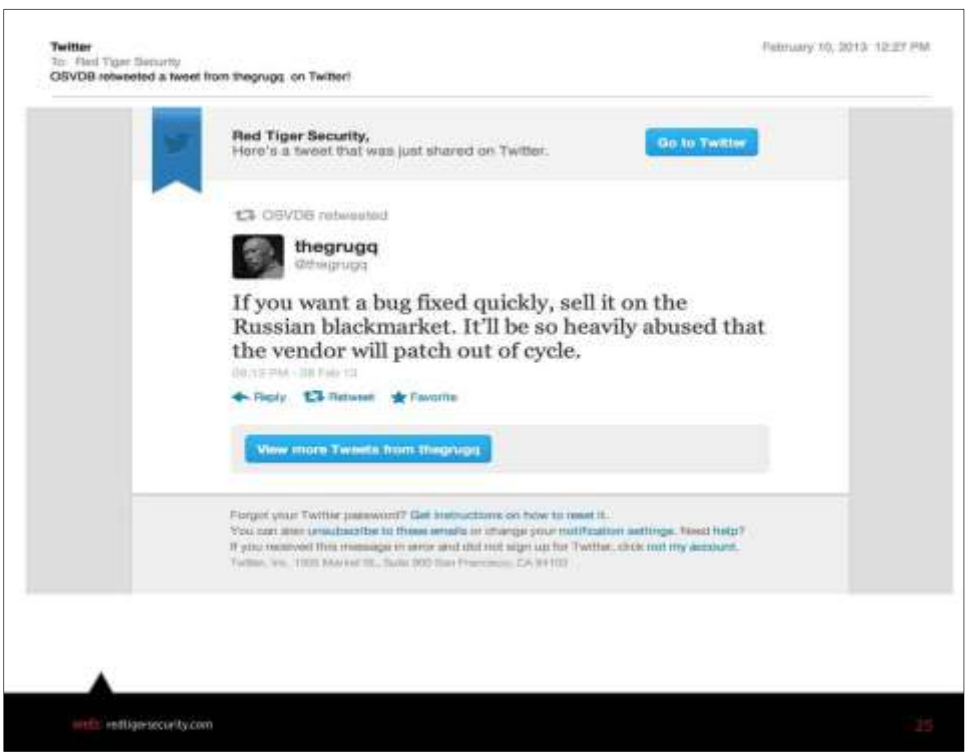

If you want a bug fixed quickly, just sell it on the Russian black market and that will just force the vendor to patch it out of cycle.

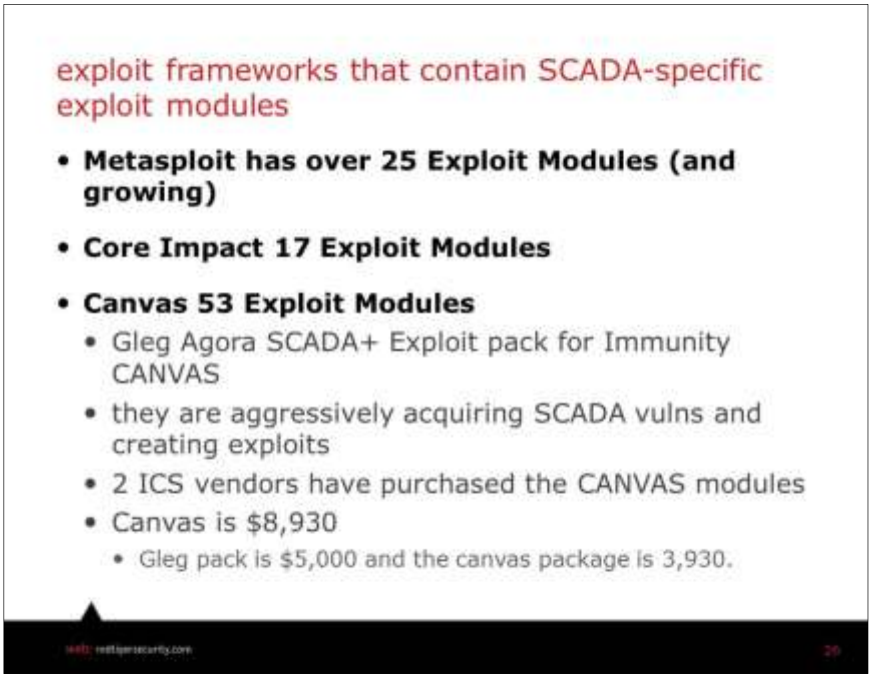

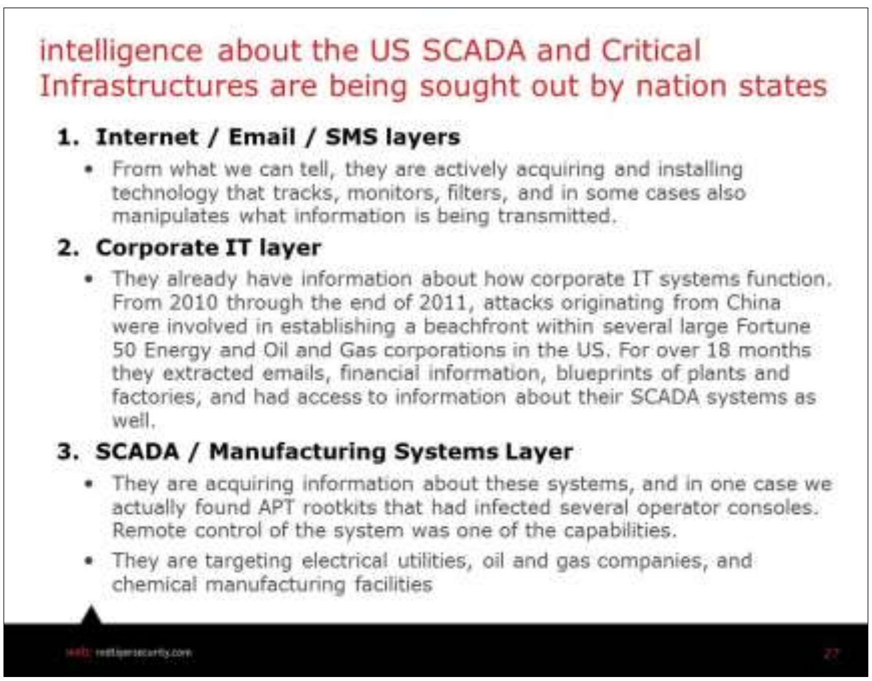

We're finding that the end result is that a lot of these exploit modules are showing up in the pen tester's toolkits, which means the hackers are getting this stuff. These toolkits, some of these, like Metasloit, are free to download off the Internet. These will cost you: Core Impact and Canvas.

I'll try to frame it for you. How many of you have traveled to other countries outside of the U.S. and your Internet is somewhat filtered? You don't really get the full Internet. I was just in Qatar and I was in the UAE about a month ago. So what happens in these nation-state countries is that they actually control all of the ISPs in their country through the government. So they actually track and monitor anything that's done over the Internet. You just have to assume that when you're in those countries, they can not only track and monitor and filter. In some cases, they can manipulate the information routing between their different ISP interfaces.

Energy, oil and gas companies, and pipeline companies have gone through a lot of attacks over the past two years.

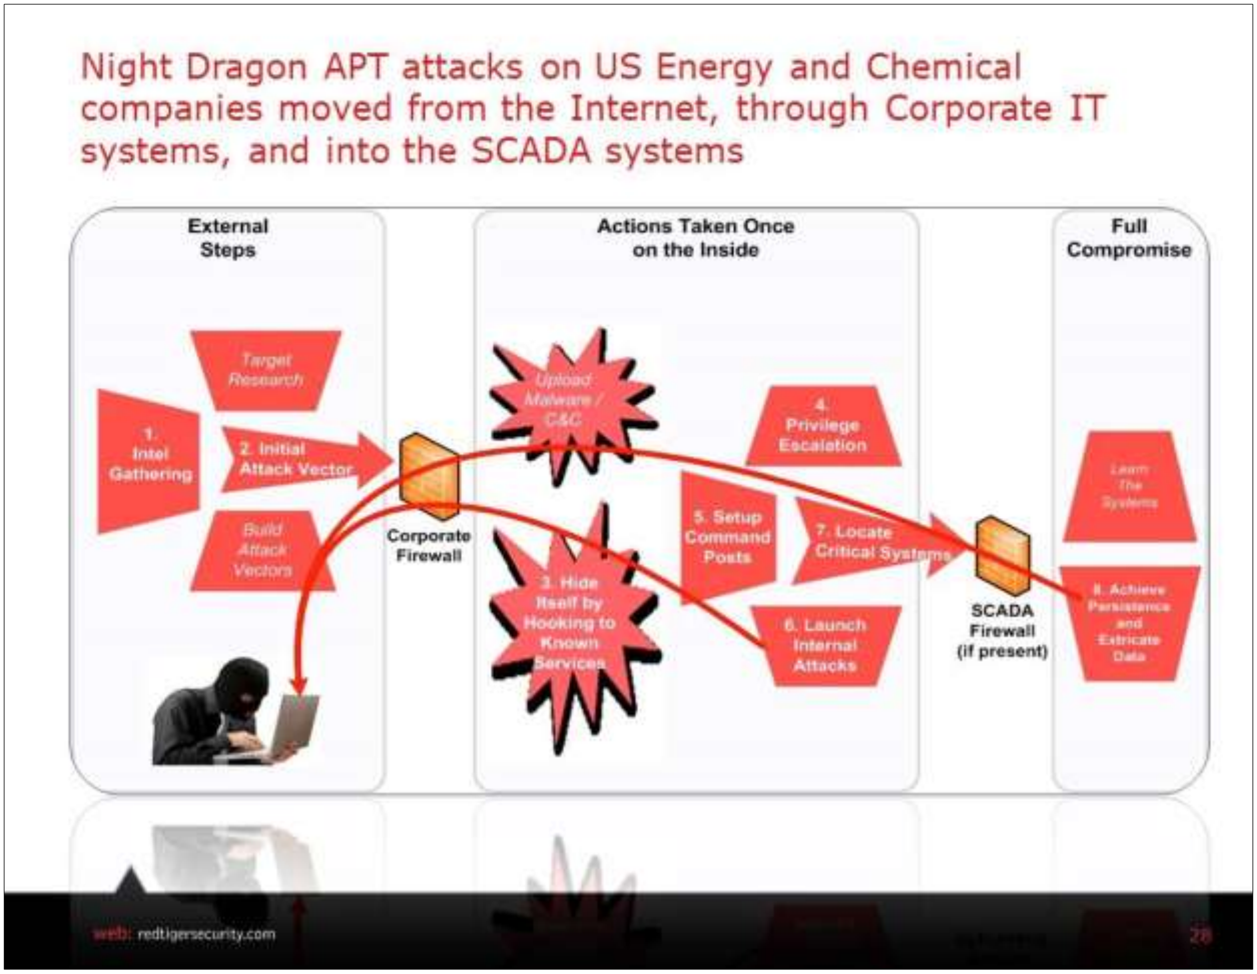

We also know that other nation states are getting involved in the corporate layer because of what's been happening in the U.S: energy, oil, and gas companies, as well as pipeline companies have gone through a lot of attacks over the past two years. Night Dragon was a huge campaign against energy companies. There was another campaign against pipeline companies that persisted for about 12 to 14 months.

There was a reason. I don't know if you kept up with what was going on, but basically, the Night Dragon attacks were focusing on chemical companies and energy companies. Then, interestingly enough, like right after all those attacks died away, we found that there were several very large acquisitions of different plants in China; that the acquisitions went to Chinese companies, and they outbid the U.S. companies within very small margins.

17 major U.S. pipeline companies were hacked through the corporate layer.

What happened is the pipeline industries? 17 major U.S. pipeline companies were hacked through the corporate layer. Targeted emails were sent to the CEOs and the C-level managers. Once their desktop was compromised, there was a dropper put on the desktop which would allow the hackers to come back into the systems each day and extract emails, different blueprints, financial information on how pipelines function, etc.

Why so much interest in pipelines? It wasn't just one or two pipeline companies; it was 17 U.S. pipeline companies. So that whole campaign of hacking into pipeline companies died off. Then, about a month later, we saw a big splash in the news when a Venezuelan pipeline operation bid went to a Chinese company. So nation states are using these tactics to better themselves, to acquire information and Intel.

We found that some of the root kits actually got onto some of the control system consoles.

We're also finding out that when we went into these systems to do the investigation and to actually try to clean up the systems; we were involved in three different efforts. We found that some of the root kits actually got onto some of the control system consoles. Fortunately, they didn't know it was a control system console because it was just a Windows XP machine, just like all the other machines. They were looking for certain file extensions; they were not looking for real-time ties. They were not looking for PLC files, they were looking for emails with financial information. Had they known that these were live SCADA operator consoles, they could have had remote control of these boxes and done whatever the operator was doing, sitting in front of them.

This is just how the Night Dragon attacks work. Basically once one machine on the inside is compromised, it would spread at night. It would pull all of these together into a TAR file, send it outbound, and then spread it. We actually found someone on the SCADA side of the firewall over here. We put a sniffer on this side of the firewall and were able to find out what the destination that was sending all this outbound. Then, we went backwards from there to find other compromised targets. This was pretty rampant in 2011 and 2012.

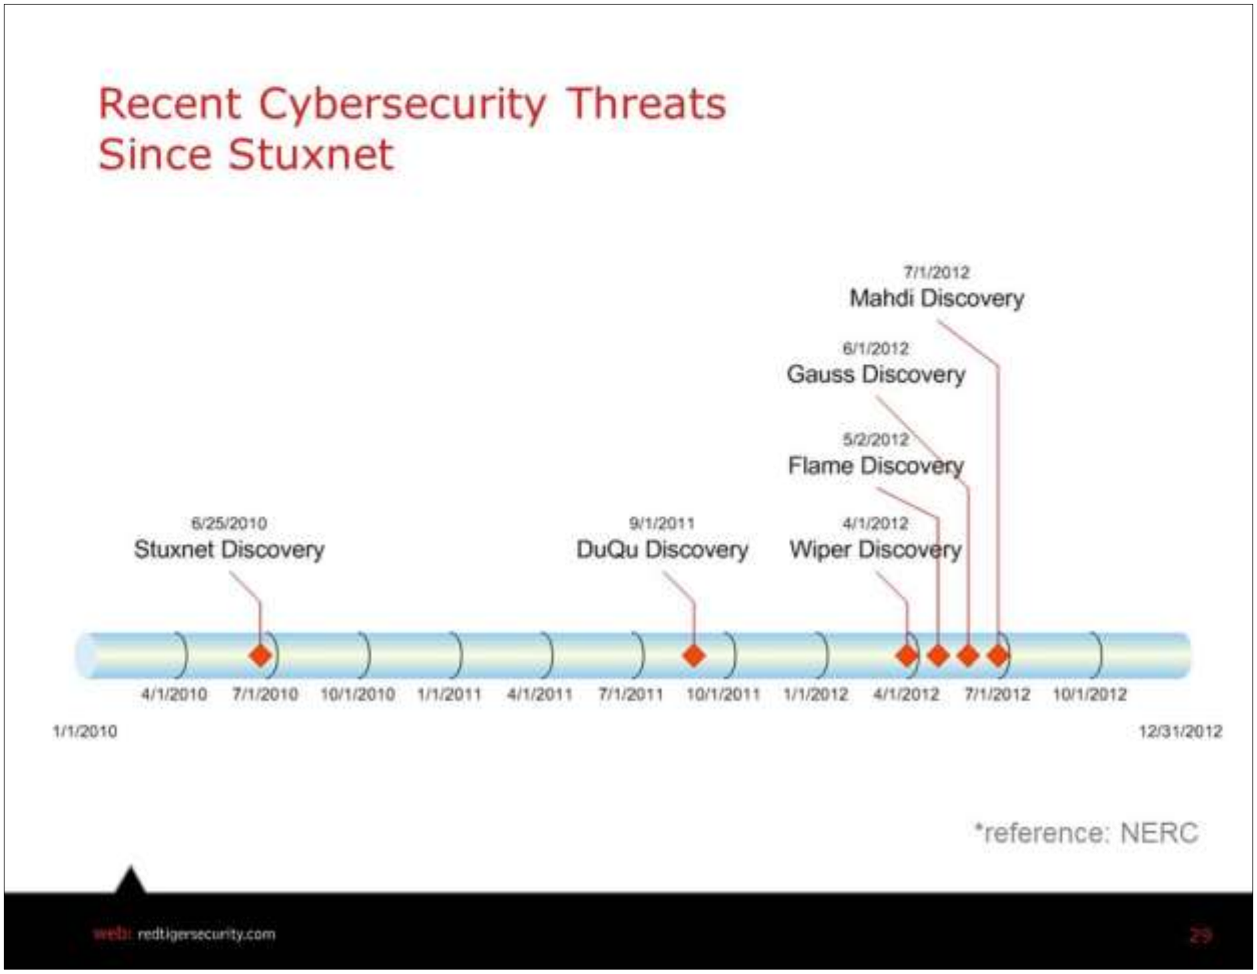

So if you see this on a timeline, we probably all remember what happened with Stuxnet. After Stuxnet, we had Duqu, Flame, Wiper, Gauss, and Mahdi. All of these attacks are ramping up.

Saudi Aramco had 30,000 computers wiped out in one day. It was an inside job. They set it all on a certain time so that it all happened on the same day.

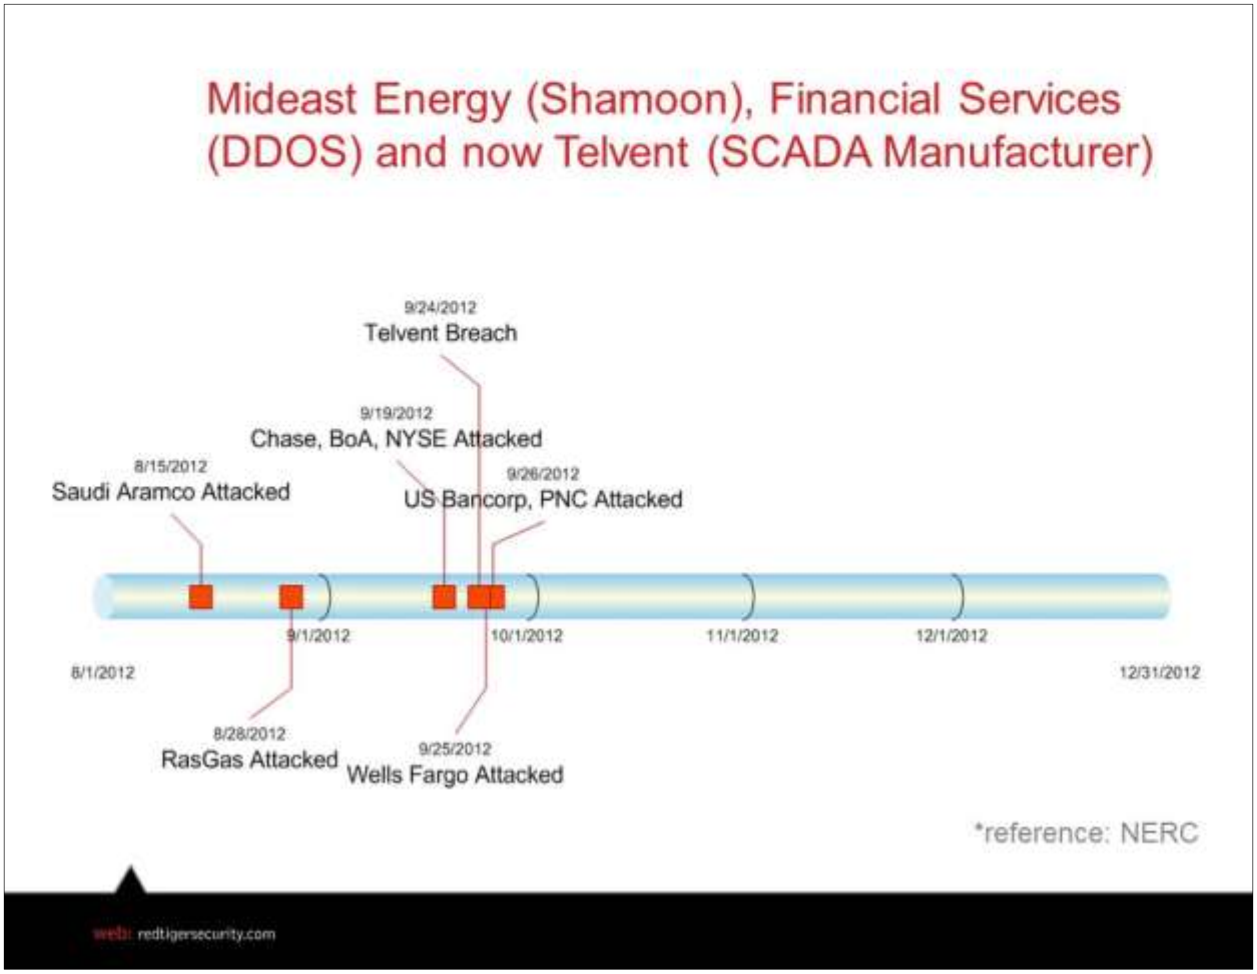

There was recently a rash of attacks in the Middle East, which is why I was up there about a month ago. Saudi Aramco, RasGas… all of these companies are now trying to figure out how to build cyber security into their programs because they weren't ready for it when it happened. Saudi Aramco had 30,000 computers wiped out in one day. An insider did it by writing a piece of code that would format the SD drive and remove a critical piece of the boot sector. Whenever the machine reboots, it can't function anymore. It basically made the computer unable to boot. Then they set it all on a certain time so that it all happened on the same day. So 30,000 computers on the corporate side were not functioning at all.

Fortunately for Aramco, they had put firewalls in front of their control systems, so the only thing impacted was the corporate side.

Now, fortunately for Aramco, they put a bunch of firewalls in front of their control systems from 2007 to 2011. They did a big push and they put in firewalls across all their plants. So the only thing impacted here was the corporate side.

In an ideal world, we'd like to have our control systems, SCADA systems, smart grid stuff sealed off. Nothing can get in and nothing can get out. That's not really practical because a control room environment for a front-end server, front-end processor that's polling all of the field devices, metering infrastructure, then this has to share this meter data with other systems. So there's usually DMZs involved where data is pooled, then sent out to other systems. And, of course, these are all connected to the Internet and therefore everything else.

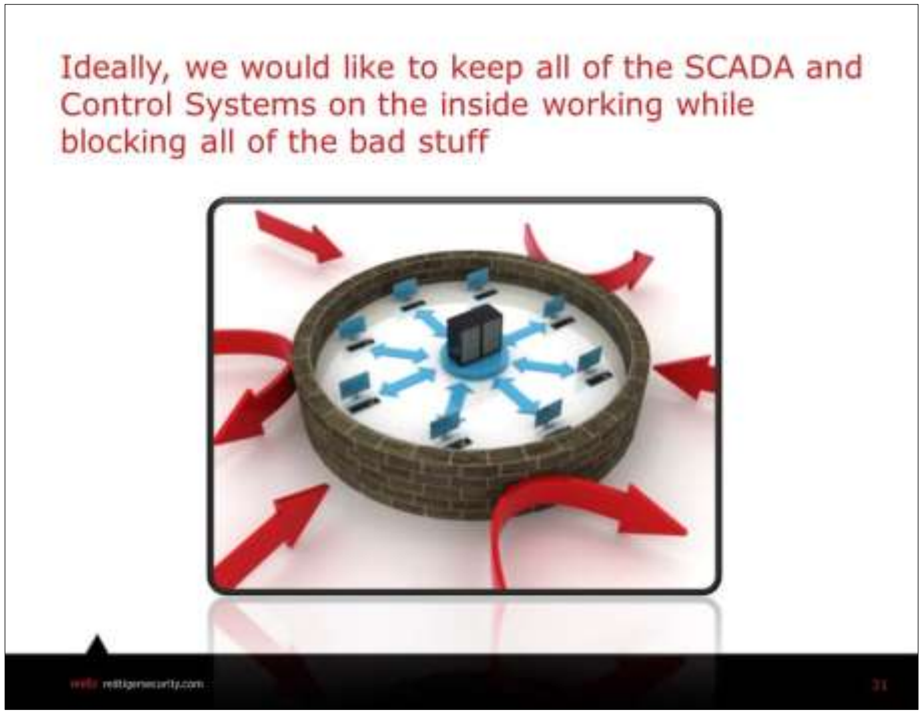

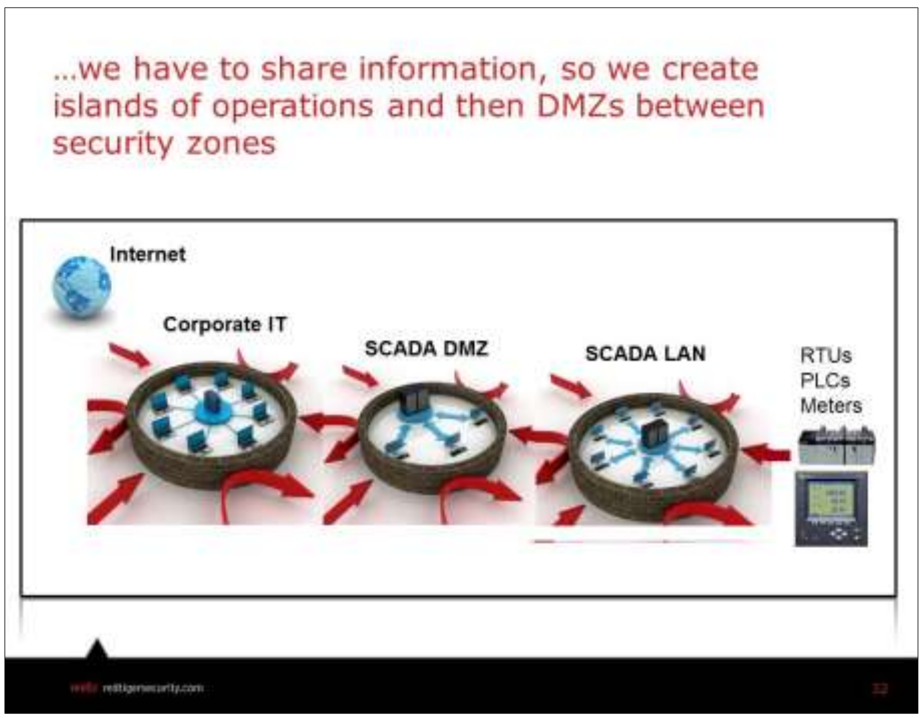

In an ideal world, we'd have our control systems, SCADA systems, smart grid stuff sealed off.

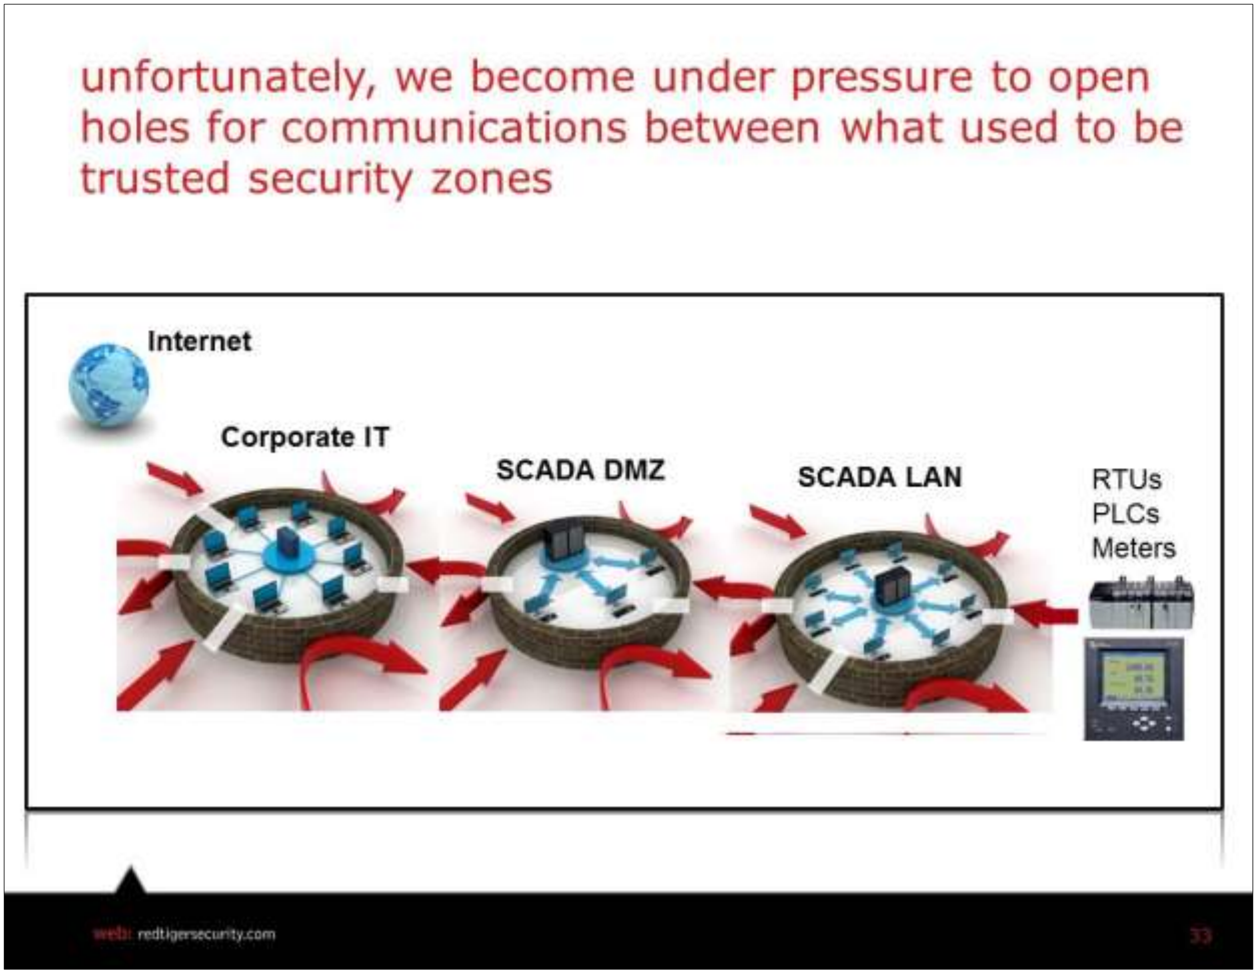

We’re under increasing amounts of pressure to integrate and open up systems, which is contrary to the nature of security.

What you really want to do is *contain* things.

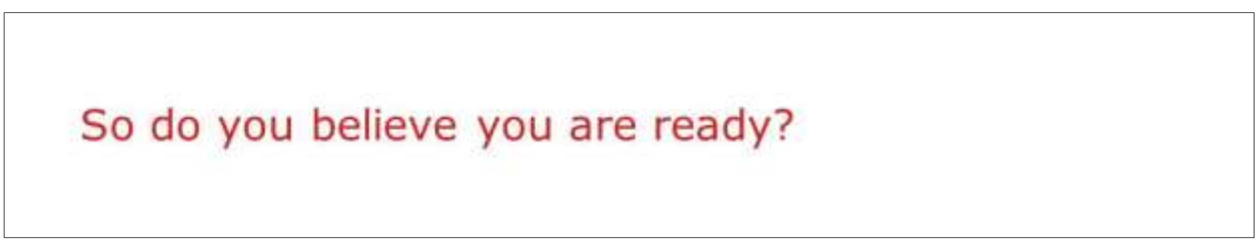

I didn’t get into the whole defense side of things. I can’t solve all of the problems in this presentation, but I really want to crack open your minds and let you see what’s happening on that side of the world is that cyber security is becoming increasingly a problem to everyone.

We’re under increasing pressure to integrate and open up systems, which is contrary to the nature of security. What you really want to do is contain things.

Now there's going to be a problem with the smart grid too. I can see that we're recreating the sins of our fathers. In the SCADA side of the house, we've figured out that the protocols were not encrypted, the end devices were too weak to handle the right level of security features, the front-end servers were not being patched fast enough, we didn't have the right security features, then we had too much connectivity with everything. What I've described to you is exactly what we're doing today with AMI and AMR deployments. We're making the same mistakes.

What I've described to you is exactly what we're doing today with AMI and AMR deployments. We're making the same mistakes.

Hopefully, that opens your mind as to what's going on out there, some of the history around why we're at this point we are in the world today. And that way when you see things in the news about cyber security, you can understand the background.

Section V
Second Panel Discussion

Second Panel Discussion — Jonathan Pollet, Harley Garrett, Jian Sun

QUESTION:

Will the cyber security response cause an increase in the need for more electricity because of more servers or more devices?

JONATHAN POLLET:

Absolutely. Anyone hear about the mountain of servers being developed in an underground bunker in Utah right now?

HARLEY GARRET:

There's an underground bunker that the National Security Agency is building right now that has an enormous server capability for national security reasons. That's going to take a lot of power.

QUESTION:

Is that the reason they're doing that out there… to put a kind of armor against an intrusion?

HARLEY GARRET:

Armor against intrusion? I know a physical attack…

QUESTIONER:

But what about a cyber-attack?

HARLEY GARRET:

Well, the agency, of course, is on top of that. God knows that if anything is as secure as it can be, this will be. You can go out there and look at it in Utah there. I forget the name of the town, but the amount of data that they can store is unbelievable.

QUESTION:

Harley was talking about how fast we can transmit data and power, how fast all these devices connect. And, of course, Jonathan scared the hell out of us how we're going to...be at risk; let's put it that way.

If we're translating data that fast, how is it possible we can protect against hacking? Or do we, are we reacting every time to a hacking? We're not preempting cyber-attacks?

JONATHAN POLLET:

We're not doing a very good job of protecting ourselves.

DR. JIAN SUN:

I might be biased, but I think the local autonomy that I referred to needs to be the first priority. We have to make every device sort of fail-safe. You can't rely on communications to do everything.

I think that's part of the tactical challenge. We want to use communication for a variety of things. But, first of all, we have to start from the local design and make sure they are working properly. If you cut off all of the lines, everything that you have, who knows what will happen.

QUESTION:

That's along the same lines I was going to ask because this is the a huge, huge movement now toward cyber security. Did we see a slide that addresses the Android? the IOS?

I was wondering, we're looking now at how we need more of a firmware- or hardware-based security approach rather than having to rely only on LANTS software because everything is going mobile.

JONATHAN POLLET:

Yes. Purpose-built appliance is definitely where things are at in security now. Five years ago you had probably a server, like an ISA, that basically had a firewall running on Windows. You don't find any vendors doing that now. It is basically purpose-built hardware and chipset just to be a firewall or a switch or whatever device it is. So you won't see as much security running on top of software platforms.

QUESTION:

In the presentation, you showed where the Android was the most vulnerable. What about IOS and Apple?

JONATHAN POLLET:

I'm not going to say no operating systems are all completely secure, I'm just saying that the way that Apple controls the code -- and they review all the code before it goes into their workspace -- prevents a lot of malware to get into the system to start.

The biggest weakness that goes on with IOS is when somebody jail breaks their device, they basically have opened the device for other non-signed, non-Apple-approved software to run on it. Then you basically put yourself on the same market as an Android device.

You still have to take the same types of precautions or run checks. I've had Apple computers since 2006. I run an AV on my computer, I sweep it from time to time; I do all the same types of checks I would if I was using a Microsoft® computer. I patch it as fast as I can.

HARLEY GARRET:

I'm a big open-architecture, open-source software fan. I do believe that with Android --which is open source, by the way — you can get the source code if you want and you can make modifications to the source code. They're not necessarily going to implement it.

If you can get the source code -- and the same thing if you run Linux. We run a lot of Linux machines where I come from. The fact of the matter is that it's sort of a social idea to say that if a whole lot of people are using it, and something goes wrong, you'll know it quicker.

That absolutely has, I think, been proven in the past. First of all, the PCs -- and you can go back in history, Bill Gates really took over the world. Back in 1975, when he decided software was going be proprietary; he was going to charge for it. So from that day on, he sold a quick and dirty operating system, Q-DOS, which he bought into, which is a group of kids in Seattle that had come up with an operating system they called Q-DOS. He bought the rights to that and… guess what? That became DOS.

So IBM bought into that. IBM did a lot of smart things, but that wasn't one of the smartest things they ever did. They sold millions and millions of PCs and they all ran DOS, which is Microsoft's operating system you need with Windows 7 or whatever.

If you go into corporate America anywhere in the world, practically, you can almost be guaranteed that all their employees are using Microsoft platforms.

Microsoft servers are actually a minority worldwide on the Internet. Most of them are Apache -- Excuse me, LANTS, which stands for Lenox, Apache, NIST, SQL, TCP. All which are open source.

There is an incidence of malware on IOS, no question about that. Not as much because of the loyalty of the people who use Macs; much more prevalent on the Microsoft side. Many other things that come out, like Stuxnet, for example, exploited four zero-day vulnerabilities in Microsoft's operating system.

A zero-day vulnerability is a vulnerability that Microsoft didn't know about. That means that nobody knows about it but you. If you exploit it; once you exploit it, everybody else will know… but you can do it first.

They exploited four of those. No Linux systems involved, no iPhones involved, or anything like that.

So that's what is going on. That doesn't mean that Microsoft is bad, absolutely not. But they still have a huge market share of all the platforms that are out there. If you're a perp and you want to attack somebody, guess what? Why bother with Apple? Why bother with Linux? You're going to get caught quicker because there are more people using it. And so much for that.

One more thing Jonathan pointed out; they can't break in unless they figure out a way to break in. That's called an attack vector. In this case, the most beautiful attack vector is phishing. You can blend that with some of your own intelligence. If you know who the CEO of the company is, you can broadcast an email to some of the employees from the CEO and, guess what? They're going to open it. That's their CEO! "And by the way, here's a PDF I want you to look at." Right. Boom. You're dead. That quick.

QUESTION:

Melding security with earlier presentations, we were discussing the general idea of the utility, us giving the utility the right to come in to turn our air conditioning and refrigerator on and off. That doesn't seem like a good thing now.

Wouldn't it be better to just, for the utility to put out a price of electricity, publish that every few minutes, and then you have your own system to say "Okay, at that price, I will turn off…"?

HARLEY GARRET:

And leave the decision to the user.

QUESTIONER:

Right.

JONATHAN POLLET:

Okay. The moment that smart meters became bi-directional, I became extremely concerned… because people like to do load shedding from a click of the mouse. If a person with a mouse can click on one or more homes and disconnect them from the grid, certainly a piece of malware can do that. Maybe not just one or ten; it can do it to a thousand all at the same time.

So I'm certainly really worried about the remote disconnect feature. A load shedding feature is built into a lot of smart meters. That's my biggest concern. In the NERC regulations, nothing is a critical asset unless you can put on or take off 300 megawatts of power at once off the grid. That's going to cause a cascading issue, right?

Well, what if I just did that to enough homes to make that an issue? Now that amassed number of homes is basically a critical asset; I just took that much load off the grid all at once.

I'm thinking like you. That I would prefer the power company either contact me by text message or email something that says "Energy is at a premium. We need to shed some load. If you shed load, here's how you much you can save on your bill."

Then I go to my system and I turn it off myself or I remote into an in-home system and turn off things myself. I'd rather be the one to control it rather than some headless system that could be taken over by somebody else.

DR. JIAN SUN:

I don't have much to add to that; just my own personal experience.

A few days ago, my Dish Network receiver didn't work. A guy came in and adjusted the dish. At the end he said, "It was a minor misalignment. It was easy."

He offered to put a device on my receiver and said, "If you accept that, next time something happens, you don't have to call us. Your network will send us a message and someone will come here automatically."

I said, "I don't want that device. I don't want something traced to my network, sending out info -- I don't want it."

But anyway, that's just my personal take on it.

Section VI, Session III
Improving Grid Stability

Grid Stability

DUSTY BECKER:

Our first speaker is Dr. Gregory Smedley. Greg Smedley is a Caltech Ph.D., Founder and CEO of One-Cycle Control, Inc. (OCC).

Dr. Smedley is a renowned researcher, educator, and industry leader with a diverse career that spans power electronics, medical devices, research and education at Caltech, and space exploration at JPL.

Dr. Smedley is a world-class technologist with a passion for paradigm-shifting technology across a broad spectrum of fields and possesses a vision for potential and opportunity. He has won awards from DOE, DOD, CEC, CIEE, NASA, and NIH. He is a recipient of the Teaching Excellence Award at Caltech in 2000 and US Department of Army SBIR Achievement Award in the Pentagon in 2010. He has delivered invited talks throughout the world.

At OCC, Dr. Smedley is devoted to breakthrough power conversion technology for renewable energy, energy storage, grid stabilization, smart grid control, hybrid power, and mobile power. The OCC team has developed a family of modular, rackable, scalable power-converter products featuring leading power density and dynamic performance to address the kW to MW market.

Dr. Smedley is focused on tackling global-scale issues in energy, electric grids, transportation, and low-cost access to space.

GREGORY SMEDLEY:

Dusty, thank you for the introduction. It's a pleasure to have the opportunity to join in this activity today. I'm learning a lot and also getting scared. I think it's touching a lot of different emotions for me. So thank you very much for the opportunity to come speak. Also thank you for your patience as I worked through getting this presentation to you.

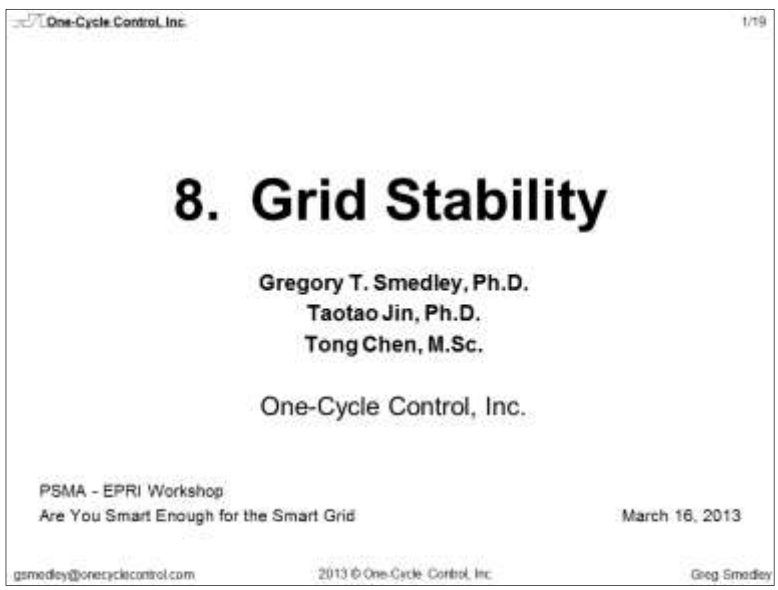

My talk today is on grid stability. At One Cycle Control, we're focused on all kinds of power converters, but one of our big passions is stabilizing the grid. I'm very pleased to have an opportunity to talk about this.

I thought it would be good to start with a definition. I looked up “grid stability” on the electrical engineering portal and I found out that it was the ability of a system, given an initial condition, to remain in a normal state of equilibrium after being subjected to a disturbance.

I thought, “That sounds a little bit wishy-washy to me because it sounds like the system is going to get kicked around and it’s going to oscillate a bunch, then maybe it’ll drop some load or generation, and it will eventually come back to a state of equilibrium.” I was looking for something a little bit stronger.

So I found one that says, “The ability of an electrical circuit to cope with changes in the operational conditions.” I thought that sounded a little bit better, but I still am looking for something stronger because I really want some stability.

I found in a medical dictionary: “The condition of being stable or resistant to change.”

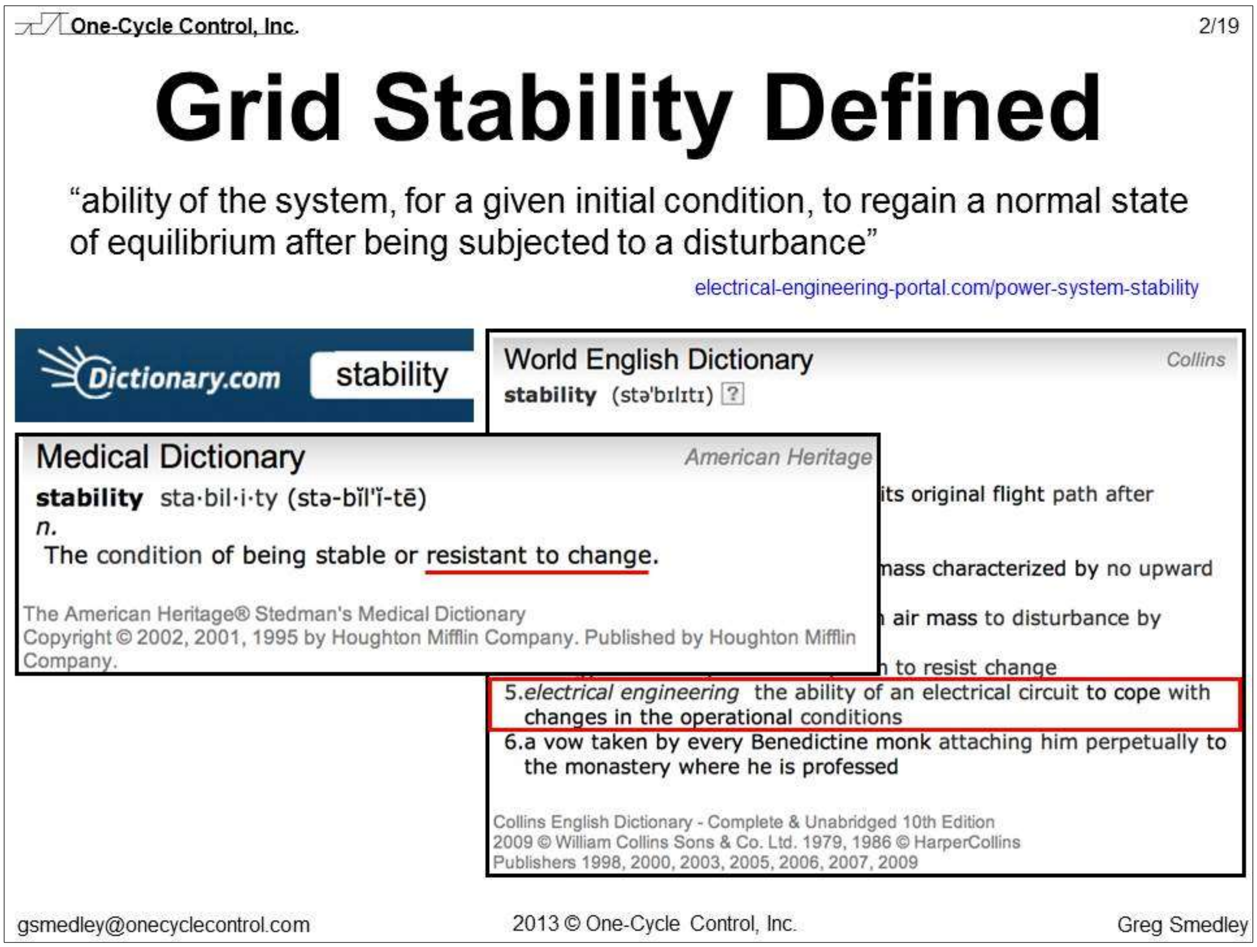

The grid is a very large system — man’s biggest machine. When that machine gets perturbed, it should stay rock stable like nothing happened to it.

I look at this grid as a very large system. As we heard earlier today, it's man's biggest machine. When that machine gets perturbed by anything; basically, most of it should stay rock stable like nothing happened to it.

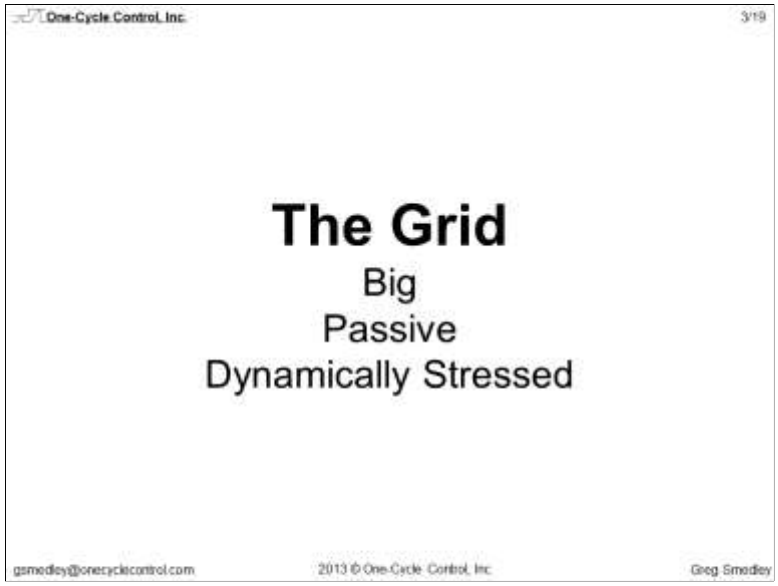

The grid, today, is a very big, very passive, and dynamically stressed system. It's mostly passive. There are some actuators in the system, and I'll touch on those in a bit, but in general, it's relatively passive.

We have dynamic sources and loads; we've got renewable energy sources that produce electrical power at varying rates (depends on which cloud goes by when, when the wind blows, how it gusts, etc.); and we've got loads that are dynamically changing at the whim of the users who randomly turn switches on and off, with no way to keep any control over their personal behavior. This is a problem because these types of dynamic sources and loads lead to a dynamic voltage and frequency on the grid.

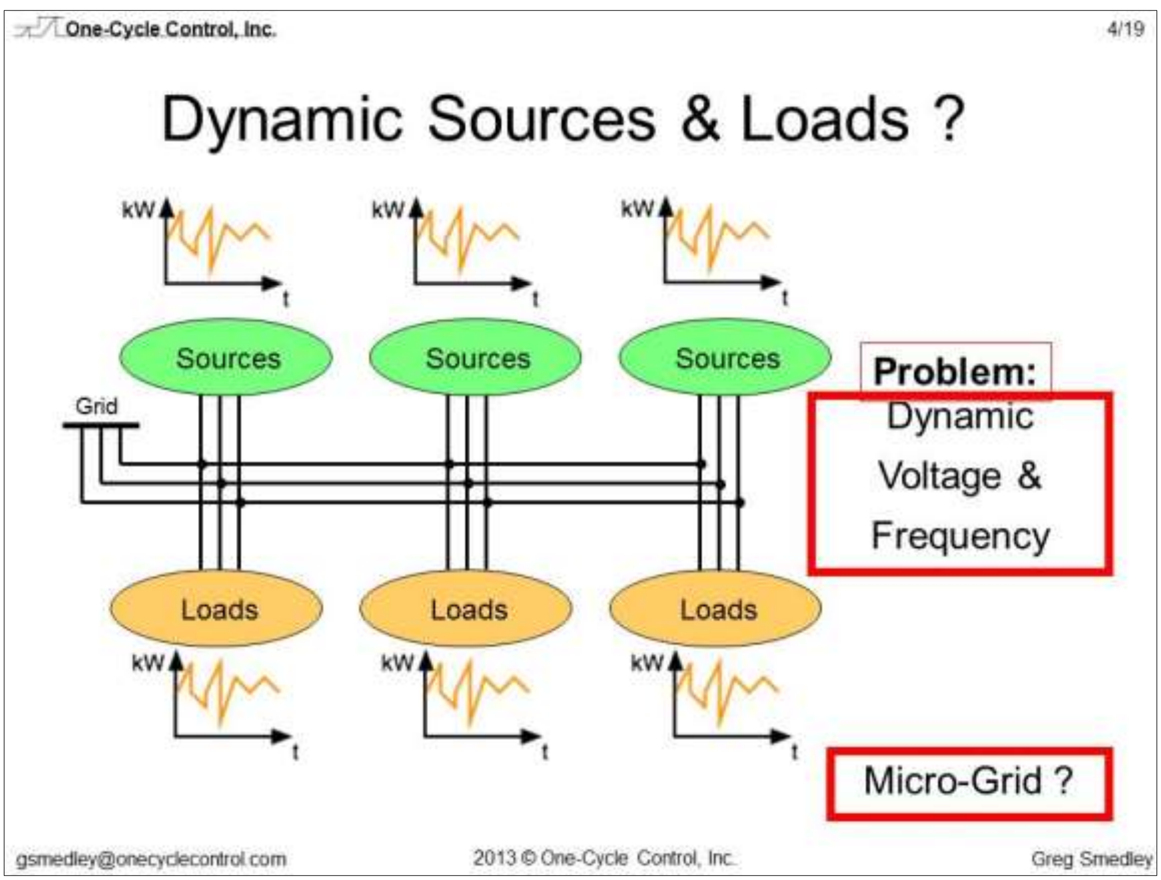

Then, of course, what about the microgrid? Because everyone talks about microgrids. So we move down to the size of a microgrid and we've now dramatically reduced the inertia of the system. It becomes even more dynamic and the speed of response to exert control, therefore, becomes even more critical.

Furthermore, if we also have dynamic sources and loads, you can imagine that the dynamics of that grid, that microgrid, is going to be even worse! Frequency and voltage would be all over the place.

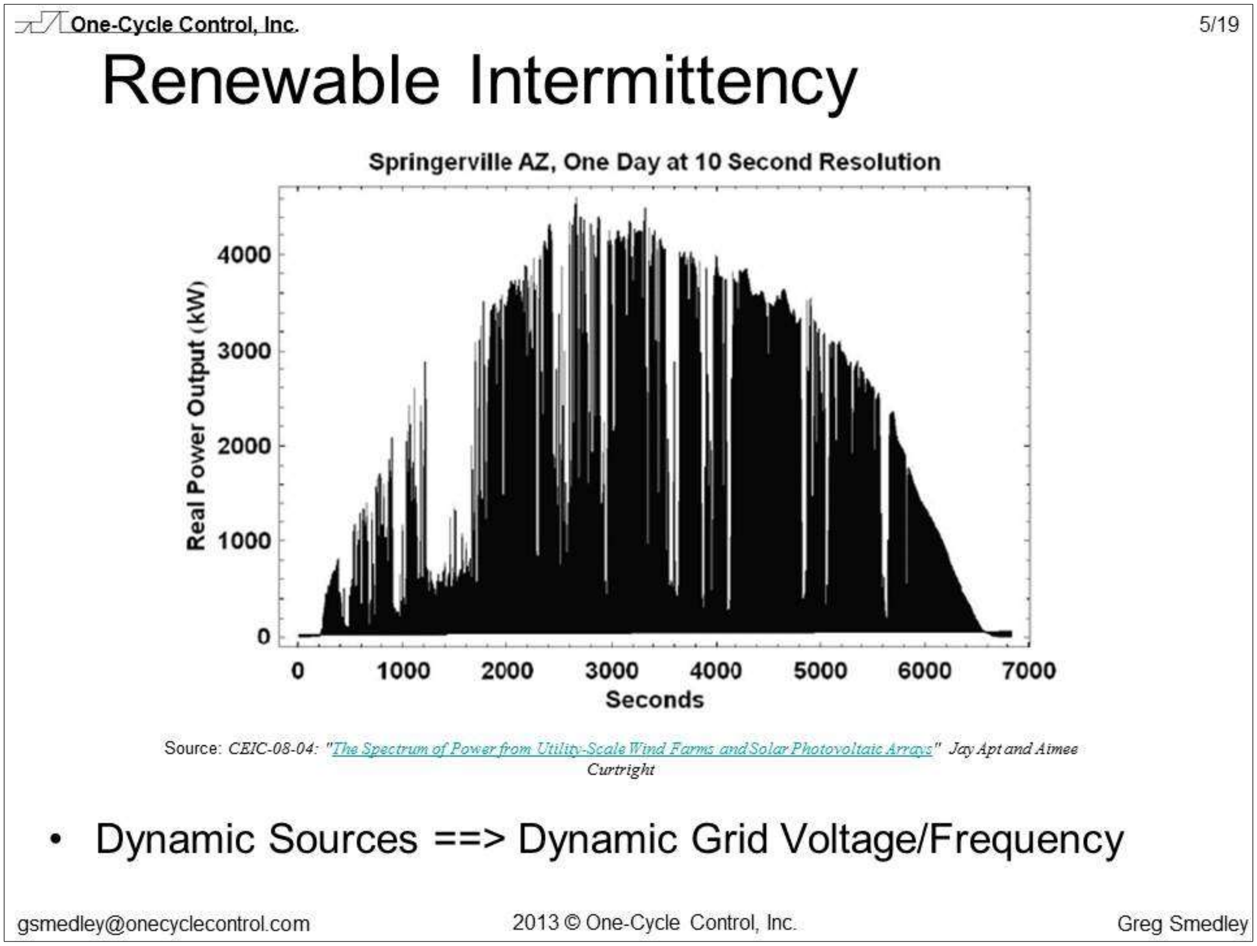

This is an example of intermittency. This is on 10-second intervals. You can see the output of this power plant goes from 400 megawatts to a few hundred kilowatts in a matter of seconds.

We were talking about ramp rates earlier; this is just one specific power generation facility with PV power. I was at a high-penetration PV forum that was put on by the DOE down in San Diego a couple of weeks back. Tom Bialek of SDG&E put up some slides where he showed the intermittency on 1-second intervals, 10-second intervals, and he kept clicking through. He said "When I get to 15 minutes, I don't even have a problem anymore." That's normally where people are measuring power.

"Dynamics" is a big issue and it's happening at rapid time scales. "At 15 minutes, I don't even have a problem anymore!" and that's where people are measuring power.

So dynamics is a big issue and it's happening at rapid time scales. The problem? Cascading outages.

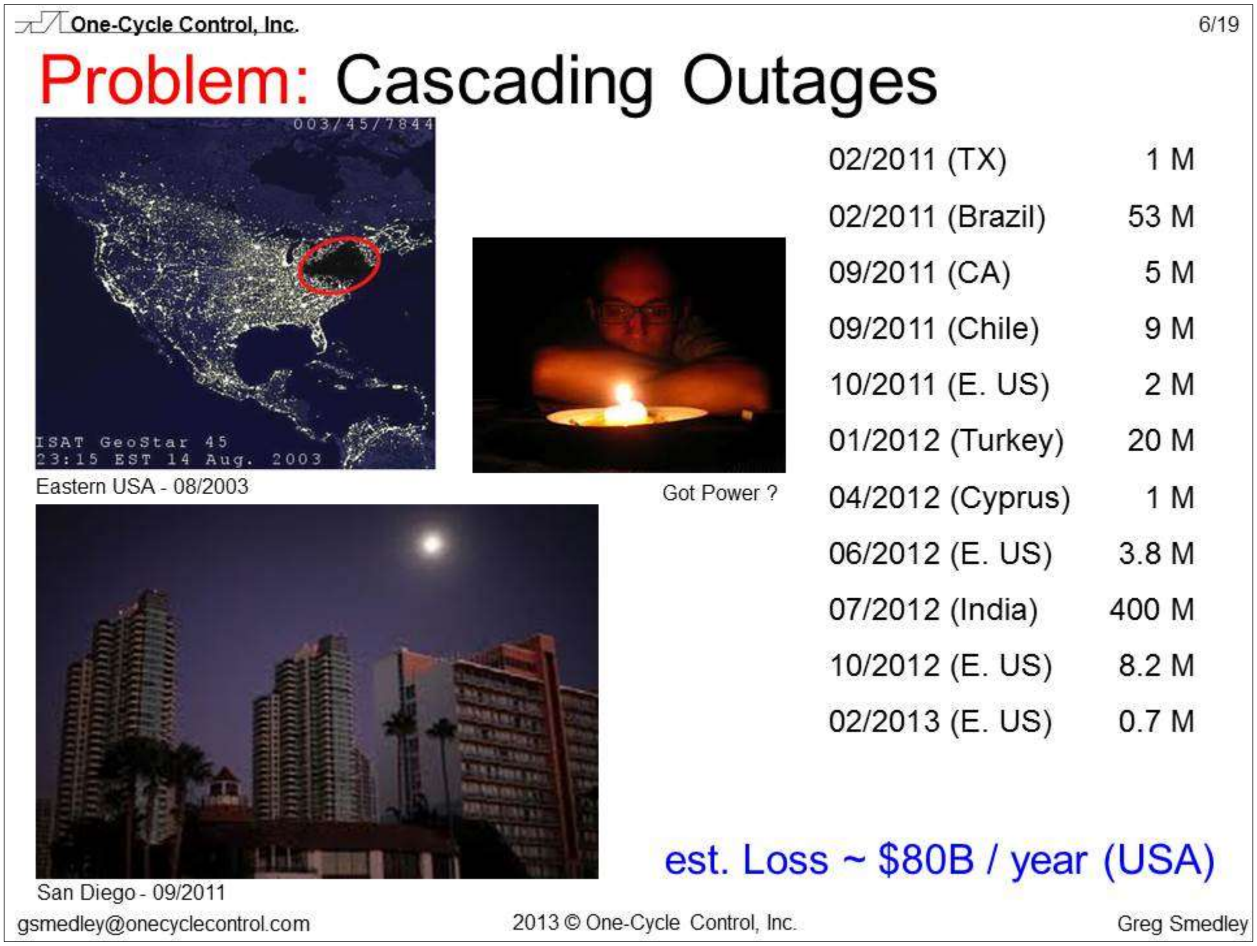

I didn't go all the way back to the one we heard about in early 1996. I just wanted to put a few on the slide, so I go back about a year or two. Texas, Brazil, and California, back in September 2011; down in San Diego; there's a picture of that one.

This is the 2003 [Eastern USA] blackout we heard about earlier… a big, black hole in the middle of a highly, densely populated area.

Here's someone sitting there staring at a candle, wondering when the power is going to come back on. Is it going to be 30 minutes? or 7 days?

Estimated loss, of course, is about $80 billion a year in the U.S. Let's say we could cut down maybe 20% – 25% of those. That would be about $20 billion a year. What could we do with $20 billion a year that could help us improve the resilience and reliability of our system? And get an ROI in about four years on some assets that could really make the difference for the United States to bring it into the twenty-first century and really keep the grid under control?

Renewable energy limits: The CPUC says "33% by 2020" and the utility says "maybe 15 to 20% on a feeder." That's a problem.

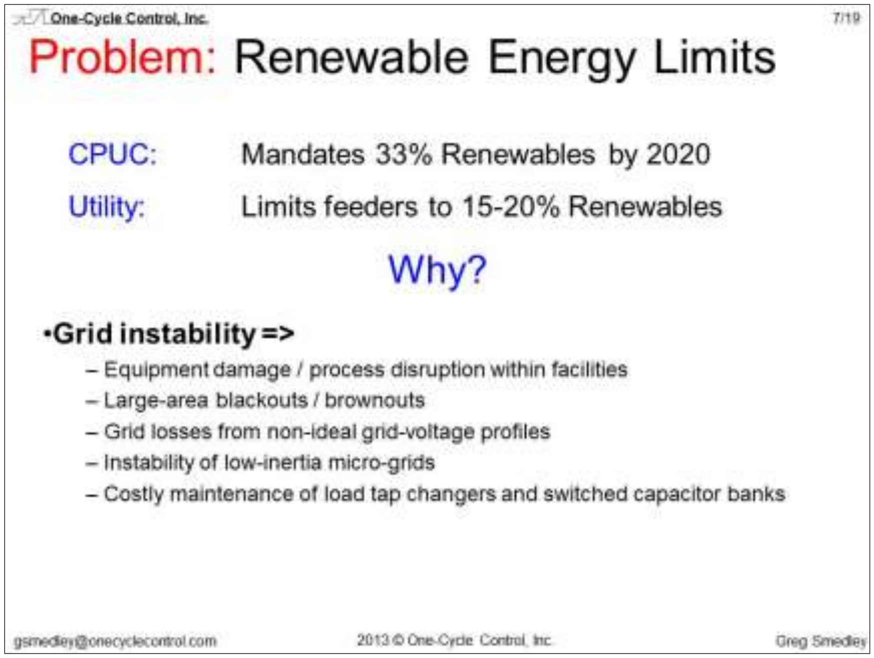

Another problem: Renewable energy limits. Everyone wants renewable energy. The CPUC says 33% by 2020 and the utility says maybe 15 to 20 percent on a feeder. That's a problem. If we want to get to 33% renewables, we need to go out to the desert, we need to pave it with solar panels and tell all those tortoises to move out of the way as we punch the holes in the ground.

Then we need to build new transmission lines or upgrade the transmission lines so we can bring it back to the load centers. Add to that, power plants with once-through cooling that are coastal and located in load pockets are going to be taken offline.

As you go out to the tendrils of the network, you see all the dynamics of that system.

We've got an issue here. Why is it that we're limiting the 15 to 20% on the feeders? As you go out into the network, and you go out to the tendrils of the network, you start to see all the dynamics of that system. Then when you wrap it all back up into the transmission system and you see this nice curve that goes on in the daytime.

Out there at the tendrils of the network, there are a lot more dynamics. We've got issues. We've got equipment damage because of grid instability. We can't keep the voltage in control. We've got large-area black outs and brown-outs happening because we've got cascading ripples going through the network.

There are grid losses from non-ideal voltage profiles. We heard about conservation voltage reduction earlier today. If we can control that, maybe we can dial it in and optimize the grid performance so that we can reduce losses.

As we go to microgrids, we have low-inertia systems.

And, as we go to microgrids, we have low-inertia systems. We heard earlier about maintenance of load tap changers and tap banks where switches are going on and off at high rates of change because of dynamics imposed on the system by the PV.

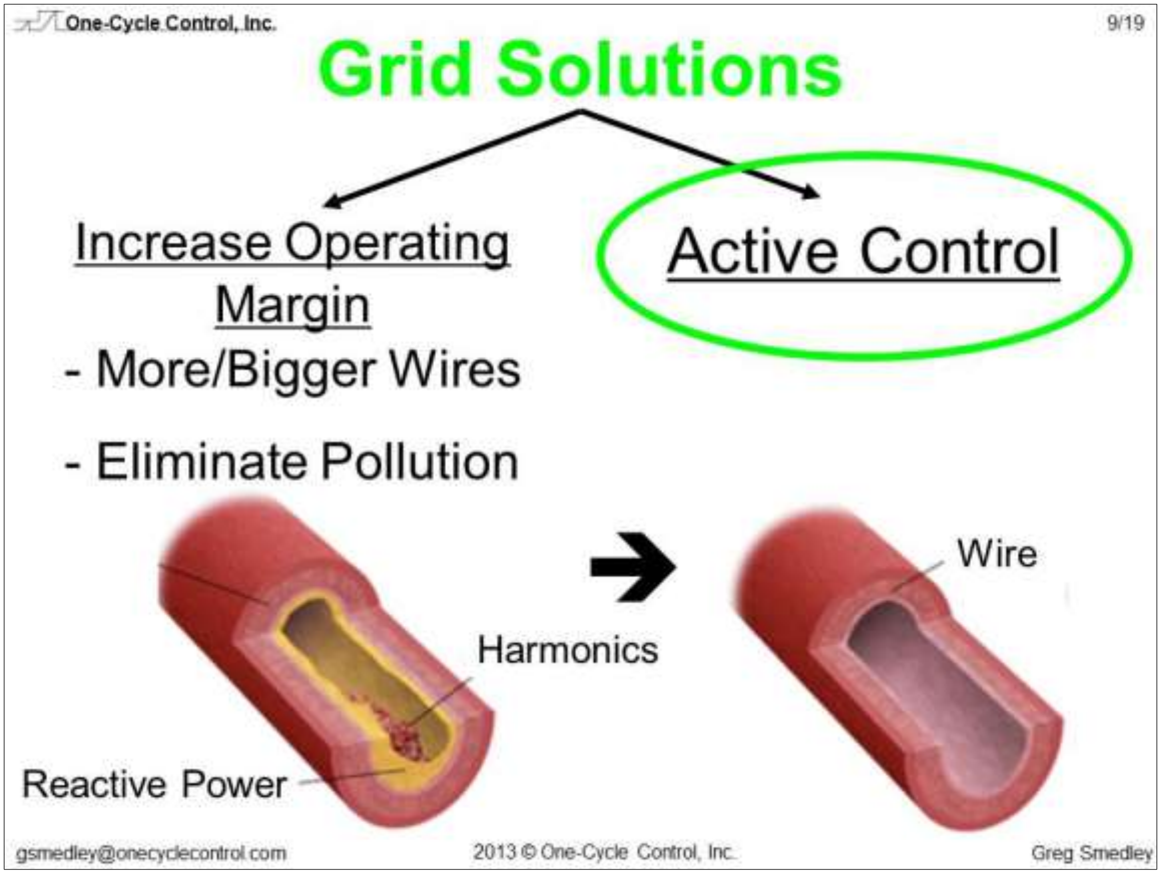

We'd like some solutions, grid solutions. Basically, the solution to the grid problems break down in a couple of different areas. One is we can take the path on the left here, which is to increase the operating margin. If we increase operating margin, the people at CAL-ISO are happier. Everybody takes a coffee break and can enjoy looking at their screens. We can do this by putting in more wires, bigger wires, transformers, etc.

We heard earlier about maintenance of load tap changers and tap banks where you have switches operating at high rates because of the dynamics of PV.

There's actually a power electronics solution to this problem as well. We can eliminate the pollution on the network so we can free up operating margin. Then we'll have a little more room to spare for these types of dynamics imposed on the system when things fail.

Here is a medical analogy to grid pollution: You've got a blood vessel here and it has some plaque inside it. If you drink lots of red wine, maybe you can eliminate this plaque. There are also some blood clots.

We can eliminate the pollution on the network so we can free up operating margin. We have reactive power and harmonics. That's the pollution in our system.

If we look at this from an electrical perspective, we have reactive power and harmonics. That's the pollution in our [grid] system. We'd like to get it out of there. If we can get it out of there, we basically free up the flow so we can get more useable power through the network, rather than all that reactive current we heard about earlier this morning. That artery is actually representing a wire on the grid.

The other solution, of course, is active control. These solutions are not mutually exclusive. These are really a bag or a tool kit, as it were. We need to apply these tools to get the system under control.

We can do active control to make better use of the [grid] assets that are already deployed. We'll be able to operate them over a wider range of performance.

We can increase operating margin and maybe we're going to have to put in some new transmission lines as we grow out the system. We could limit some pollution and free up some bandwidth there. We could do active control to make better use of the assets already deployed and we'll be able to operate them over a wider range of performance and serve more customers with the same infrastructure. This is what I'm going to focus on for the rest of my talk.

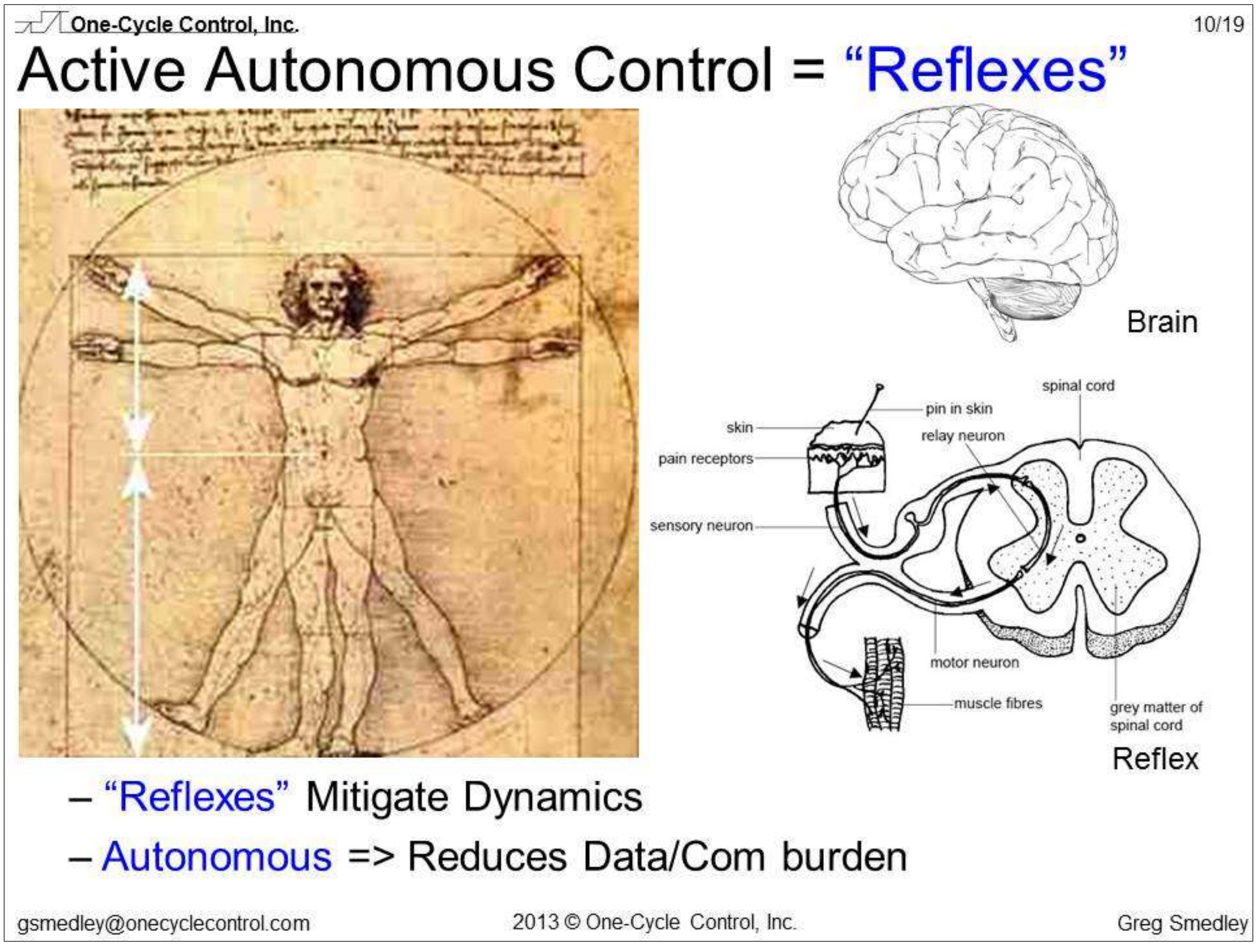

If we're going to do active control of the grid, it's a very complex system, and we heard from our friend at RPI, Dr. Sun, talking about autonomous control. Autonomous control is something that really becomes very essential in any kind of a large, complex system. We need to have nodes that are distributed throughout the network and these nodes need to know what to do. You give them a command [set point] and they immediately try to control to that particular control point.

The human body has fast-acting systems that help us to free up bandwidth in our brains so that we can do other higher-level functions… our reflexes.

Just as the human body is a very complex system, I think of the grid as a very complex system. They're comparable. When we think of the smart grid, we usually think of a very large server network somewhere and that's going to be the brain of the whole system. It's going to collect all of this information, it's going to know what to do, and it's going to tell everyone what to do.

The human body, actually, is not just a brain and a body. There are some very critical, fast-acting systems in our body that help us to free up bandwidth in our brains so that we can do other higher-level functions. Those are our reflexes.

If we reach out and touch something that is too hot, we immediately pull back. The command sent to that muscle didn't go to the brain: it just went to the spinal cord and back to the muscle.

Reflexes can mitigate dynamics down to the tendrils of the network. They need to be autonomous so that we can reduce data and communication burden.

So these types of reflexes, if we can build them into our power distribution network and our power grid, can maintain a much stronger stability so that the people at CAL-ISO, the big brains, can keep the rest of the system under control.

Reflexes can mitigate dynamics down in the tendrils of the network. They need to be autonomous so that we can reduce data and communication burden. Ideally, they should know what to do if those data-com links are cut off, so the system will continue to operate in a relatively stable fashion.

Let's take a look at the kind of control equipment we have today. Today, as I said at the beginning, we have a passive grid. But, in fact, there are some active elements in the grid already. They've been there right from the very beginning.

Controlling the grid is all about controlling voltages. We can put in a capacitor to boost it up or we can put in an inductor to drag it down.

We have capacitor banks and we have inductors we can put into the system. Really, controlling the grid is all about controlling voltages. We can put in a capacitor to boost it up or we can put in an inductor to drag it down. We can, of course, add load or shed load.

This is just an example. This is a static-VAR compensator. It takes up a very large space at a major substation. Basically, what you've got is the thyristors we heard about earlier, attached to these various passive elements. They get switched in and out of the network.

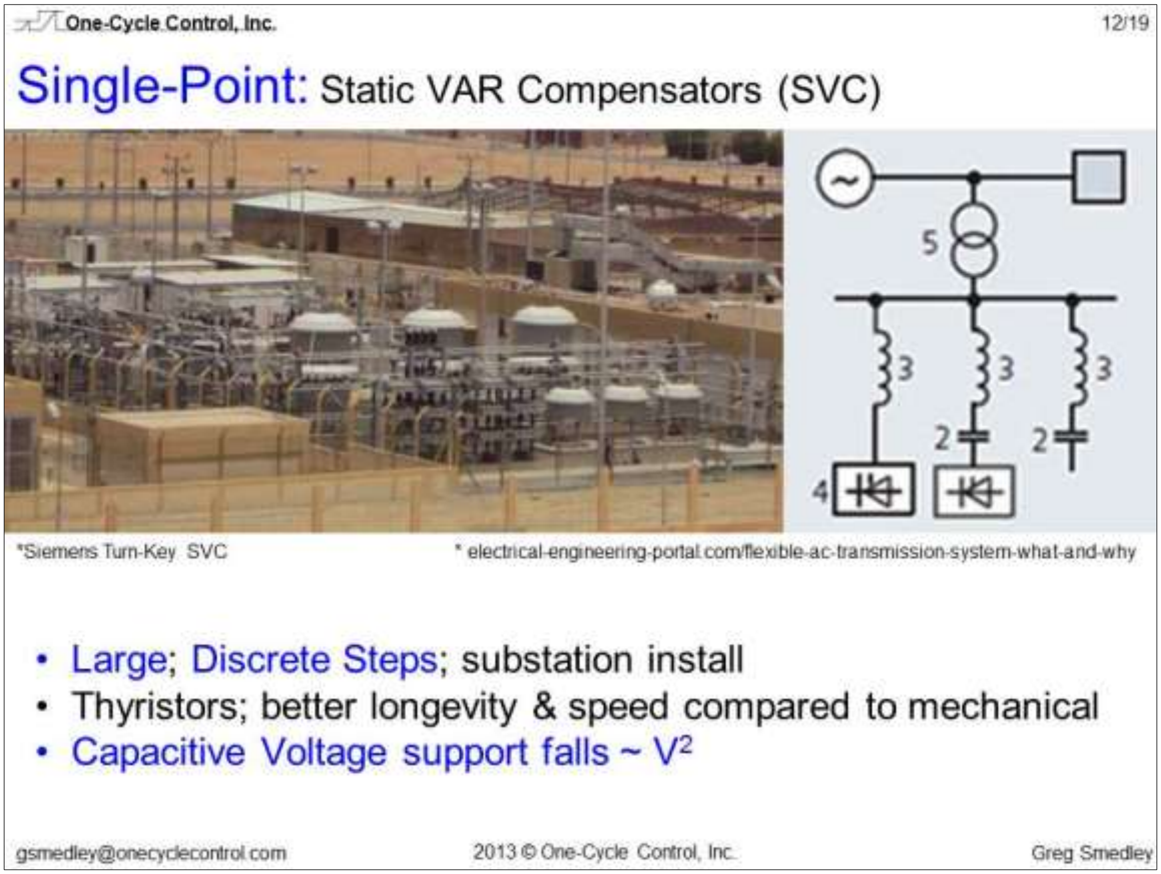

Well, that's very nice. That enables us to control the voltage up and down. But, as we all know, whenever you switch a capacitor into a circuit or switch an inductor into a circuit, you get a dynamic response. Basically, every time we fire off one of these switches, we're giving the system a kick and we're going to see what happens. Hopefully that kick just ripples and dies out passively and we don't have a problem.

The delay in operating big equipment can actually be a problem on the network.

If we've got some other dynamics in the system that are causing these to be exercised more often, the delay in operating this type of big equipment can actually be a problem on the network.

There are large, discrete steps taken that are typical in these types of systems. They're usually installed at substations. Thyristors, of course, have much better speed and longevity than typical mechanical switches. In the past, this would be done with vacuum-type circuit breakers and things like that with much slower response.

One of the key things to remember is that the capacitive support you get out of the capacitor, which is something that you rely on as the grid is crashing; that capacitive support goes down as a square of the voltage. That means it's falling off very quickly. You better use it very early in the process because, if you use it later on, it's not going to provide you with as much support. Or you have to put a lot bigger ones there to get the job done.

Power electronics enable us to create something that looks like a capacitor or an inductor and control it smoothly; the static, synchronous compensator or "statcom."

Rather than switch those capacitors in and out with these thyristors; another way is to use a power converter. The power electronics enable us to smoothly create something that looks like a capacitor or something that looks like an inductor and move back and forth through these different types of behavior. We can automatically control this without having those big step changes imposed by those mechanical or thyristor-type switches and passive elements.

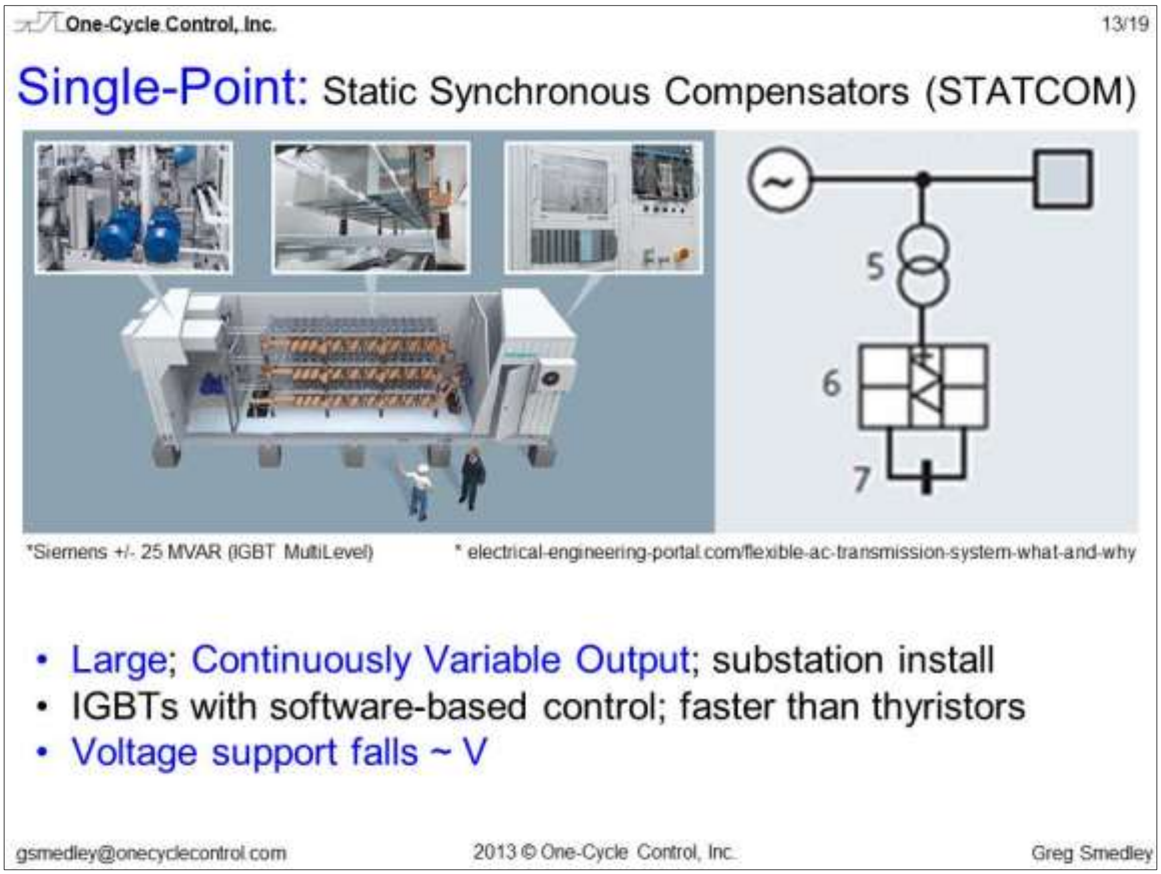

This is a static, synchronous compensator or "statcom" and this is typically a large system. It has continuously variable output, which is an advantage, and faster response. It's installed at the substation that uses IGBTs. It's faster than thyristors, typically, and switching at higher frequencies.

The other nice thing about the power electronics is how it falls off with V. As the voltage is dropping, you get a linear reduction in the amount of capacitive or inductive support you can provide to the network.

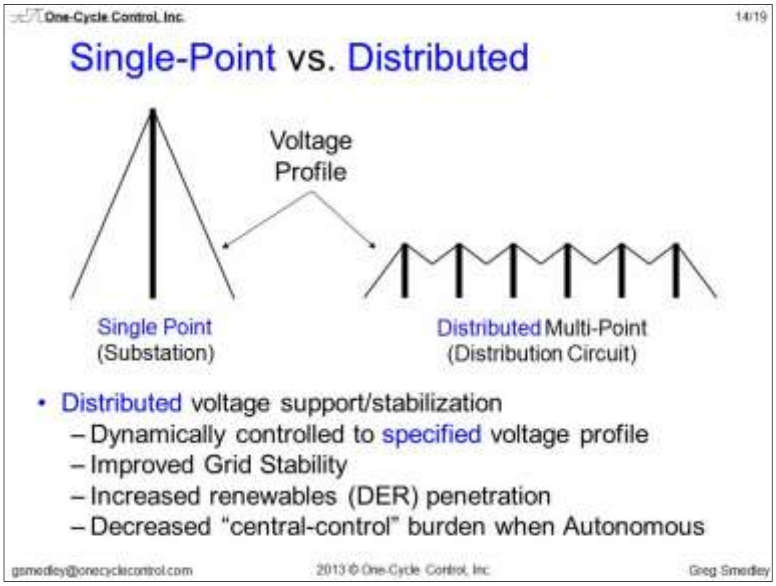

Let's take a look at this. I showed two single-point solutions. Those are solutions or tools in the toolbox we can put at substations to provide single-point support on the network. Think of that as a big tent pole. You can put this tent pole in the substation. That provides you a single point of support. You can control the voltage at that point. The rest of that is the voltage profile. As it goes out across the network, that voltage is going to fall as it falls, based on the load being applied and, of course, there's the detail of load tap changers that give you some additional refined control.

I would encourage distributed, multi-point support to create the voltage profile you want.

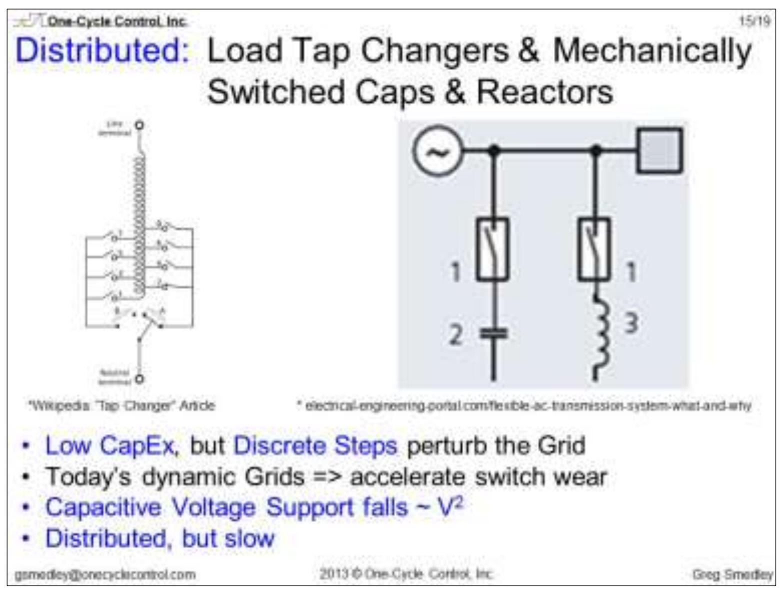

What I would encourage is distributed, multi-point support. Now you have smaller tent poles, but you put in a lot of more of them. Now, you can create the voltage profile that you want. If you imagine this in a two-dimensional space across a geographic region, each one of these tent poles can be adjusted in height so you can exactly specify the voltage profile for that network and change it at will.

With distributed voltage support with stabilization, we can basically control to a specified voltage profile. That voltage profile can be sent out and be controlled too. We can improve the grid stability because now we have multiple points that can dynamically respond to events. This enables us to increase renewable penetration so we can get up to those 33% levels we're targeting. We can decrease the central control burden because these are autonomous elements. The amount of data transfer is basically a simple set point control.

We can decrease the central control burden because these are all autonomous elements.

The good news is that we actually have distributed multi-point control in the network today in the form of load tap changers that enable us to set the voltage. We heard about that earlier today when we saw the plots of the voltage profile and how we can move that up and down with the load tap changers at the substation, then distribute along the network.

We may have some load capacitors and those load capacitors have mechanical switches. A lot of systems on the distribution network actually use mechanical switches, but others use thyristor-based systems.

We can do all of this, but the thing to realize is that this is happening on relatively slow time scales. When you add in the dynamics of PV and wind, these start to get over exercised, so they wear out too quickly and need to be replaced.

So what does the user do? The user goes out there and says "Oh. Every time a cloud goes by, the voltage drops, then this capacitor switches in. The cloud goes away, it jumps back up again, and it switches back out. Clouds seem to go by every five seconds or so. Why don't I just turn up the delay on this machine?" Let's turn it up and make it delay, maybe open up the range so it only adjusts if the voltage goes out to these other limits.

They basically are slowing down the control response of the system to not wear it out.

So they basically are slowing down the physical response, basically the control response, of the system. They're dialing it back so as to not wear it out.

We heard about power electronics before. We can do this with power electronics on the grid. These things [legacy devices] have advantages because they have relatively low capital cost, but they have discrete steps. These discrete steps are perturbing the grid down in the distribution network. Today's dynamic grids accelerate the wear.

We have, again, a capacitive voltage support that falls like V^2. We do have a distributed solution deployed, but it's slow… and they made it slower because they're worried about it wearing out.

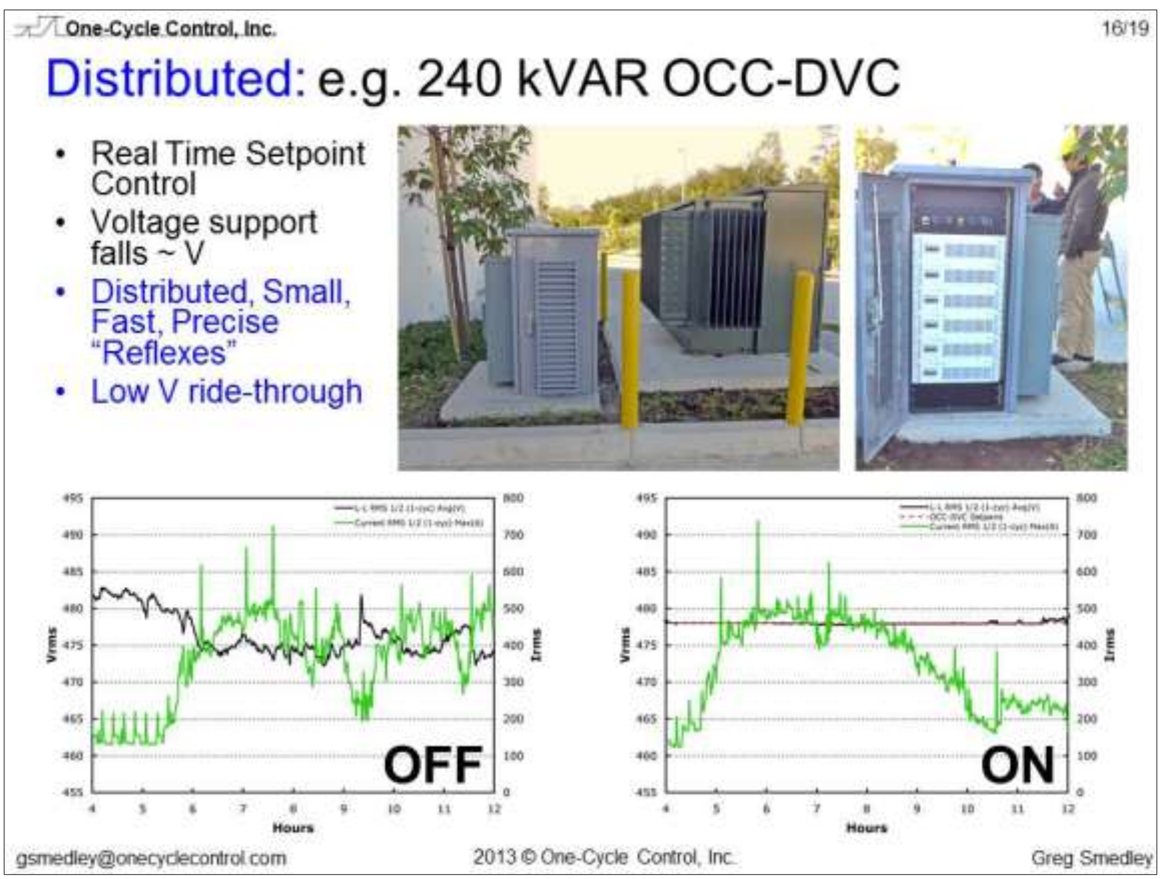

This is a great opportunity to use power electronics. This is an example of a distributed solution that we installed in a U.S. utility earlier this year. In this particular case, the system is located near a distribution transformer close to the customer site, and enables us to control the grid voltage.

This is a great opportunity to use power electronics. The grid voltage changes dynamically throughout the day. Once we install the dynamic load compensator, we can set it where we want.

In this case, on the left, we've got dynamic load at this site. It's a manufacturing site and the grid voltage changes dynamically throughout the day. When we installed the dynamic VAR compensator and set it to 478 V, you can see it's completely flat. This means we can take this voltage and we can set it wherever we want.

So in the morning, we want it at 478 V. In the afternoon, we want it at 482 V. On the weekend, we want to set it at 485 V because there's no load there. Whatever you want it set to, you can.

You actually do control the voltage. You specify it.

Then put many of these across the system and you can set all of them. Each of them will control the voltage. It's not just a matter of setting it to some approximate voltage and letting the voltage go wherever it wants based on how the load changes or the supply changes. This is a case where you actually do control the voltage. You specify it.

You have real-time set point control. Voltage support falls with V again because it's a power electronics solution. Now we have distributed, small-scale, very fast, and very precise reflexes.

Remember we talked about reflexes earlier? These are reflexes that are built right down inside the network and they respond to keep the grid under control. As the grid rides down, of course, these can just ride down with it and continue to support as much as possible across the whole dynamic event.

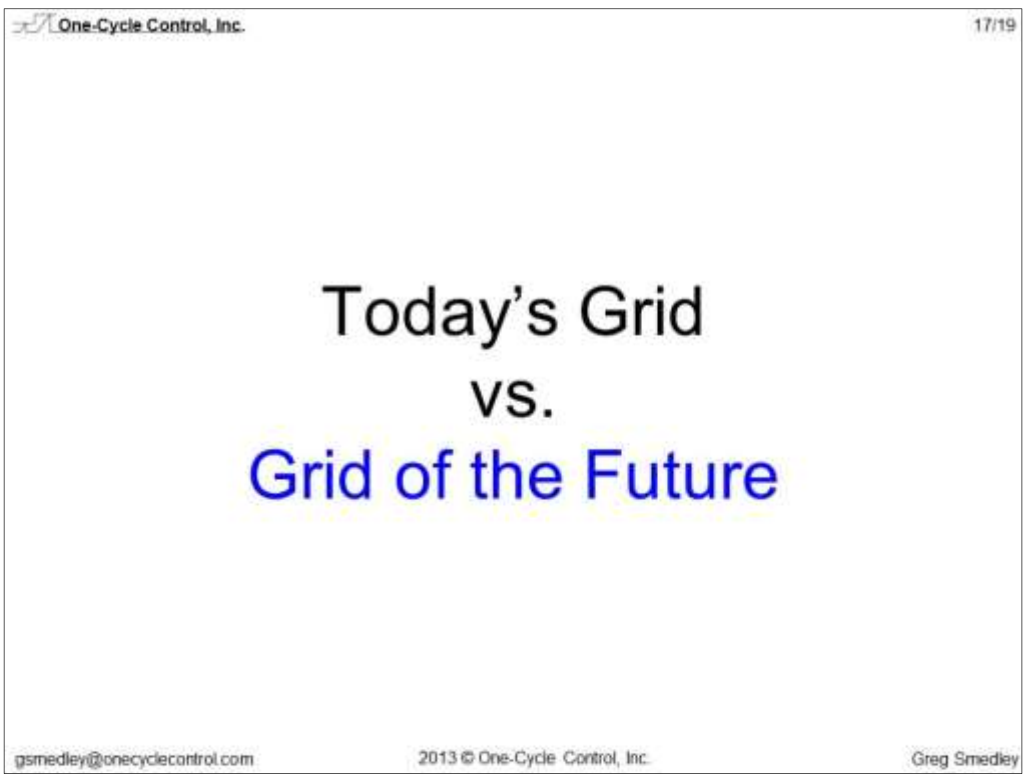

So we have today's grid and we have the potential for the grid of the future. Let's just go back in time to September 2011, San Diego. I don't know if anyone in the room is from San Diego, but I want to tell you a tale of two cities.

On the left, we have the city with a passive grid. It's pretty black. How did this happen? There was a guy in a substation in Yuma, Arizona, and he pulled a switch. It's just a switch. But it happened to be a switch that was connected to a very large piece of load. When he did that, the whole grid lost a big piece of load.

When a big piece of load falls off, the whole system's voltage jumps up. Assets on the grid all have protective relays on them. They disconnect.

When that big piece of load fell off -- I actually went and downloaded the CA power-demand profile from Cal-ISO and you could see this — the grid load was going up and up and suddenly there was just a step. It was a big drop on the profile for that day. For California, it was visible at the state consumption level.

So that was a big load drop. As soon as that load dropped off the system, the whole system jumps up. The voltage jumps up and travels like a tidal wave across the grid. When that voltage jumps up, certain assets on the grid (they all have protective relays on them) say, "That voltage is too high. I can't deliver power anymore." or "That voltage is too high. I need to disconnect this load." There's going to be a response in the system.

Imagine the grid voltage goes up. San Onofre says, “That voltage is too high. Obviously, there’s too much power on the network.” It shuts down. Maybe it was too much power to let off the system because maybe five or ten other gas plants shut off at the same time. Now the whole grid crashes down. As it’s crashing down, of course; if anything stays connected at the point, it’s pretty much a surprise. The grid is dead.

Now let’s look at the grid of the future. The city with the active grid over on the right remains connected and active and operational. When they drop that big piece of load in Yuma, Arizona, and the grid tries to jump up; all of those distributed, active nodes immediately pull down. They all pull down on the grid to try to bring the voltage back where it should be. The ripples don’t go out across the system because they’re being damped by these active nodes as it tries to spread.

So it probably doesn’t even reach San Onofre. They get a signal at the control area and they say, “You know what? All of those active nodes out there are working at 95 or 100% to try to keep the voltage down and we still have the voltage up at plus 8%. Let’s shut down some generation.”

You call up these guys and tell them to shut down a few plants. They’re offline. Those active reflexes in the system start to relax, relax, relax. Now you’ve got the system all under control.

Remember about the trouble of trying to reconnect a load after it has been dropped? If you have an actively controlled grid, you can do that very quickly.

Remember we heard earlier today about the trouble of trying to reconnect a load after it has been dropped? If you have an actively controlled grid, you can do that very quickly. Rather than worry about “Oh, let’s see. If I start with this grid generator, the grid voltage might go too high. If I close that re-closer over at the substation, it might put too much load on the system.” And you’re kind of doing this dance.

Instead, you start throwing switches. If you add generation, the voltage tries to go up and all the reflexes pull it down. Then you go add some load; that tries to pull it down and the reflexes relax. Maybe they support a little. Then you throw on more generation. You can very quickly bring the system back up.

The grid of the future will be active, autonomous, distributed, and stabilized. Control nodes will be fast, precise, and compact so you can put them where they need to be.

This is what I believe is our grid of the future.

Active, autonomous, distributed stabilization nodes that are fast, precise, and compact so you can put them where they need to be. These power electronic nodes can enable us to create the best grid in the world right here in the United States!

Thank you.

Energy Storage for the Electric Enterprise: Opportunities and Challenges

DUSTY BECKER:

Our next speaker is Dr. Satish Rajagopalan. He is a Project Manager in the Power Delivery & Utilization Sector at EPRI, where he is responsible for power conversion and energy storage research, testing, and development. Satish Rajagopalan received his PhD in Electrical Engineering from Georgia Institute of Technology, an MS in Electrical Engineering from Iowa State University, and a BE in Electrical and Electronics Engineering from the University of Madras. His present areas of research are in utility energy storage systems, electric transportation, and advanced distribution automation. He has authored and co-authored numerous technical papers and reports and holds five U.S. patents and has two pending.

DR. SATISH RAJAGOPALAN:

Thank you very much, Dusty. We have a lot of people talking about energy storage today, so I thought I'd give you a more in-depth view
of energy storage for the utility grid. We can look at the opportunities and challenges that we face.

So before I get into that, I'll just give you a brief overview of EPRI and the energy storage program. We are a $6 million program, a very small program inside EPRI, considering the size of EPRI. Our main focus areas are understanding storage technologies; identifying and calculating the values of storage – that is, monetizing the value of storage; looking at technology development; looking at specifications; developing laboratory and field demonstration systems; trying to understand their performance and shortcomings, and so on; and developing grid-level solutions -- how you control it, dispatch the storage systems, control systems, so on and so forth.

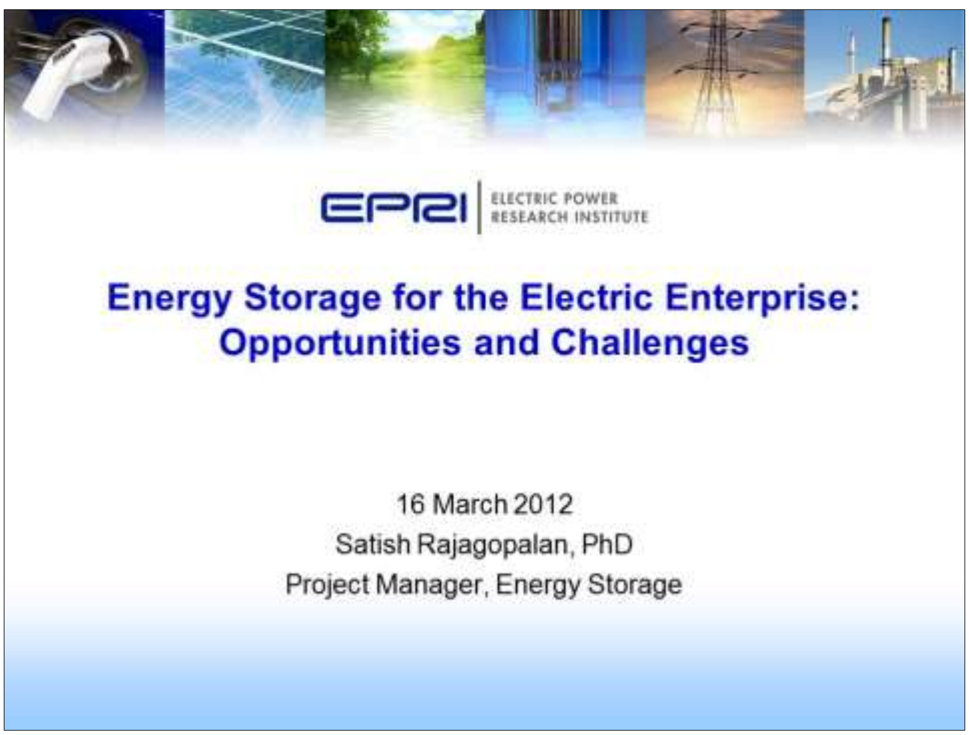

EPRI Energy Storage Program Mission

Inform stakeholders about the capabilities, costs, values of energy storage, and facilitate availability and implementation of storage options for the grid

- Understanding storage technologies
- Identifying and calculating the values of storage
- Specification and testing of storage products
- Implementation and deployment of storage systems

EPRI Research Portfolio

- Maintain Technology Awareness and Provide Guidance
  - Understanding the underlying technologies
  - Advise on Performance Targets
- Facilitate Availability of Storage Products
  - Functional Requirements
  - Standard duty cycle development
  - Testing and Evaluation
- Facilitate Grid Integration of Storage
  - Grid deployment objectives and logistics
  - Tools and methodologies for planning and operation
- Analyze Uses, Benefits, and Costs
  - Cost Projections
  - Business Cases
  - Regulatory developments

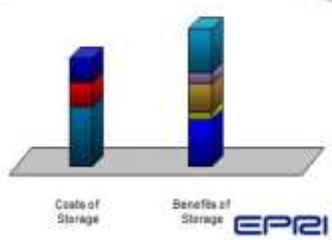

In more detail, we maintain technology awareness and provide guidance. We keep a constant watch on various battery-storage technologies that are out there. We advise on performance targets. One of our main goals is on technology development, developing functionality requirements, duty-cycle developments, testing and evaluation, grid control, and dispatches in the area. That is something in which we are very deeply involved, especially as regulatory utilities have not much insight into how we can use and control and develop these storage systems on the grid.

Finally, we look at the economics, the cost predictions, so on, and so forth. Our main focus is basically on stationary energy storage for grid applications. We also maintain a secondary role, looking at regular storage technologies, secondary storage, grid applications, focusing in that area.

How can you monetize energy storage and sell it?

This is an outline of my presentation today. Briefly talk about the drivers for energy storage solutions; overview of the applications that are out there; the various energy storage options available; the value of storage; how you can monetize and sell it; and, finally, the research needs regarding storage.

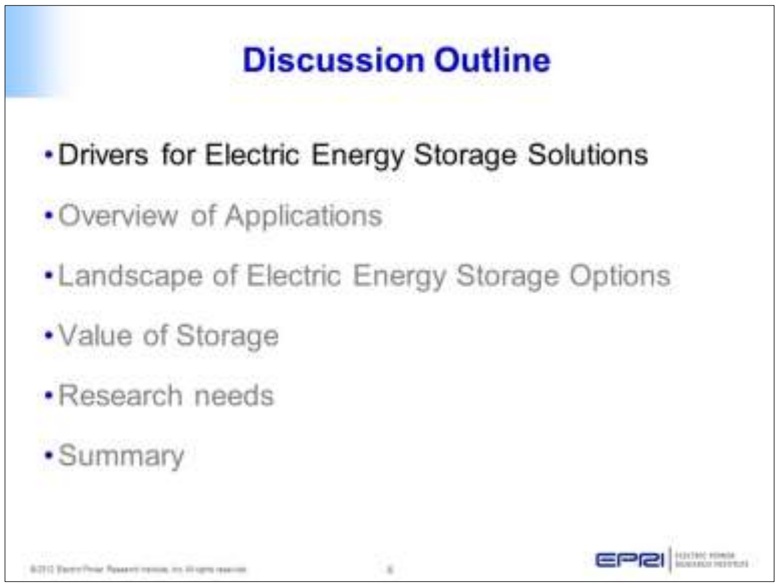

What are the roles of storage on the grid? Well, you can pretty much put storage anywhere on the grid. You can start off with bulk level. Bulk storage is used for transmission assets deferral, distributed arbitrage, maintenance of system stability, frequency regulation, and so on.

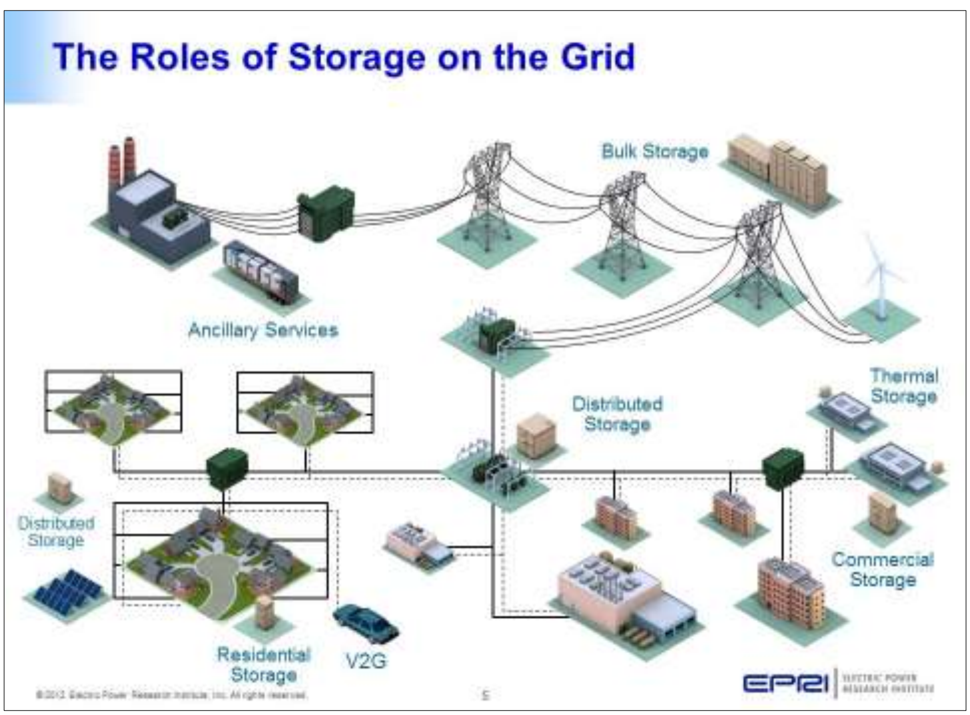

Looking at ancillary services, voltage and frequency regulation; as we keep coming down from the transmission to the distribution scale, we can look at distribution storage. We can put storage systems at the main substation, mostly for distribution and applications. As we keep coming down, we can put distribution storage at a substation level or a community level.

The size and type of storage are dependent on application and value of putting them on the grid.

The key thing to keep in mind is that each of the storage technologies, depending on their role, is very different. You'll be looking at different kinds of storage technologies, different sizes of storage technologies. For example, when we are talking about bulk storage, these are usually pumped hydro kinds of storage. These are large storage systems, typically, with 100 or so megawatts. As you get down to the distribution storage system, distribution storage scale; you are mostly looking at electrical storage, maybe some thermal storage as you get to the distribution levels. So the size and the type of storage are very dependent on the application and value of putting them on the grid.

Typical distribution storage applications at the substation are anywhere from 1.0 to 5.0 mW/hr in range. As you get down into the community-level storage systems, you are talking about 2,500 kW/hr; and as you get further down to the residential scale, you are talking about anywhere less than 10 kWh system.

Our storage is pretty much stationary most of the time. It's used for less than an hour in any 24-hour period.

There are other storage options out there, a little bit futuristic, which may be very good for the grid. As wind and PV penetrations increase, there's a big potential for using those batteries for grid ancillary applications. Keep in mind that our storage is pretty much stationary most of the time. It's used for less than an hour in a 24-hour period. So you have this whole fleet of assets that are basically sitting there in place. If you can communicate with them, if you can break it down; you can take it to ancillary services.

There are a lot of issues and challenges with that. Who pays for the battery? Because you are going to wear down the life of the battery. Those are issues that need to be sorted out. It's a technology that we are looking at, but it's a little bit more futuristic… another 10 years down the line, I would say.

The other basic important storage is in commercial locations, distributed storage, and other external locations. Instead of single-phase systems, these are typically 3-phase systems. These could be the electrical storage systems or thermal storage systems, systems that are not based on electrical storage.

There's more and more penetration of wind and PV, which are variable sources.

So what's been happening in the grid that's driving the need for storage? There's been an increased penetration of variables, more and more penetration of wind and PV, which are variable sources. There's increased need for flexible grid, ancillary services, and balancing as the other gentleman from California ISO was talking about. We are now dealing with variations and generations that stress the grid, so the complexities are increasing.

The cost of assets is increasing too, so anything we can do to defer the need for capital will help improve the viability. Mitigation is another driving factor. Peak energy management and end-use energy management at the residential level are some of the applications that are driving storage, but they are mostly in response on the smart grid level, as Dennis discussed.

So a lot has been happening in the last five or ten years that has really put energy storage in a prominent position, as far as its role on the grid is concerned.

We have identified ten key applications along the electric value chain.

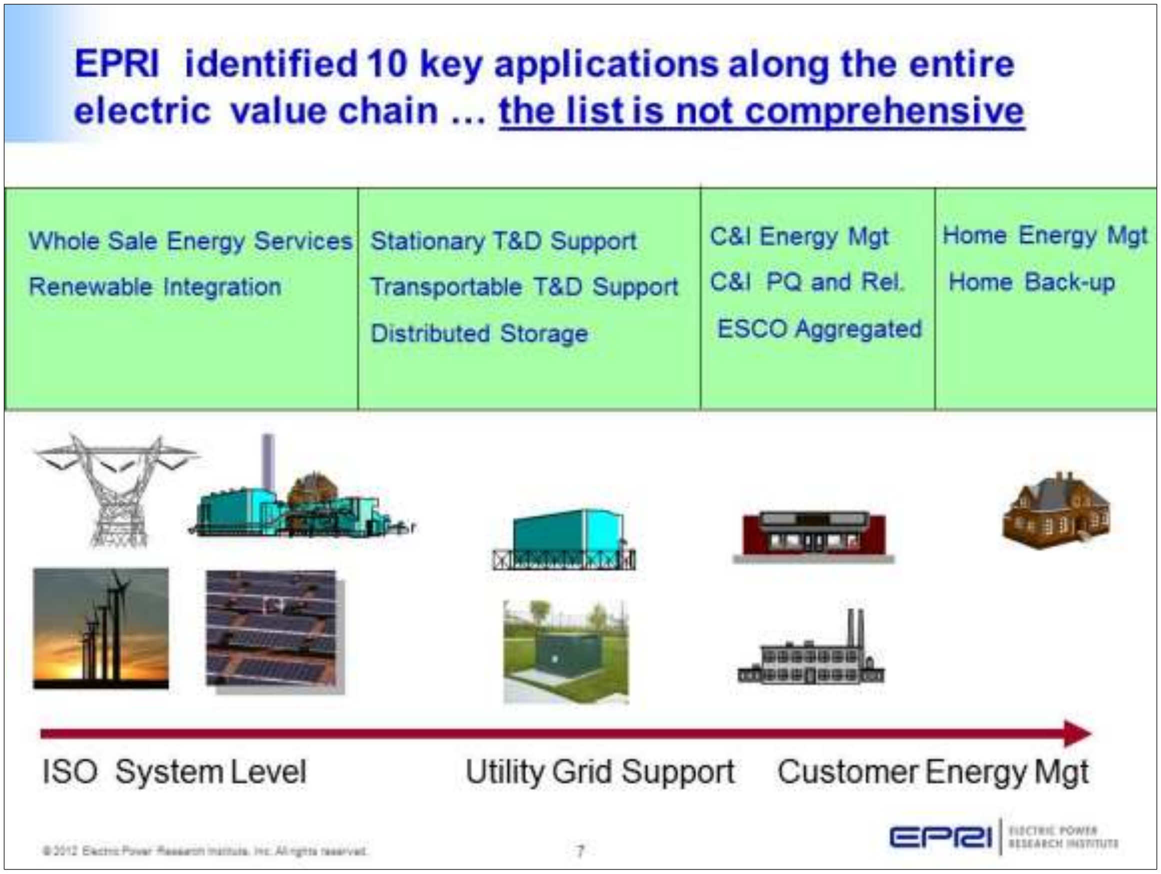

So EPRI has been working with the electric utilities and we have identified ten key applications along the entire electric value chain. There are a lot of applications for storage. You need to prioritize because it cannot work for each and every one of those applications.

So at ISO system level, wholesale energy services, energy arbitrage, renewable energy integration, and frequency regulation are the most important applications. Going down into the distribution levels: stationary T&D support, primarily for capital and deferment of distribution resources; distributed storage in support of lower-level applications; and supporting variables at the distribution scale. Those are some of the key applications. Further down into the local and industry scale, energy management and peak-level control are very critical.

As we get into the small-scale systems, which are pretty much the small-scale distribution systems as well as residential systems, there could be several business models. It could be utility-regulated models, it could be a third-party company. It gets all the energy storage assets and then employs the service.

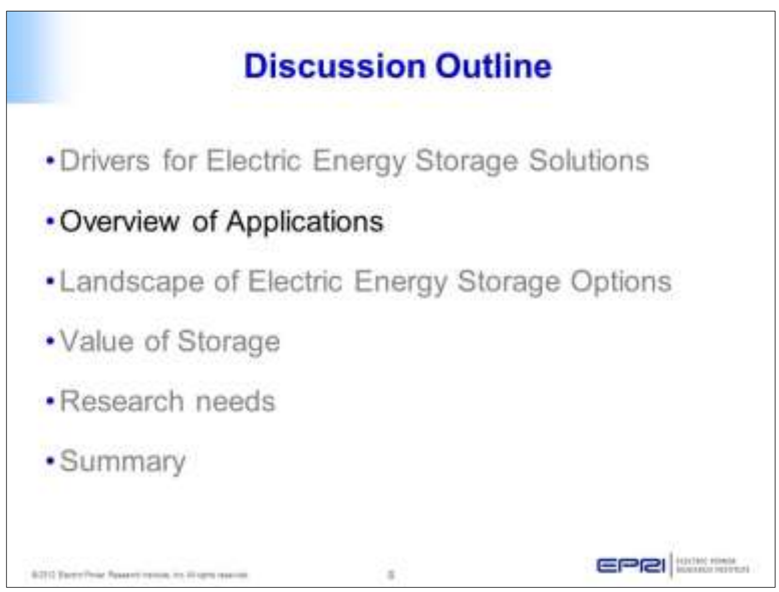

There are various ways we can monetize energy-storage applications.

So there are various ways we can go about monetizing these applications, but these are the primary applications that we think would make the case for energy storage in the near future.

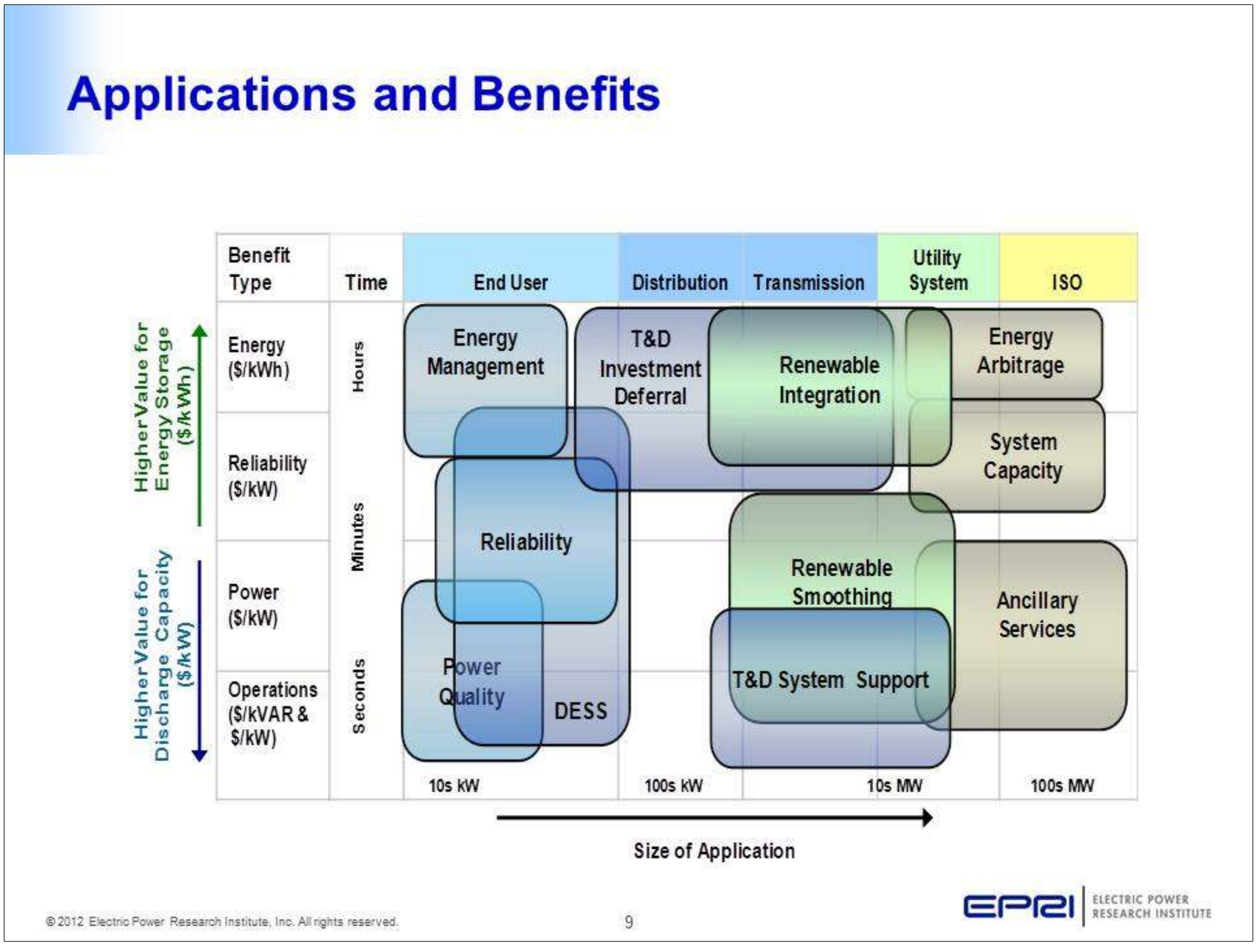

Looking a little more in depth into the applications, this chart shows the whole set of applications. Your Y access is basically showing time, that's the response time that will be needed for the energy storage system. The exact response time depends on the size of a system. Depending on how fast you want your energy storage system to respond and how big your energy storage system needs to be, you'll need to pick an energy storage system suitable for it.

Large-scale energy storage systems are energy arbitrage systems, mostly pumped hydro.

There are applications all over the map here. Large-scale systems are pretty much energy arbitrage systems, which are basically hydro pump kinds of systems, which are slow response systems, slow to respond. Ancillary services and frequency regulation are typically faster on the response scale.

Then you come down to applications such as power quality and availability. You need much faster-responding systems, especially at the power-quality level.

Another thing to keep in mind is these applications are very divided: energy-intensive applications and power-intensive applications. This is another key thing to keep in mind because you can pick a lithium-ion battery, for example, a power-intensive battery, or an energy-intensive battery; and they are completely two different pieces. They are not the same in terms of cost, operations, or design.

Energy storage applications are very divided: there are energy-intensive applications and power-intensive applications.

So your goal, your size, your application, your response time, and whether it's a power- or energy-intensive application, will basically determine what kind of storage you pick for your solution.

Let's just look at a few of these applications in a bit more detail. Peak shaving for asset upgrade deferral and transmission distribution level aspects is one of the key applications for energy storage. Most of these applications are peak shaving kinds of applications where you're basically cutting down the peak and ensuring you that can defer substation upgrades for a few years, maybe 3 to 5 years. Typically, a substation costs millions of dollars to upgrade. If you can postpone that upgrade for a few years, that could be really helpful.

There are several systems deployed right now. There are exploring substations for peak-shaving, load-leveling kinds of applications. This one is basically a sodium battery by AEP. This is a flow battery that is applied by DTE. This is another flow battery, PacifiCorp. There is a whole range of technologies for distribution asset upgrade.

Peak shaving for asset upgrade deferral is one of the key applications for energy storage.

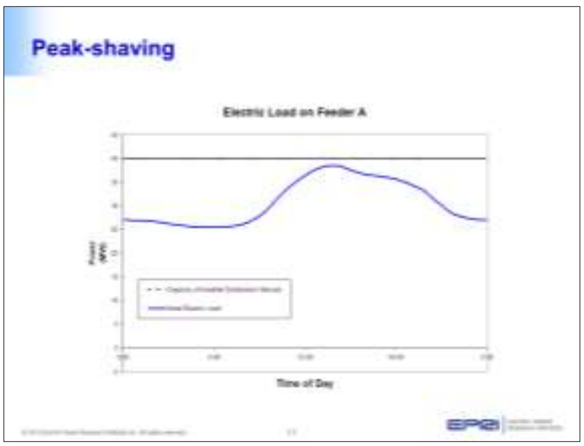

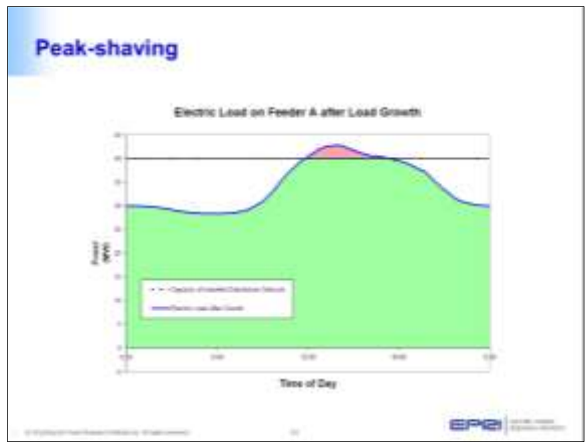

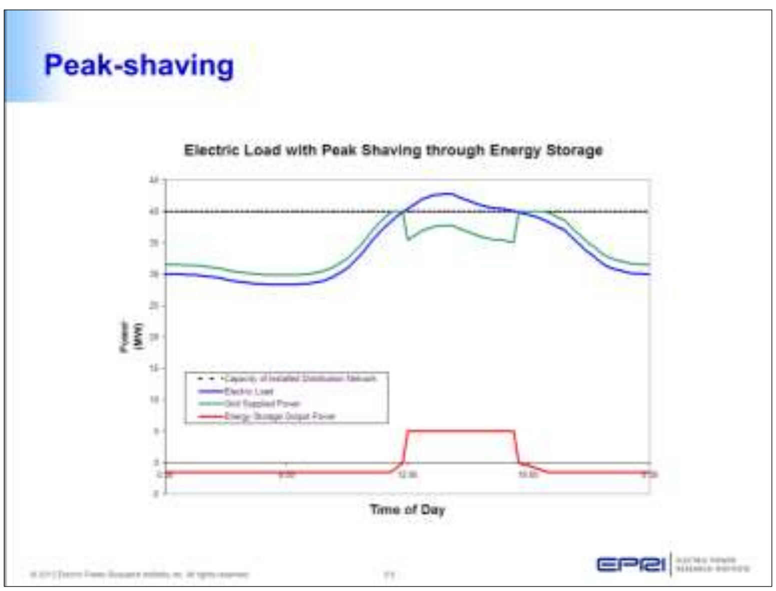

A peak-shaving application is a very simple application; it's an energy-intensive application, not a power-intensive application. All you are doing is making sure that the peak on your feeder is basically contained. When the peak goes above a particular threshold, the battery kicks into action and supplies energy or stores energy. You can think of it in either way. So in clipping your peak load, you are basically deferring the need to upgrade your substation for a few years.

As I mentioned, there are large-scale penetration of variables right now on the grid. The penetration is at bulk level, transmission level, and distribution level. Each of those cases presents unique problems. If you have a large-scale penetration of variables on bulk level, typical problems you're going to encounter are stability issues, frequency regulation, so on, and so forth. These are especially critical in islanded grid systems; like, now, basically, we are at the mercy of whatever happens with the wind or the sun. There are large-scale variable penetration through PV and wind. Anytime the PV or the wind fluctuates, the intensity basically fluctuates because there is a small increase. They have large-scale energy storage systems to mitigate and stabilize the grid.

If you have distributed PV on a distribution system, there are a lot of voltage-control issues.

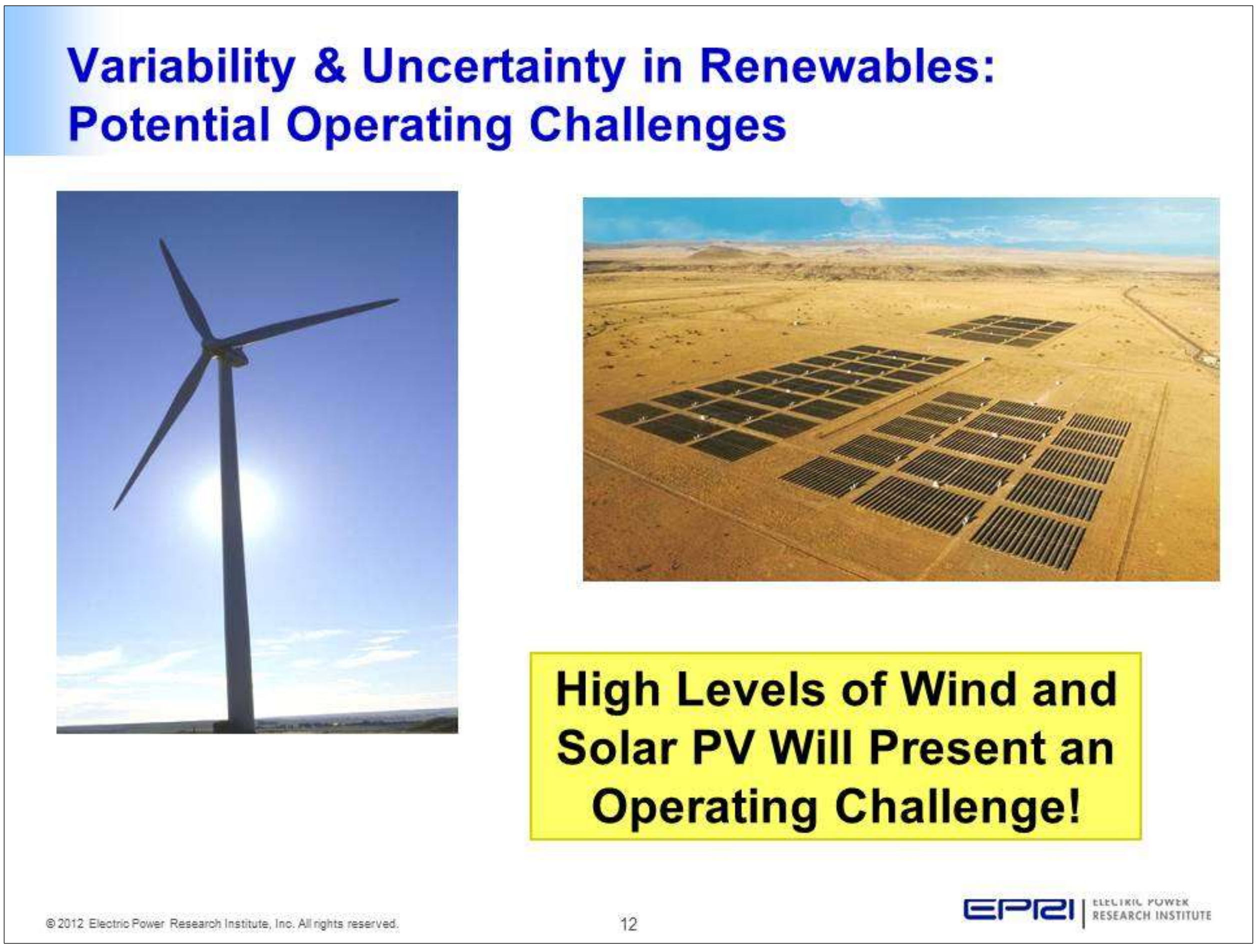

At the distribution scale, the problems are quite different. If you have distributed PV on a distribution system or a small-scale PV, there are a lot of voltage control issues on the distribution feeder. Typically, your work area needs to be maintained at 120 volts, plus or minus 5%. If you have a large-scale PV penetration on your distribution system, there are massive voltage deviations and fluctuations. These fluctuations actually may go beyond the 5% limit. Once that happens, feeder regulators and capacitor banks are called into action to correct those abnormalities.

You'll see they are looking at several sites with very dense PV penetration. You will see the tap changes to the capacitor banks are operating hundreds of times a day, trying to keep the voltage within limits.

This is not a typical design specification for feeder tap regulators or capacitor banks designed to switch on maybe once or twice a day. You're forcing them to switch hundreds of times a day, basically wearing out all the systems. So we can use energy storage to smooth out those critical problems and ease the burden on those feeder tap changers. Of course, it is a cost issue and a balance issue. We are doing the monetization calculations to see where the value benefits lie in terms of cost for energy storage.

We are doing monetization calculations to see where the value benefits lie for energy storage.

That is basically the wind generation fluctuation that you will see typical in a wind farm. Just like someone said this morning; if you go and stand in the wind, you don't feel anything. If you actually measure the wind output, you'll see that it's all over the place. That's almost impossible to categorize or classify. Wind power is a cubic function of wind speed, so even small deviations of wind speed cause massive deviations in power.

At the bulk scale, this causes large-scale regulation frequency issues because supply and demand have to very quickly match. In cities, you'll see that many of the wind plants, wind farms that are installed typically get installed with energy systems, because you want to smooth out all of those fluctuations.

This is a case of a PV output very similar to the picture that Greg was showing. As the cloud goes over, you see massive fluctuations in the PV. This is PV at a distribution scale. Once you have all these kinds of fluctuations, you have corresponding massive voltage fluctuations on your distribution grid. So you want 120 volts, but you could be fluctuating anywhere between 110 to 130 volts. To correct this, capacitor banks are switching really fast to correct these abnormalities. This is not what they were designed for and you're going to wear them out very quickly.

You can use energy storage to smooth out power fluctuations.

So you can use energy storage to smooth out these power fluctuations. Let's say that the energy storage is doing some other real power functionality and you have some kind of VAR capability where you can even use that VAR capability to smooth out fluctuations. So you can deal both with the real and reactive power level. This is one case of having four-quadrant energy storage systems. You have both the capability to build your VAR control, as well as your real power control, to solve these issues.

High levels of wind and solar PV keep presenting operating challenges, especially as the penetration of variables continues to increase on the grid.

This is an example of distributed PV without energy storage and with energy storage.

What is happening is that, in some cases, you have an excess of power generation. You have the voltage going up and vice versa. You basically use your energy storage to solve that so that you can keep your voltages within the spec. It's pretty straight forward.

Distributed PV without Energy Storage

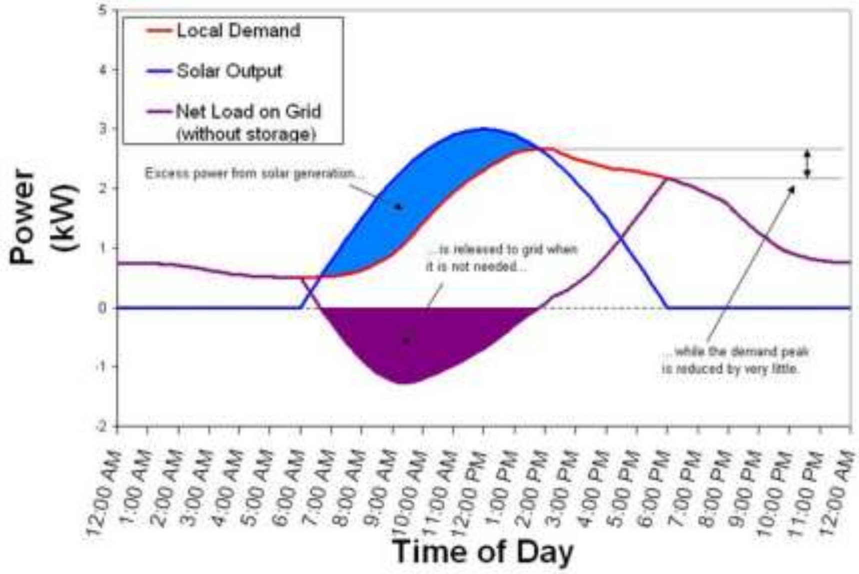

13
EPRI | ELECTRIC POWER RESEARCH INSTITUTE

Distributed PV with Energy Storage

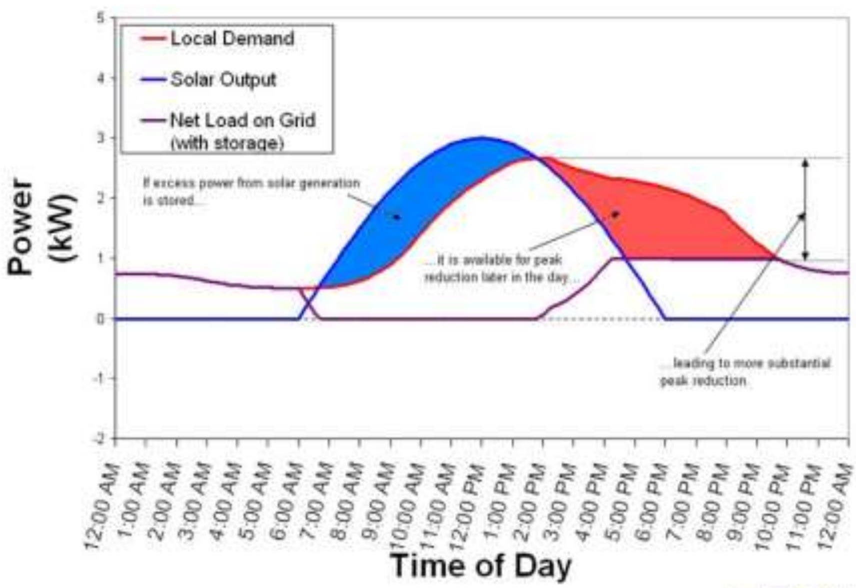

14
EPRI | ELECTRIC POWER RESEARCH INSTITUTE

Frequency regulation from variable loads is a key problem. You need to balance the dynamic load versus the dynamic generating resource.

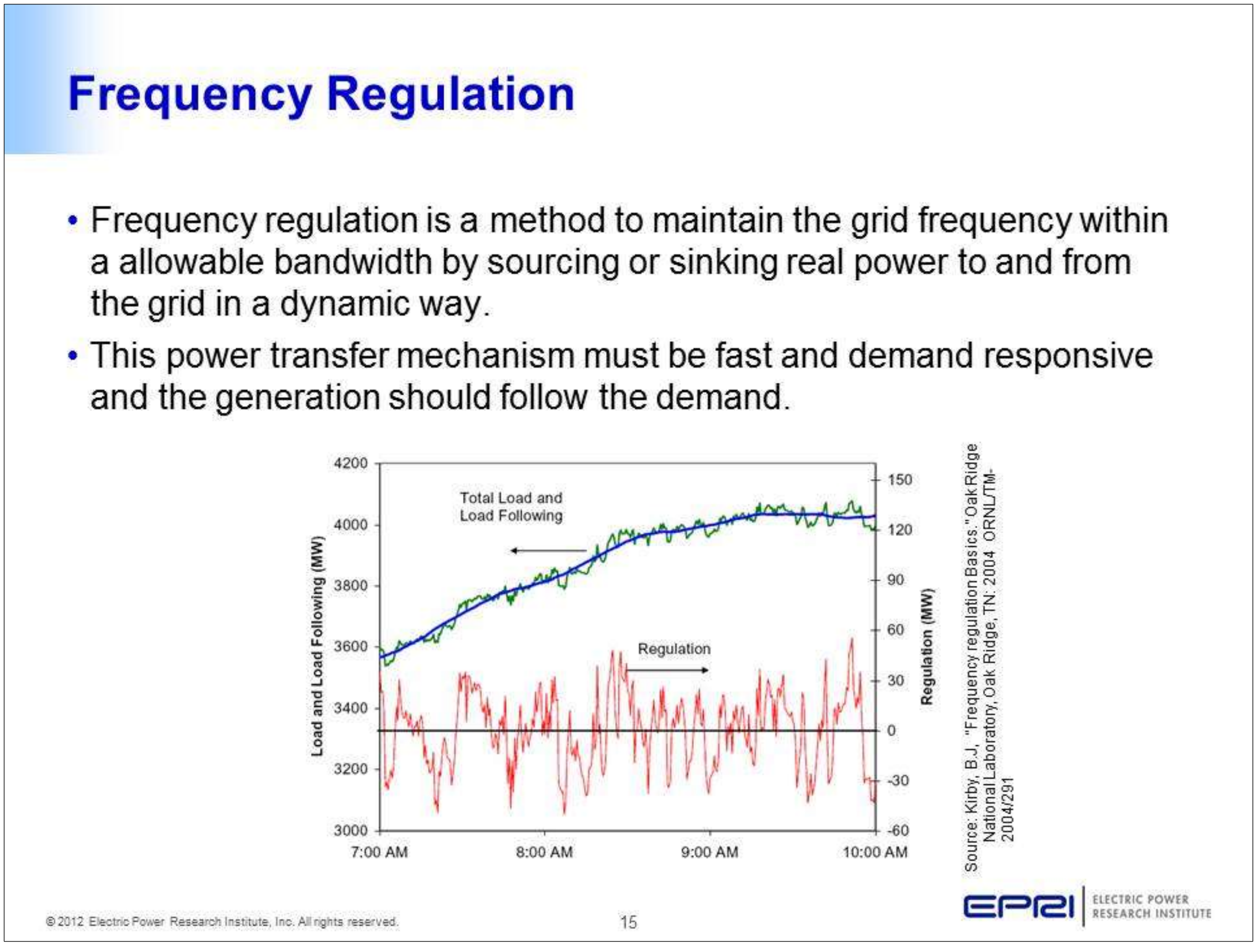

Getting into frequency regulation... frequency regulation from variable loads is a key problem. You need to balance the dynamic load versus the dynamic generating resource. Frequency regulation is also needed, even in the case of non-battery generation. In any case, you need to manage the generation.

Frequency regulation is basically a high-power, fast demand-response application. The requirements of the battery are very stringent. The operating battery will be put through a very dynamic duty cycle detrimental to the life of the batteries, very subject to the actual charge and discharge rates that needed on the battery.

Keep in mind frequency regulation is probably one of the most well-documented monetized cases for the use of energy storage today because the wholesale market is much more defined in terms of valuation than the distribution market. The important thing is that the energy storage is being pressured by gas-based applications, which are now coming in and offering their services at a much lower cost. They are fast-response systems. That is intense competition as to what kind of resources to use for frequency regulation. Gas is really cheap. They are providing very stiff competition into other commercial as well as non-commercial means of frequency regulation devices.

Some examples of large-scale systems: This is an AES system. I think it's in Pennsylvania. It's an A123 battery system. It's a 32 megawatt, 8 megawatt-hour energy storage system.

This is a 1 megawatt, 15-minute flywheel system; in a similar frequency-regulation kind of application.

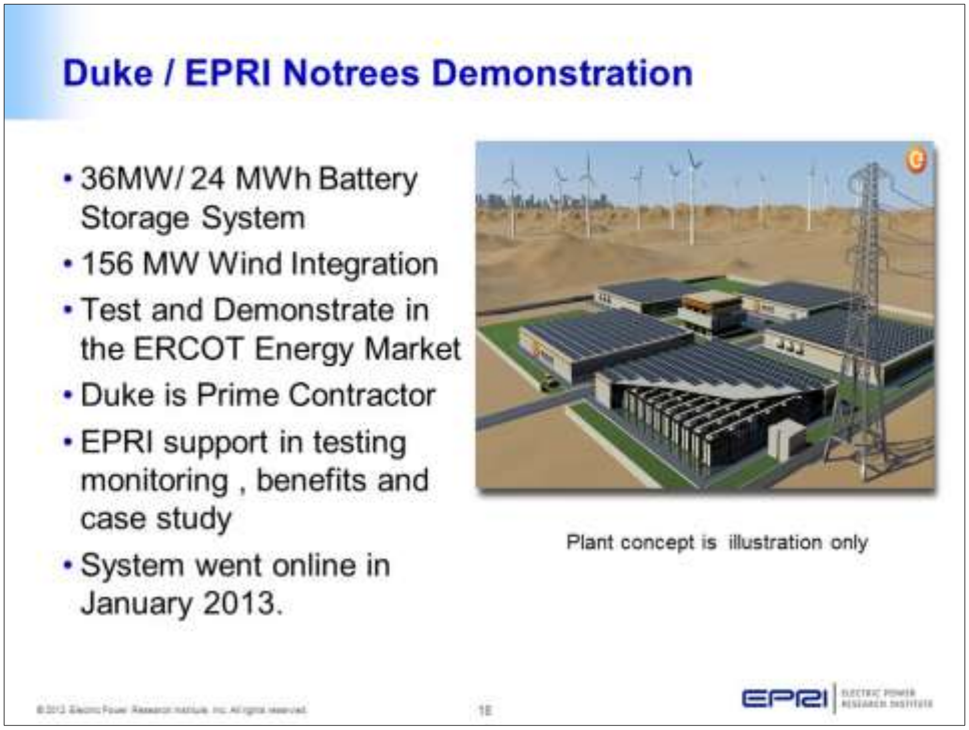

We just had a 36-megawatt / 24-megawatt-hour lead acid-based battery system go online in support of a 156-megawatt wind installation in northeast Texas. That battery will be used to test and demonstrate the value of energy storage, not only in support of wind integration, but for the whole energy market. I think it's the largest battery system in terms of power and energy together in the continental United States. Although there is a larger system in Alaska that is still in operation.

Each application has its own, individual characteristics.

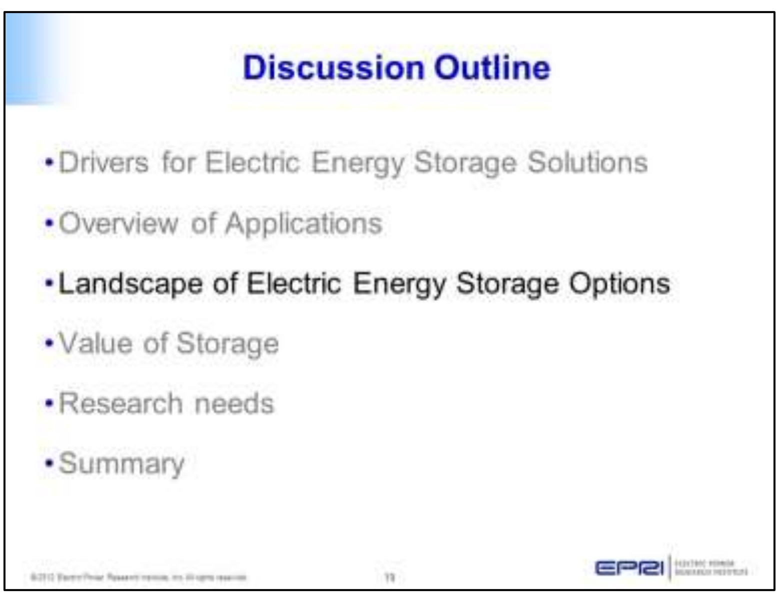

So those are representative of the applications out there. As I mentioned, each application has its own individual characteristics: the amount of power, the amount of energy demanded, the response time. All of these actually dictate the kind of solution you are going to use for the application.

There are a lot of energy storage solutions out there. Just to give you a brief overview, this is a pumped hydro of facility operated by TVA in Chattanooga. This is a compressed energy storage system in Alabama. This is the only compressed-air storage system in the U.S. right now. You have other battery-based storage systems that vary from sodium-sulfur based chemistries, flow batteries, lithium-ion batteries, fly wheel systems, and so on and so forth. These are all typical distribution- or transmission-level, large-scale systems.

Coming down into the small-scale distribution systems, you have sodium-metal-halide-based systems coming out now. A number of lithium-ion systems are being deployed anywhere from substations down to the residential scale. You have a lot of options out there. It all depends on the actual application that you are trying to target.

Lithium ion systems are being deployed from substations to residential scale.

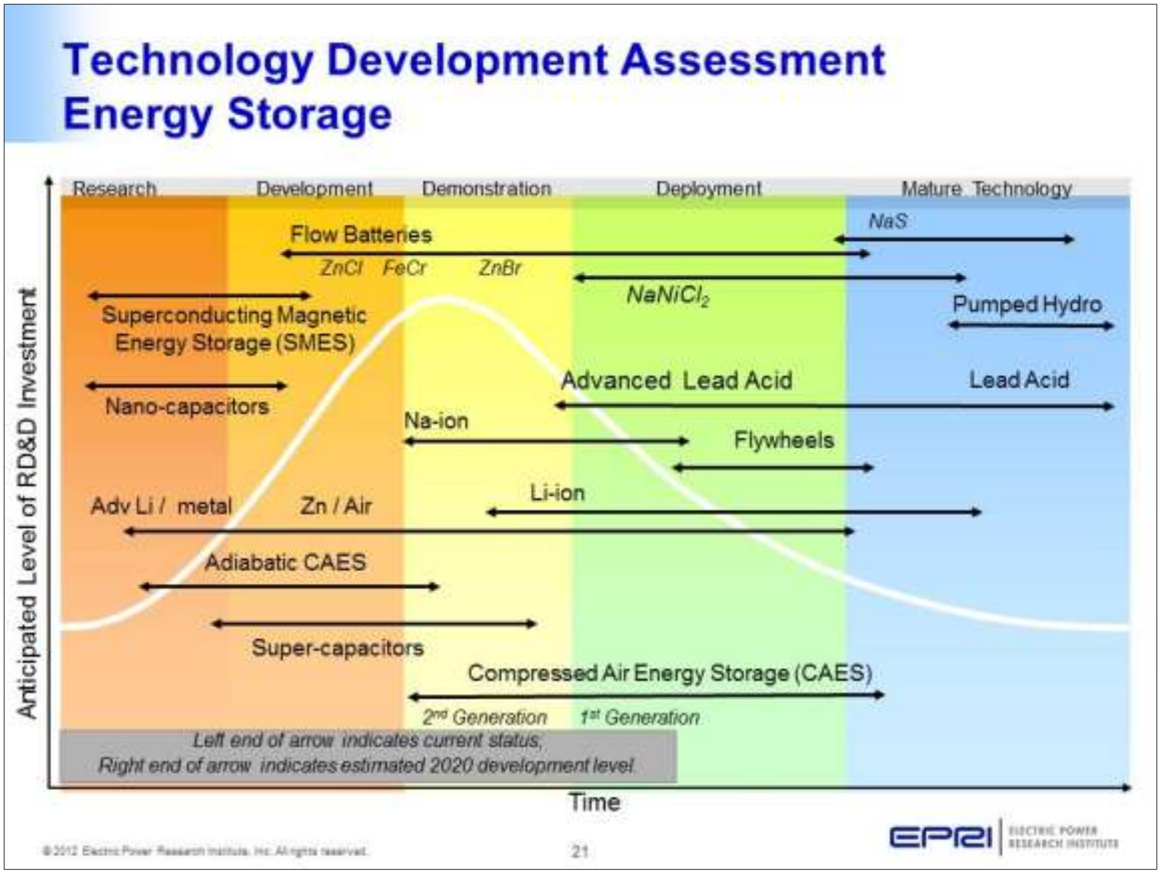

What is the state of the technology of all of these energy-storage options? There are a number of energy-storage technologies out there. There's a lot of research happening in the inventory space. This picture basically gives you a broad idea of where various energy storage technologies stand today. I'm not going to go through each one of them in detail, but if you look at main categories, each them have several sub-technologies and subcategories under them. For example, there are several dozen different types of lithium-ion battery systems out there.

Most of the technologies today in distribution space have been in the lithium-ion. This is because of development of EVs and all of the money that has been put into the development of lithium-ion batteries.

Stationary storage is benefitting from this. Advanced lead-acid-based systems are probably the most mature, but they are very limited in terms of energy density. These are large systems in terms of size; their cycle lives are low; there are environmental issues regarding lead; and many of the utilities are removing lead-based storage systems nowadays. Pumped hydro is probably the most mature of all the bulk storage systems.

If you look at other transmission and distribution applications, flow batteries are in the deployment stage right now. They are still not mature, but getting there. Sodium nickel chloride chemical systems are probably the most mature system, but these are typically used for large-scale transmission applications.

Each technology has its own advantages and disadvantages.

If you look at this particular picture, it looks like there are several solutions that are almost mature. Keep in mind that each of these technologies has its own advantages and disadvantages. That's why you'll see enormous research into other chemistries and other technologies still happening. Notably, if you take lithium-ion, for example; yeah, we have problems to a point. These batteries are much safer and we have been able to increase the energy densities, but we still need to go much further.

So there is basically a large amount of government-funded research ongoing right now thanks to a lot of money being put in by DOE. They are very much into all of the financial matters. I will talk a little bit about this later in the next few slides.

99% of energy storage is pumped hydro.

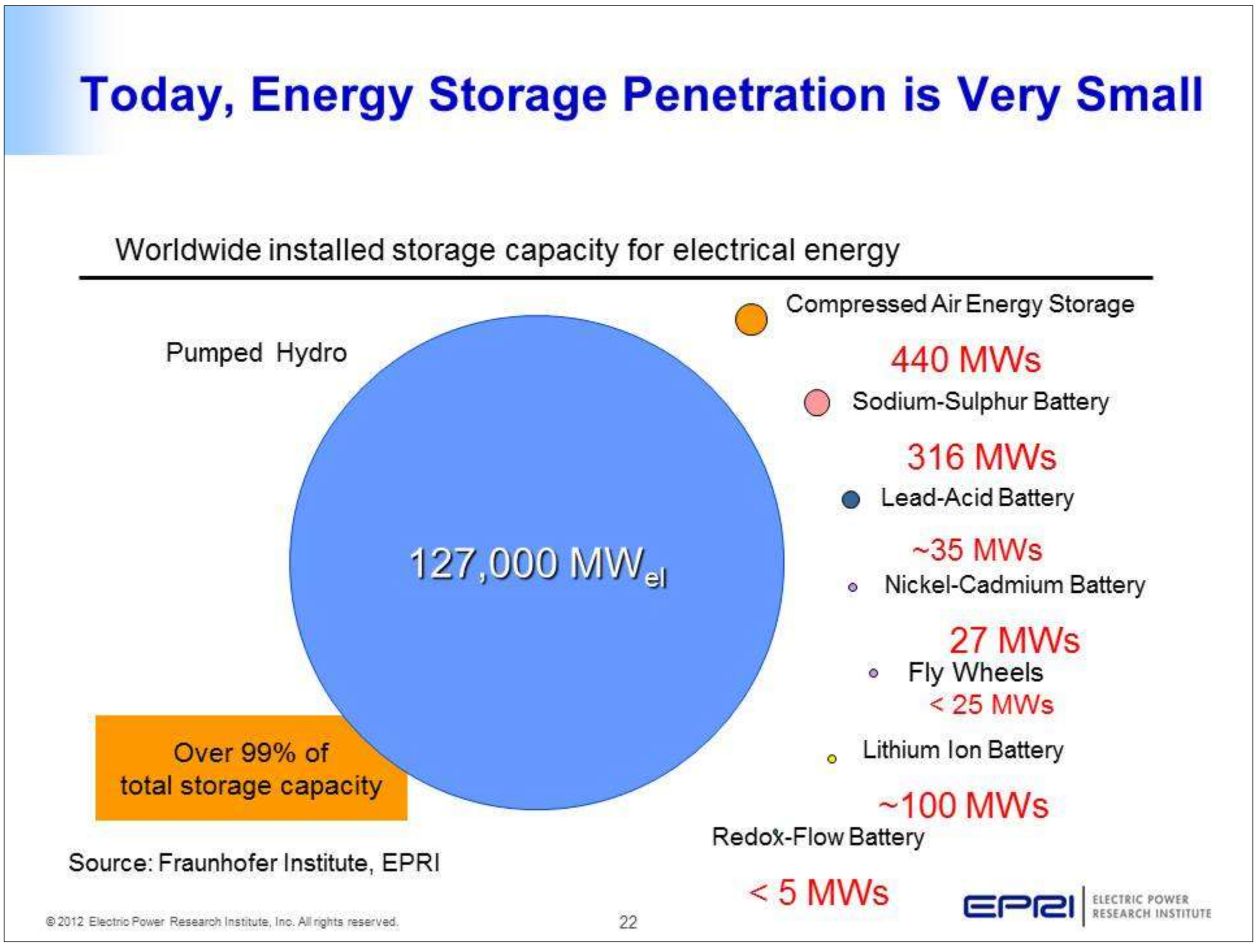

If you look at the overall landscape of energy storage, it is 127,000 Megawatt hours worldwide. If you look at it, 99% of it is pumped hydro. 1% is all the other stuff, basically composed of compressed air energy storage, which is bulk. There are other storage systems just coming up that are basically a fraction of all the other energy storage. This is rapidly increasing as electrical energy storage gets more and more mature and we expect this to go farther as there are a lot of opportunities right now.

Our main focus now is on electrical energy storage at the distribution level. We are also working quite a bit on compressed-air energy storage. But this landscape is our main focus: electrical storage at the distribution level.

This picture basically shows the areas of lithium-ion demonstrations that are out there. This is a little bit dated; there are several more deployments coming out.

This shows the number of field tryouts for demonstrations has rapidly escalated in the last few years. All of this has just come up within the last couple of years. We see a massive surge in storage-based applications and demonstrations happening.

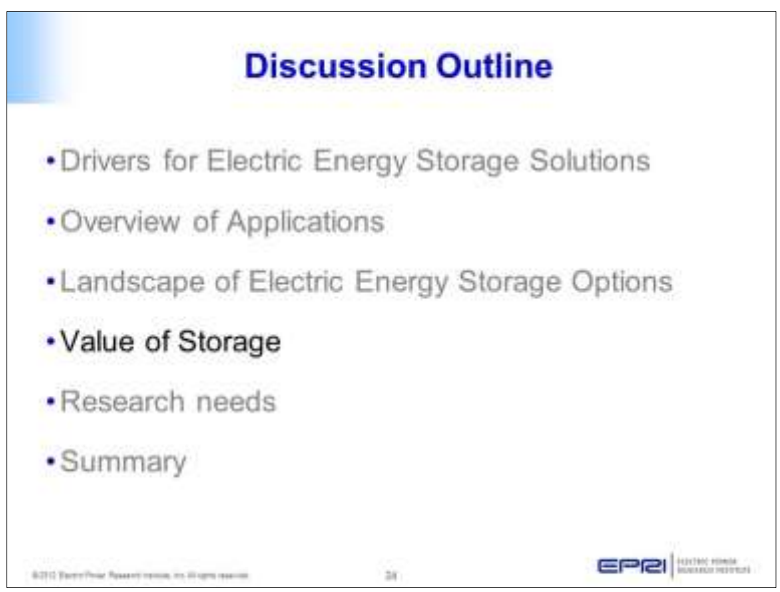

Energy storage is really expensive. Monetizing storage is one of our biggest challenges.

Clark mentioned today that it is the Holy Grail of solutions. Yes, but only if it is cheap. Right now it's really expensive. Monetizing storage is one of our biggest challenges.

If you look at energy storage, and you want to deploy for just one particular application, I can pretty much guarantee that you will never make money on that. The only way you can make money from a deployed energy-storage application is by stacking multiple applications. That is a must. These applications need not run simultaneously, but a single energy storage system must be able to perform several of these roles through stacking if you want to really get your money back.

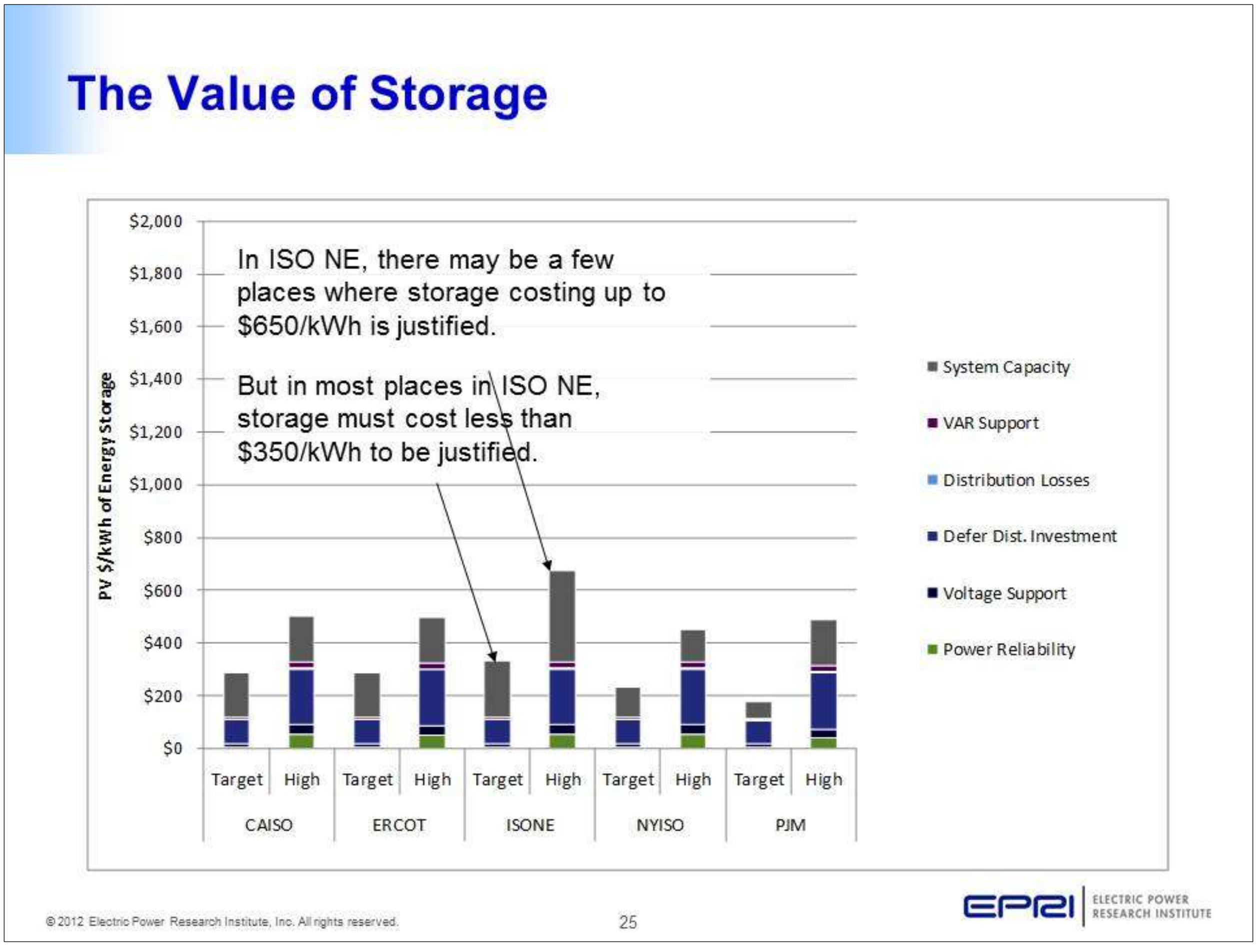

If you look at each of these conservation markets and if you try to see the cost of stacked energy, you basically see that, barring a few places, like Northeast ISO, you'll need your energy storage to be somewhere below $350 per kWh to make your money. Typically, they are about $500-600 per kWh. This chart basically shows that if your energy storage system is costing $1,000 per kWh, you're never going to make any money. That threshold needs to be somewhere nearer $200 per kWh or $100. Anything below that is going to be your profit margin where you'll make money.

Most of the efforts going on right now are trying to drive the cost of the storage price down; not just a magic price, but in terms of the value of a plan.

If you consider regulation services, your monetization case becomes much stronger.

If you tag on regulation services to all of these applications we just mentioned, your case becomes much stronger. Like I mentioned, regulation is probably the most well defined market. If you can penetrate the regulation market, your case for storage becomes much stronger from a value perspective.

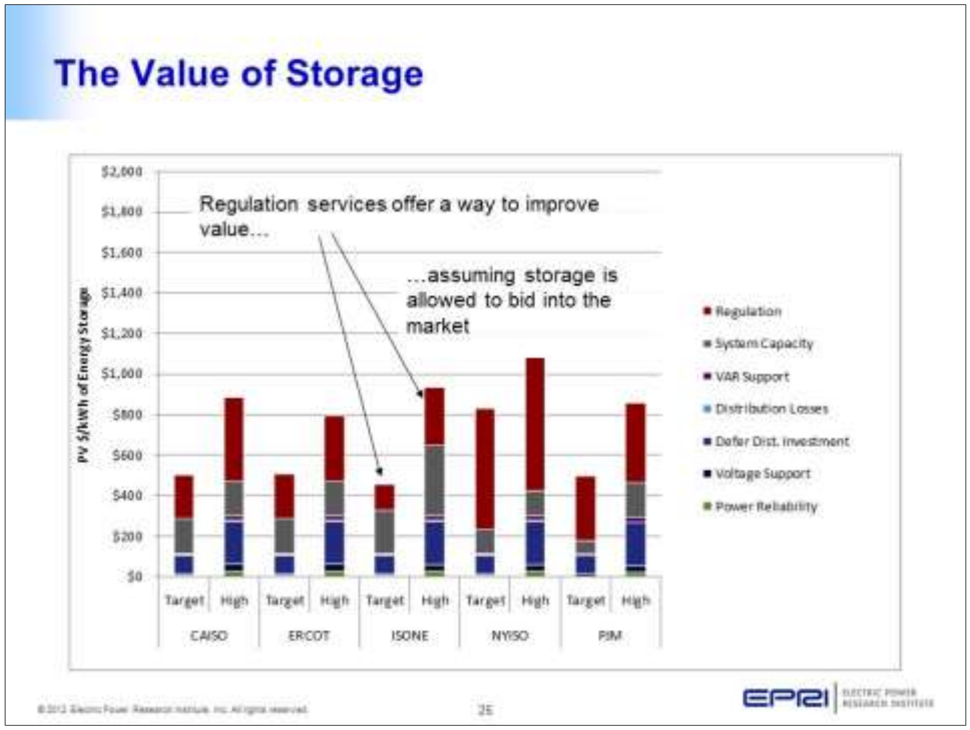

This chart is very similar to the last couple of charts. They were looking at storage from an application perspective, depending on creating a wholesale market, which is the bulk market. If you look at applications based on size of the system, you see that – let's take transmission application based systems –and you tag on regulation, even if your energy system is typically anywhere from $600 to $700 per kWh, you have a good prediction for making money.

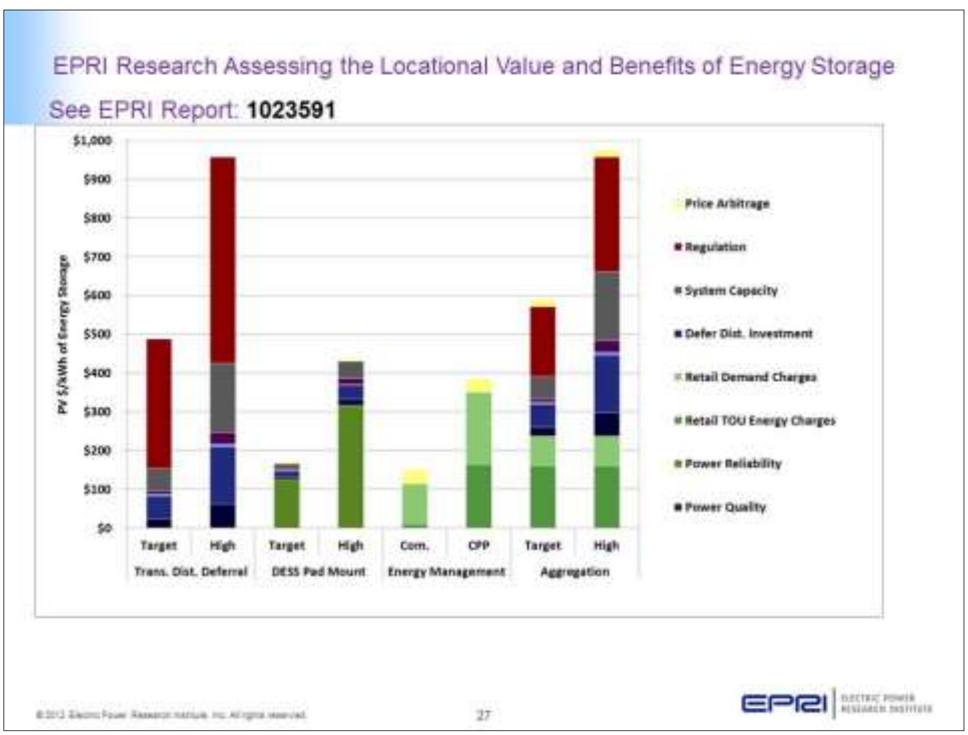

Once you get into the distribution storage markets, however you stack your system, your system costs need to come down really low before you can make any money out of it. That's the biggest challenge.

The larger energy storage systems are much less expensive per kWh of storage.

The other factor to keep in mind is that the cost of energy storage systems is very dependent on these things. The larger systems are much cheaper, compared to the small-scale systems. If you look at distribution scale storage, you can typically buy a system for $600 to $700 per kWh. You can probably justify your case for energy storage for transmission and distribution applications. Once you get into small-scale systems, like 20 kWh systems and even smaller 4 kWh systems, it could be anywhere from $1,000 to $5,000 per kWh. You're in real trouble. It is really hard to make a case for community- and consumer-based systems.

For these applications, you need to look at a few different models. You need to take a fresh perspective as to how you are going to value in those energy systems. That's going to be the biggest challenge.

What are the research needs? A lot of challenges remain. We have made a lot of progress in energy storage systems. If you asked me five years back, I would have guaranteed you that there wouldn't be one plug-'n'-play energy storage system. There wouldn't be even one energy storage system that you could buy from a vendor, plug it in, and it could work. We have been testing these systems for the last 10 to 15 years and we have never had even one storage system come in where you plug it in and it works. This was the situation five years ago.

We have never had a single energy storage system where you plug it in and it works.

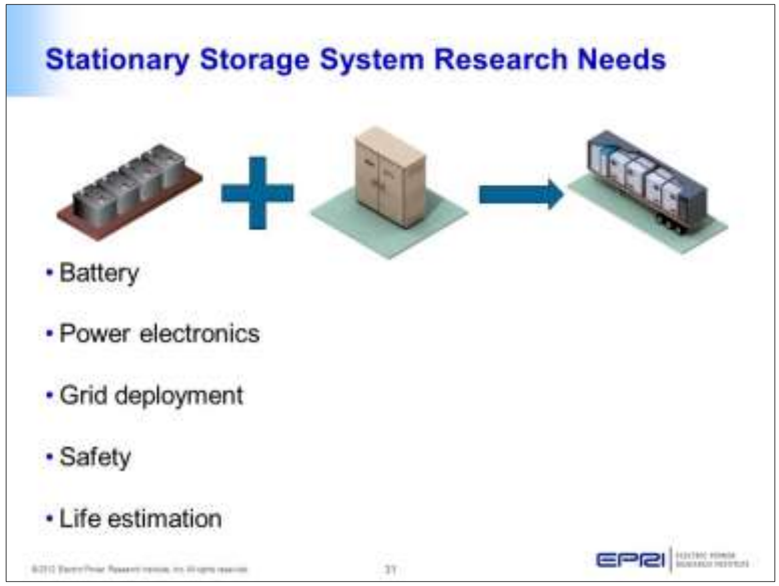

The situation has rapidly changed in the last five years. More and more players are devoting more and more money and R&D resources into energy storage. Energy storage has come much further, it's much more mature. But there are a lot of challenges, challenges about the technology, the deployment, the control. We can take a look at each of those cases in a little bit more detail.

The electric utilities, when they are looking at any energy storage system, are basically looking for rugged and reliable systems. Typically, a long life of 15 to 20 years, at least 10. Minimal maintenance; nobody wants to send ground or utility crews to maintain the battery system. It's expensive. They really don't want to do that.

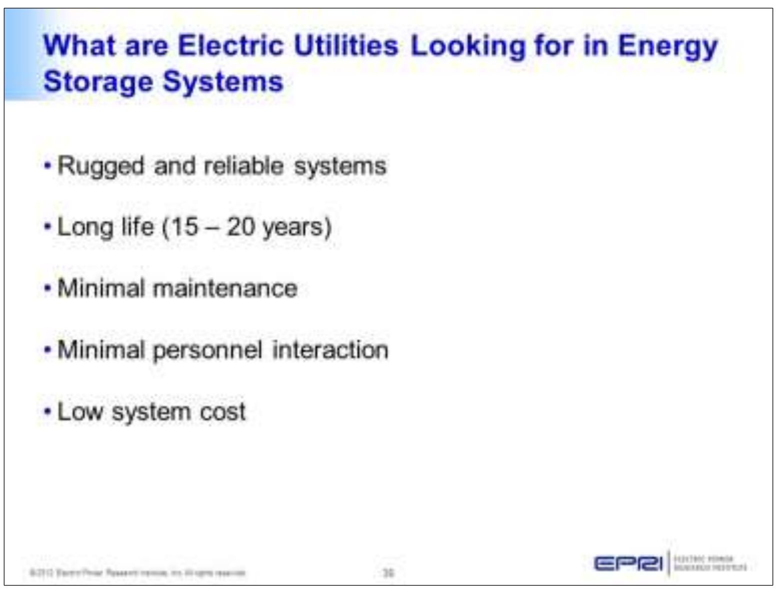

They want minimum personnel interaction. If you look at many of the substations, people think that there are people sitting in substations, all operating there. In most cases, these substations are autonomous. There's nobody there. It's just operating on its own. Some may have a "skeleton" crew of people just to make sure that the place is operational. You want to have minimal interaction with the equipment unless it's imperative.

Low system cost. Most of the utility systems are low-cost systems. That's a big need for any system that you're going to put in any energy. These are the factors that apply to any energy storage system. Energy storage is still a maturing technology. Many of these pose a big challenge. We are still in the process of working out most of those variants.

> *We have had some high-priority fires in the last year or so.*

Let's just take a simple energy storage system; a battery-based system. What is it comprised of? It's comprised of the battery and the power converter. We basically package them so that you have all your controls there. The various research needs related to any of these systems are in the battery, power electronics, grid deployment -- how do you deploy it on the grid and control it – and safety is a key issue. We have had some high-priority fires in the last year or so that have gained unnecessary attention. Bad publicity for energy storage. We want to make sure that safety is critical.

Life is really critical, like I mentioned. We need at least 15 to 20 years of life out of any of these systems, which is a big challenge. Many of these energy storage systems were not designed for long lives. Especially with the kinds of brutal applications you want: large ramp rates, large amount of power to be put in or charged in the system. The battery just doesn't hold onto those kinds performance requirements.

Long life is really critical. We need at least 15 to 20 years of life

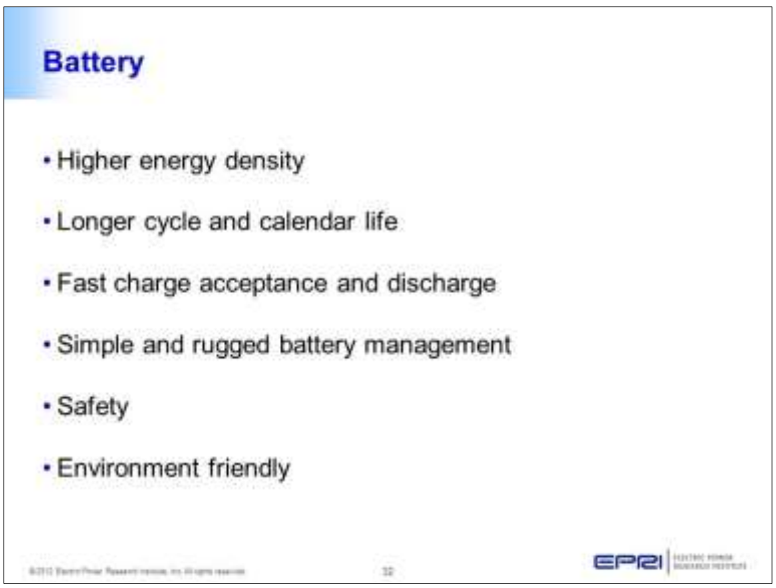

Let me start off with the battery and give you a brief overview. This is the biggest area of research right now. A large amount of money is being put in by DOE, OEMs (which includes auto OEMs), as well as battery manufacturers. From a battery perspective, the biggest thrust has been the design of high energy density batteries with long cycle in calendar life and which are safe.

If you look at an EV battery today, which is primarily the lithium-ion battery today, there are several chemical industries out there. One may ask the question, "Why are there so many chemistries?"

The reason is that not one single chemistry is an answer to all of these requirements. You have lithium-ion phosphate-based systems; there are these and many others. These systems are really safe, they have very long cycle lives, and they are low-cost batteries… but they have poor energy density. If you want higher energy density, the chemical you go to is cobalt-based chemistries, manganese-based chemistries. Once you get there, you are inherently getting into safety issues because cobalt-based chemistries and other similar chemistries are flammable.

There are a lot of challenges yet. A lot of the research has been in the chemistry area. As far as energy density is concerned, a significant amount of research is in the cathode, the anode, and the electrolyte. The idea is to push the cycle voltages higher and higher to get more energy density in those batteries.

Typically lithium ion today are 3.3, 3.8 volts. The majority of the research is pushing it to 4.5 volts. Once you try to push to 4.5 volts, there are challenges in making the blends needed for the processes you need to have for a stable cathode. What kind of electrolytes that can be stable at those voltages. So each and every aspect of battery design comes into play as you push those densities higher and higher.

The other big challenge is battery management systems.

The other big challenge is battery management systems. Especially with lithium-ion-based chemistries, for example; active balancing is absolutely critical. Maintaining diagnostics of each of those batteries cells is critical. Overcharge or undercharge is detrimental, so you need to manage them.

Finally, you need to have environmentally friendly solutions. Lead acid systems are probably the best suited for a variety of applications, but getting a government permit to put in a lead system today is really challenging. We have had discussions with our European members and, if you look at European regulations, lead acid is done. You cannot even think about lead acid applications.

You have to look at all of these advanced chemistries. Once you look at these advanced chemistries, each of them have their own problems regarding safety, balancing, control, so on, and so forth. There is a lot of research in the battery area right now.

Power electronics, surprisingly, is an area that needs more research.

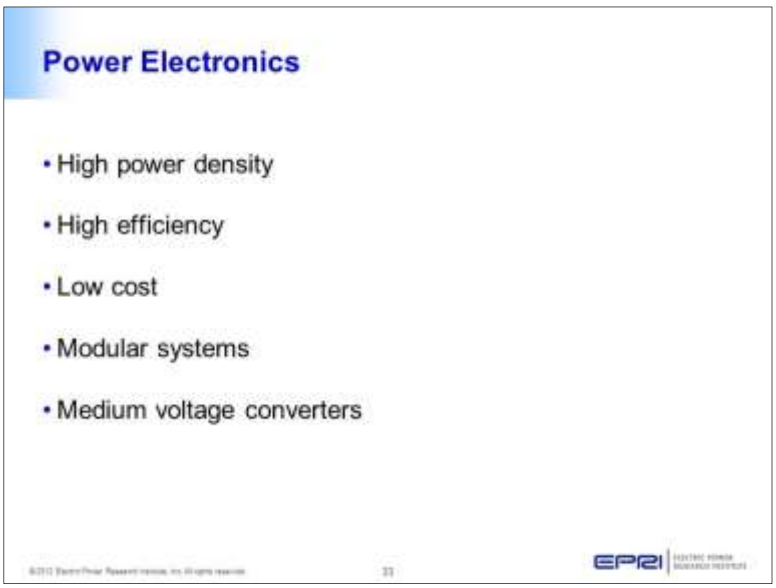

Power electronics, surprisingly, this is an area that needs more research. Right now, virtually zero amount of research is occurring. The main reason behind this is that all of the focus today, as far as utility storage, is in trying to cobble together these systems and putting it onto the field and seeing if they can actually solve any of the issues.

All of the focus is basically about deployments, not really on system integration and technology development. What we have been noticing is that, with all of these deployments happening, a lot of issues in power electronics needs to be improved. More research... the needs are there.

Dr. Sun was presenting negative impedance loads. I'm telling you, we are not even at that point. We are just at the point where we are trying to put the systems on their feet to see if we can even solve any of the problems that are out there.

We need high-power-density systems. This is really critical.

From what we are seeing, we need high-power-density systems. This is really critical, especially if you are looking at those community energy based systems and residential energy based systems. You need to minimize the size of bulk converter. If you look at a community energy storage system, a typical distribution transformer is about 3 x 3 x 3 feet, about 27 cubic feet. A distributed energy storage system is comprised of two boxes, each very similar to those distribution transformers, 3x3x3. You're going to put this on somebody's front yard? Would you be amenable to that? That is a fundamental question.

So you are going to have three boxes sitting on your front yard: one is for the battery, one is for the bulk converter, and one is for the transformer. As far as the battery in a small space, it is beyond the ability of this group, at least. I would say so because there are a lot of people working on this, trying to push back energy densities. Power converter densities can really be improved. There is no reason why you need to have all these big boxes out there for power converters.

The reason why we're not putting much thought into this is because right now the goal is to get all the systems and see what even works.

So high power density is critical, we need high efficiencies. Battery efficiencies are inherently low, so we need to maximize power converter in those efficiencies. Each and every loss element of the overall energy storage system needs to be minimized. Efficiency is a critical aspect.

50% of the energy storage system cost is the power converters.

The cost overall needs to come down. When we talk about an energy-storage systems, people think that the battery is the most expensive thing. If you look at a stationary utility storage system, it's 50% of your system. 50% of the system cost is the power converter; 50% is the battery.

It is a wrong notion that the battery is driving up the cost. When we look at reality, 50% of the cost (even more when you look at small-scale or medium-scale systems), in the power electronics would be about 60 or 70% of the cost for the converter.

As we get into distribution applications, medium-voltage power converters are another critical area we need to examine. If we look at many of the distribution applications today, they're all 480-volt-based systems. However, we can put in a 5 megawatt-scale system. You're still at a 480 volt system. You use the step-down transformer to basically step up to your needed voltage. That means you basically need a small substation to put in your storage system. You need a pad-mounted transformer to change the 480 volts to medium voltage. Anything we can do to minimize the size of that will be helpful.

Modular systems are very important.

Modular systems are very important. Anything we can do to create building blocks will work for a power converter standpoint as well as a battery standpoint. We can stack them to get any voltage and power that we want. That would be very helpful right now.

These are the critical elements in power electronics that we need to look at immediately. As we get more and more experience, we can look at all of those negative impedance issues and those other issues. But right now, as I said, all of our concentration is basically focused on system assembly and integration.

There are a lot of challenges. Primarily, the challenges arise because we're seeing how to quickly put together these systems. You have a battery system that knows only the battery, which is a chemical system. The battery vendor does not know where the power goes. The power electronics vendor has no idea how to manage these batteries; how to put in battery monitoring systems; or how to ensure safe charging. These are things that the power electronics community is not familiar with.

System integrators sometimes have no idea of power electronics or battery management systems.

People are trying to put together these systems. System integrators sometimes have no idea of either power electronics or battery management systems. They get into the business; they want to buy a converter from a power converter vendor; they want to buy a battery from a battery vendor. They want to cobble it together and they expect it to work. It doesn't.

There are a lot of problems, a lot of challenges, in system integration. This is the greatest barrier we are facing right now. This is attributable to a lot of issues: people don't know what they're doing; there are no standards in place; and no standardization. Each power converter and each battery management system will be talking their own language.

The electric utilities have no experience putting these systems on their feet.

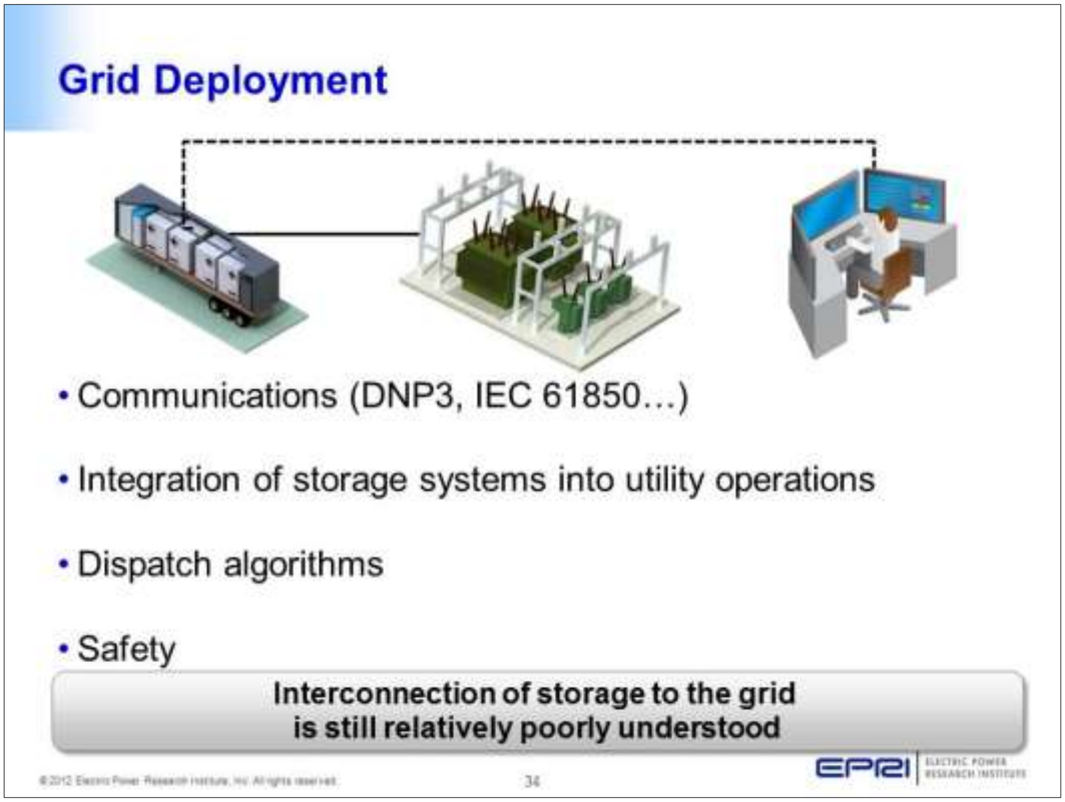

So all of the problems that I was talking about are just at the storage level. Putting it onto the grid, it's a story of errors. The electric utilities have no experience putting these systems on their feet. There are issues with installing the systems. You need to get permits that are really not written for energy storage systems. Commissioning the systems is new. How do the systems communicate with the utility grid in a secure way?

There's a lot of work going on in the field of communications right now. I define communication as being the battery and the utility systems. How do you integrate the storage systems and grid operations and distribution and management systems? What are the dispatch algorithms and control algorithms needed to control these batteries? This is still under research.

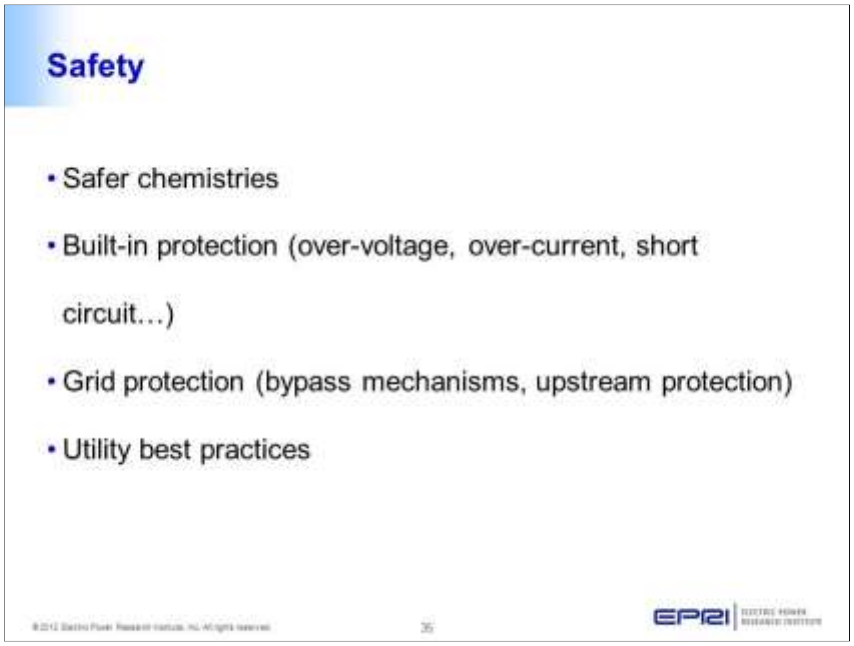

Safety is still a big concern. The foremost question that comes up from any utility is "Are these batteries safe? What would we do if there was a fire? Just run out? Throw water at it?" All of these questions are still unknown. All of these must be dealt with.

Safety. A lot of research is going into safer battery chemistries. Protection of power converters, especially if they short-circuit protections; external protections from the grid perspective; safety precautions from grid perspective, utility perspective… these are all under research, but more work is needed.

It's the power converter that appears to be the source of many of these fires.

While people may think that the batteries are inherently unsafe; in the case of the recent large-scale fire incidents that have been reported, many are leaning towards the power converter, not related to the battery. A really surprising thing. You would think that the batteries would be the main culprit. Surprisingly, it's the power converter that appears to be the source of many of these fires. So there needs to be a lot more failure mode analysis to figure out what's going on. Power converter reliability needs to be improved.

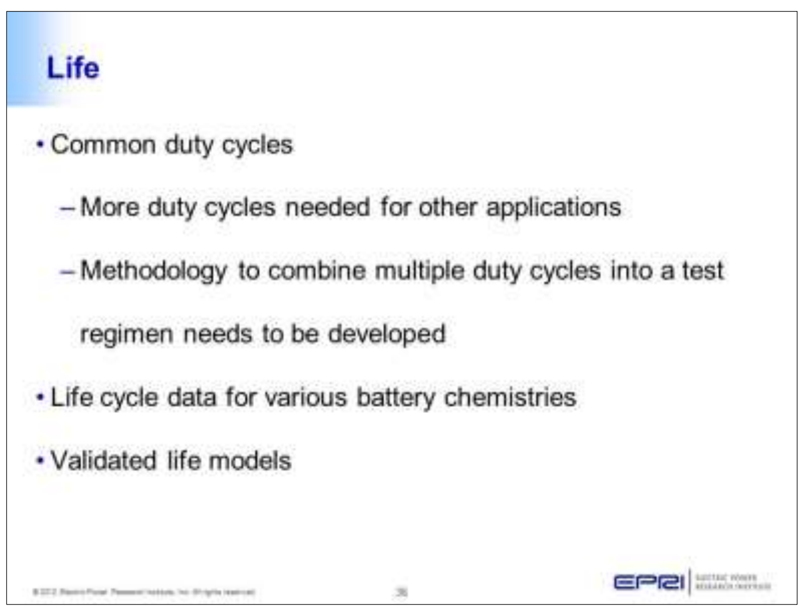

Life, as I said, is another critical factor. You need at least 15 to 20 years of life. Life of energy storage systems is a completely unknown quantity. If you ask me: "I'm putting this lithium-ion battery system for PV application. How long will it last?" I'll tell you I don't know. I cannot even answer. There are no duty cycles that have been developed that you can test these batteries through. Nobody has cycled these batteries or done any accelerated life tests through particular duty-cycles.

Nobody has cycled these batteries or done any accelerated life tests. Without data, how can you create life expectancy models to predict life?

So this is a big unknown. A lot of research we are now pulling in is toward development of these standardized duty cycles that represent different utility applications. How to obtain data, by cycling through accelerated life test on batteries, if you don't know those duty cycles. Without that data, how can you create life expectancy models that you can use to predict life? Those are the main areas of research we are developing to answer those questions.

The summary is that there are a lot of opportunities for energy storage and there are a lot of challenges. We need research in every aspect of energy storage. It is still open; the area is ripe for research. More development needs to be done. We are getting there; slowly, but steadily. We have had a lot of hiccups along the way, but it's promising technology with lots of applications. A lot of development work needs to be done.

With that, I complete my presentation.

Thank you.

Autonomous Response

DUSTY BECKER:

Our next presenter, Ed Herbert, received his Bachelor's Degree of Electrical Engineering from Yale University. He worked as a design engineer, a project engineer, an engineering supervisor, then as an engineering manager until 1985. Since then, he independently promotes patented technology for licensing.

Within PSMA, Ed is a member of the Board of Directors, Co-Chairman of the Magnetics Committee, and Co-Chairman of the Energy Efficiency Committee.

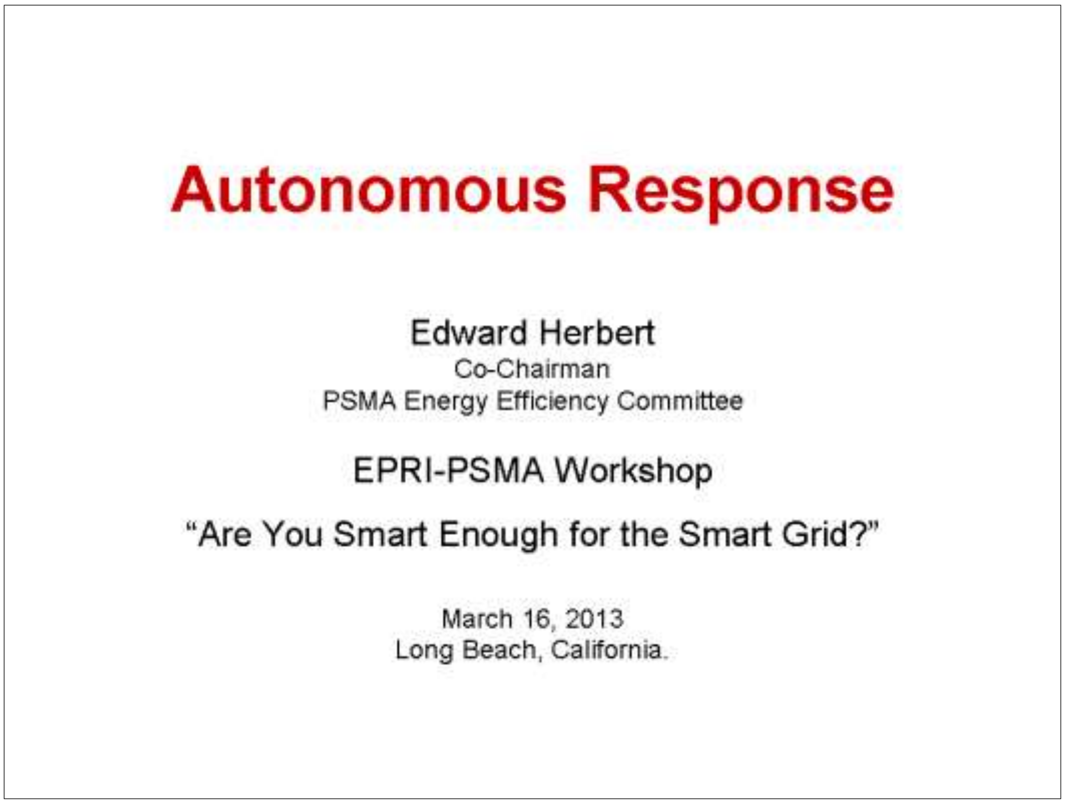

ED HERBERT:

Autonomous response has been mentioned in quite a few of the earlier talks. I think it is key to successful and reliable deployment to the grid. I'm approaching input impedance from a much simpler point of view than Dr. Sun did in his earlier presentation, because he knows how to do math and I don't. I approach it more from a nuts-and-bolts point of view.

> *"Autonomous response" is equipment responding to local conditions measured at its input, such as voltage, impedance, frequency and/or phase angle.*

- **Definition of "Autonomous Response:"**
  - "Autonomous Response" is the response of a power converter (or other load) to local conditions, especially line voltage V, but also dV/dt, phase angle φ and/or dφ/dt and other inputs. Autonomous response does not use a data link.
  - "Autonomous Response" is contrasted to "Commanded-Response" — which relies upon a data link — as may be used in a "Smart Grid" environment.
  - "Autonomous Response" is enabled by modifying the input characteristics of the power converter (or other load) to meet specific criteria under specified local input conditions.

Fig. 1

The definition of "autonomous response" basically is equipment responding to local conditions that can be measured at the input. By response I mean, largely, input current control. It is contrasted to command response, which requires a data link. Autonomous response usually is enabled by modifying the input characteristics of the power converter or other similar load.

A well-implemented autonomous response may make a data-linked demand-response system unnecessary; along with all its firmware and software costs and its security, warranty, and liability concerns.

Not having a data-link, autonomous response is immune to hacking.

When would you use autonomous response?

1. If you build equipment that requires a data link and someone says, “Well, what happens if the data link fails? What should the equipment do?” You might need a default mode that does not rely on the data link and still keeps working in a stable, reliable way.
2. For security: You might have equipment that absolutely must be immune to hacking. You can’t have a data link coming in no matter what.
3. For stability: Modifying the input response may provide better system damping.
4. For economy: A well-implemented autonomous response may make a demand-response system unnecessary, with its firmware and software costs and its security, warranty, and liability concerns.

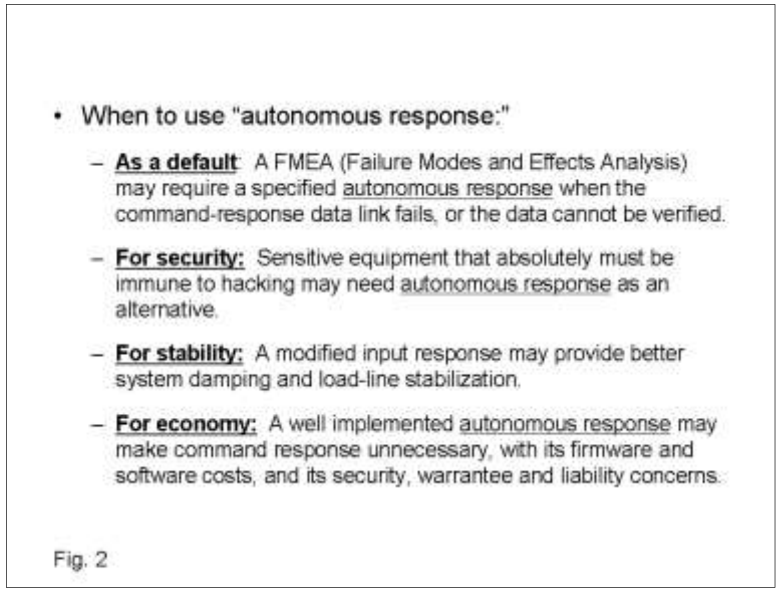

Fig. 2

We have been using autonomous response in power converters for years – it is known as power factor correction (PFC).

Modifying the input characteristics of power converters requires input current control, as shown in Figure 3. Before anyone says “Oh, that’s a great deal of added circuitry... We can’t do that!” We’ve been doing it for years. It’s called power-factor correction, or PFC.

Autonomous response is basically power-factor correction on steroids.

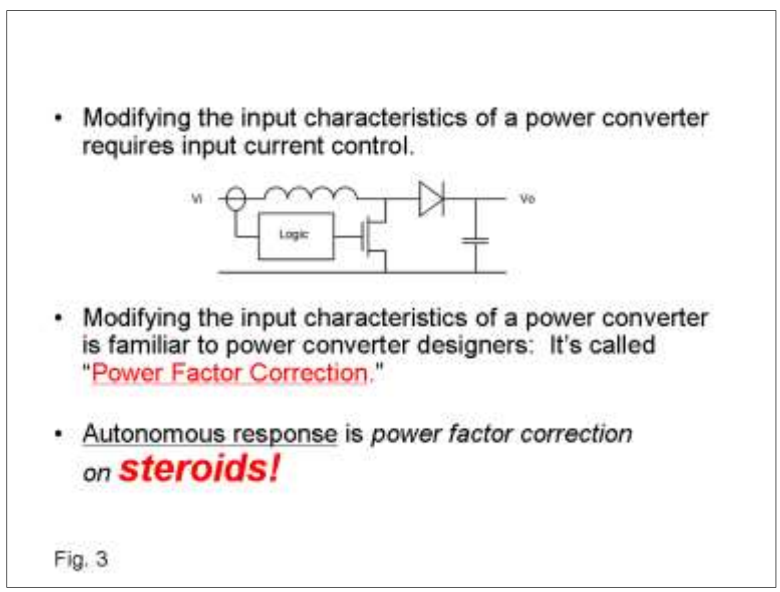

Fig. 3

Autonomous response is basically power-factor correction on steroids!

- **Modified input characteristics for autonomous response:**
  - There is great flexibility in modifying the input characteristics of a power converter.
  - **Any algorithm** that can be implemented in logic or a DSP can be applied to the input current control.
  - Both the scaling of the input current and its dynamic response can be controlled by the algorithm.
  - Designing the logic to control the gain and dynamic response of an algorithm is second nature to a power converter designer, using techniques adapted from feedback control.

Fig. 4

There's great flexibility in modifying input characteristics of a power converter. Any algorithm that can be implemented in logic or a DSP can be applied to the input current control. Both the scaling of the input current and its dynamic response can be controlled by the algorithm. Designing the logic to control gain and dynamic response in an algorithm is second nature to anyone who has done feedback control for a power supply.

- Autonomous response possibilities
  - Power factor correction—common place now
    - makes the input "look resistive" at two times line frequency.
  - Load-line modification
    - Improved conservation voltage reduction
    - Improved system damping.
  - Soft start and soft stop
    - Limit input di/dt.
  - With a bi-directional front end,
    - VARs generation
    - Volt/VAR optimization (VVO).

Fig. 5

Autonomous response possibilities are summarized in Figure 5. Obviously, it can be for power-factor correction; we've been doing that for years. Load-line modification can greatly improve the ability to use conservation voltage reduction. That's where dropping the voltage drops the power.

It can improve system damping.

If you have input current control and you have sufficient energy storage, you can soft-start and soft-stop.

We heard a talk earlier about perturbations on the system. If a very large load, like a ten-megawatt data center, were required to ramp up its input current over many seconds and then, when it goes offline, it ramps down the current over many seconds; it might be very helpful to grid stability. This should be no problem at all with the large energy storage that most data centers have.

If there is a bi-directional front-end, we can probably do volt/VAR optimization.

Autonomous response in not just for loads; distributed power sources (PV, wind) can be analyzed as negative loads.

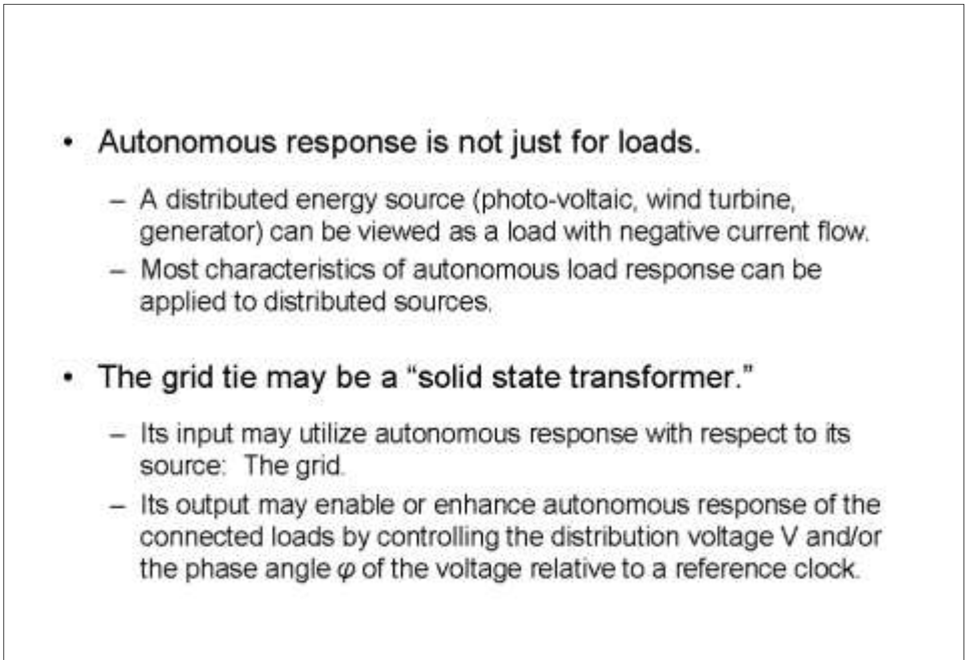

Fig. 6

Autonomous response is not just for loads. The small distributed generators, like PV or wind turbines, can be regarded as a load with negative power. Just about any characteristic of autonomous load response can be applied to distributed sources.

Anything that regulates the voltage on a power line, irrespective of what the grid is doing, converts the entire downstream system to a constant power load, with all of the problems that this presents.

Grid tie may be a "solid state" transformer. This is interesting, because its output may enable enhanced autonomous response to its loads. Its input may run autonomously with respect to its source, which is the grid.

However, another significance of a solid-state transformer, or anything that regulates voltage on the line, is that if the output of a transformer regulates the voltage, irrespective of what the grid is doing, it has converted the entire downstream system to a constant power load, with all the problems that this presents.

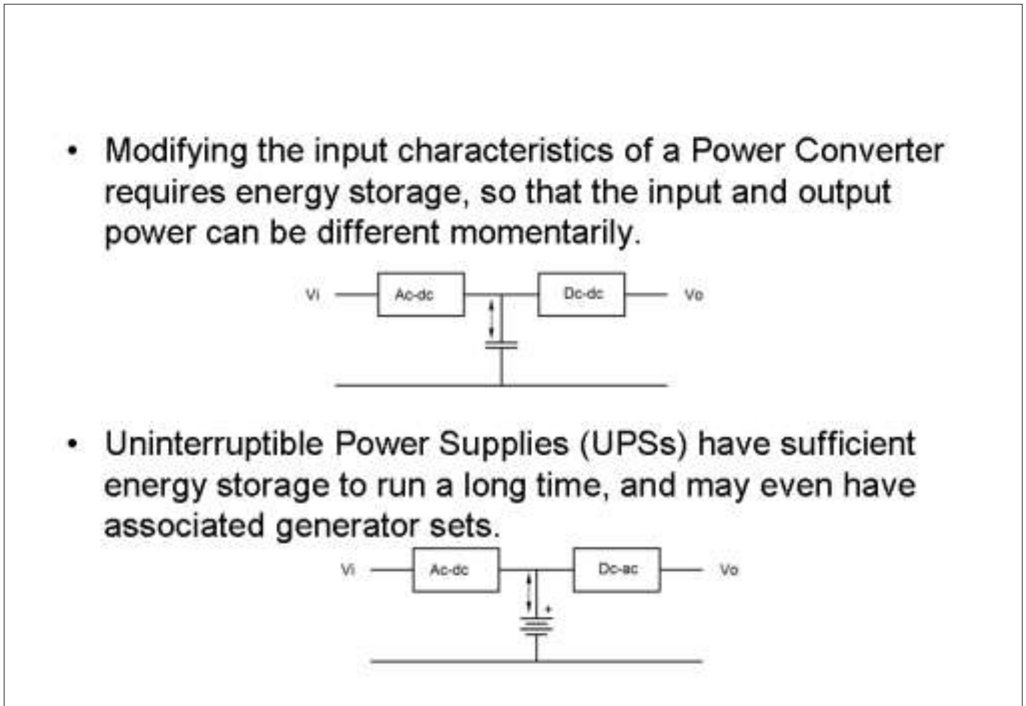

Fig. 7

Modifying the input characteristics of a power converter requires energy storage. This is basically a block diagram of a PFC control. The uninterruptable power supplies have very large energy storage. They may even have associated generator sets that might run days without anything from the input.

Energy storage is necessary whenever the input power and output power differ.

Energy storage is necessary whenever the input power and output power differ. Modifying the input power or input current, while still satisfying the requirements of the load, require that difference. If you draw a graph like that and take the area under the curve; that's the amount of storage that you need just for the input compensation.

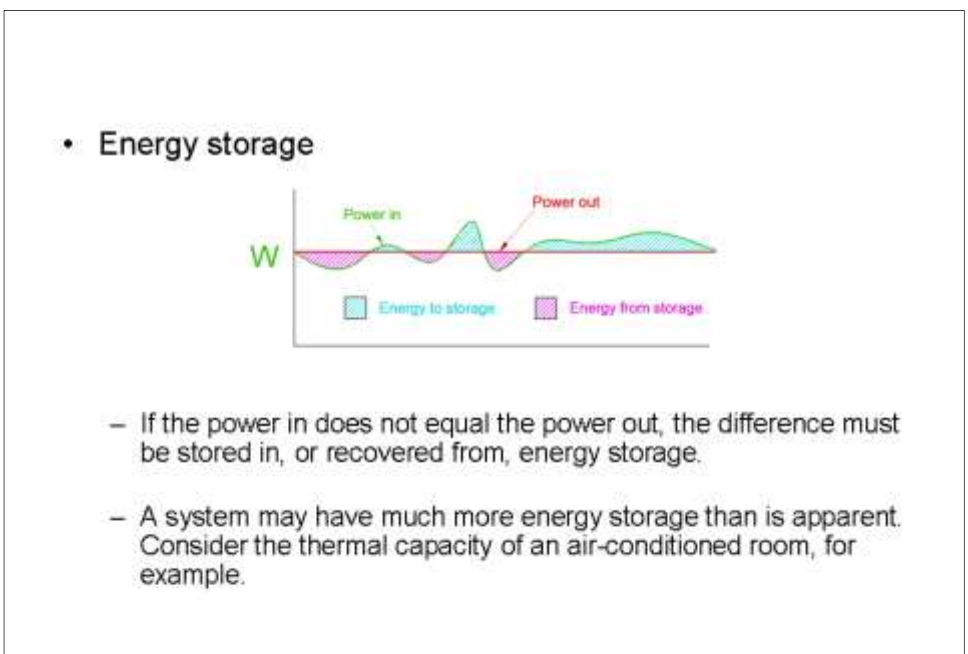

Fig. 8

The system may have much more energy storage than is apparent. When you're talking about heating and air conditioning loads, the thermal capacity of the environment has tremendous energy storage and it allows you to interrupt the load for long periods or even increase the load for brief periods. If you have ice storage, you increase that many orders of magnitude.

The effect of constant-power loads is to make conservation voltage reduction ineffective. CVRf→0.

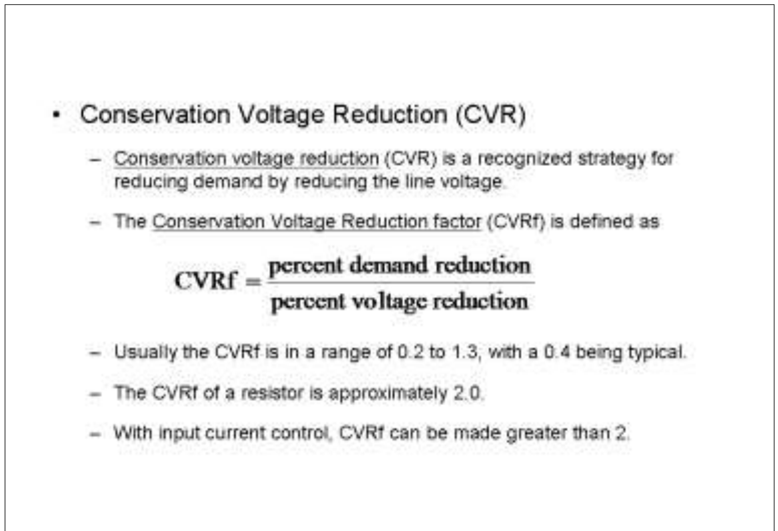

Fig. 9

Conservation voltage reduction has been covered in earlier presentations, so I'll go very quickly. It's the percent power demand divided by the percent voltage reduction. Usually it is in the range of 0.2 to 1.3, with 0.4 being typical. Earlier slides presented by other speakers said that 0.6 was typical.

The CVRf of a resistor is approximately 2. Within input current control, it can be greater than 2. I was wondering about the Hydro-Quebec example mentioned by another speaker, with the low-CVR ratio being blamed on resistive loads. I don't think that can be right because resistive loads are approximately 2. Maybe they're regulating branch circuits and converting these resistive loads into constant power loads.

[**Note**: On further consideration, I realized that heater loads have thermostats that control the power. It's the thermostat control that makes the resistive load into a constant power load when averaged over time.]

The effect of constant-power loads is to make conservation voltage reduction much less effective. In fact, it's totally ineffective with a constant-power load. I mentioned that CVR is sometimes combined with volt/VAR compensation. Autonomous response can include volt/VAR optimization if there is a bi-directional input.

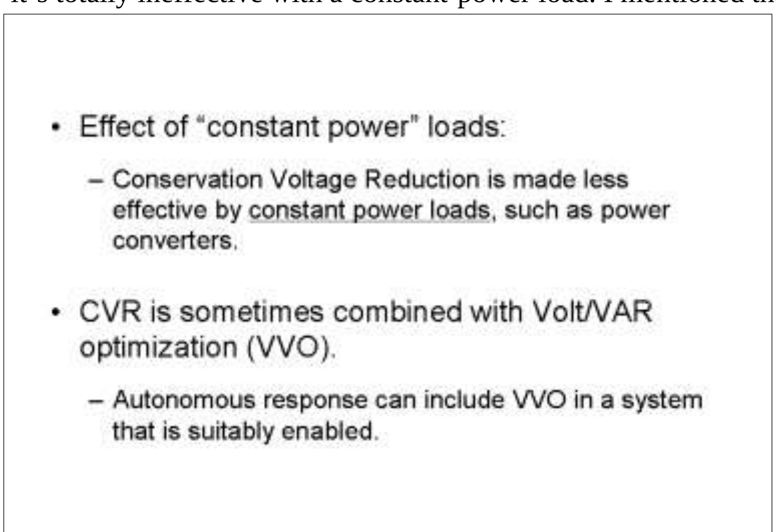

Fig. 10

As the conservation voltage reduction factor CVRf→0:

$$\Delta V = \frac{\Delta P}{CVRf}$$

$\Delta V \rightarrow \infty$.

The voltage becomes unstable.

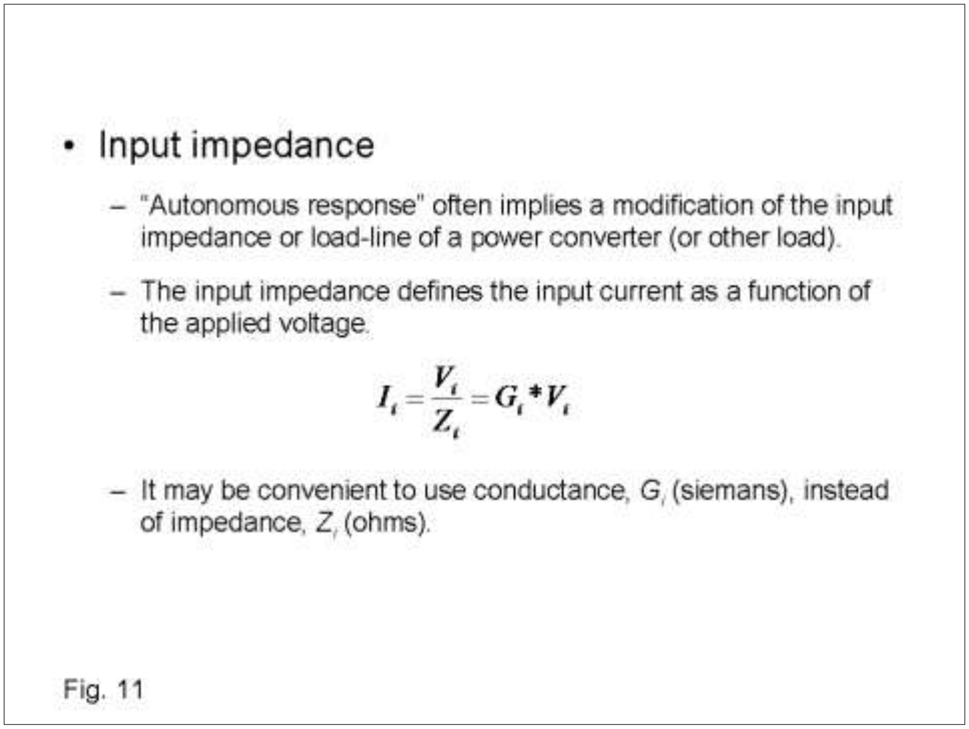

Fig. 11

Input impedance [corrected]: I define input impedance much more simply than Dr. Sun did in his presentation. Basically, the input impedance, Z_i, is the input voltage, V, divided by the input current, I, or:

$$Z_i = V / I .$$

For 3-phase:

$$V = \sqrt{V_a^2 + V_b^2 + V_c^2} \quad \text{and} \quad I = \sqrt{I_a^2 + I_b^2 + I_c^2} .$$

The dynamic input impedance is:

$$Z_i = \frac{\Delta V}{\Delta I}$$

It may be convenient to use conductance G_i instead of impedance Z_i:

$$G_i = \frac{1}{Z_i}$$

The dynamic impedance, Z_i, of most power converters is negative, but it can be modified to be positive over a dynamic range.

- The dynamic input impedance of most power converters and other constant-power loads is negative.
  - This means that if the voltage droops, the input current rises, and *vice-versa*.

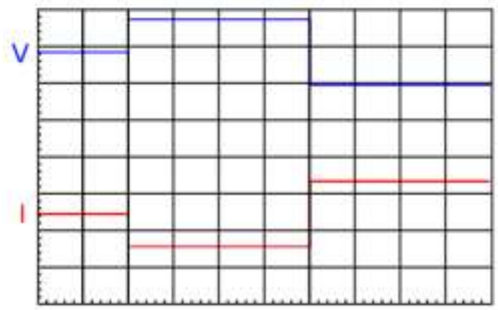

  - This is NOT beneficial to a grid that is under stress, and tends to defeat CVR.

Fig. 12

The dynamic input impedance of most power converters and other constant-power loads is negative. This means that if the voltage droops, the current rises and vice versa. This is not beneficial to a grid that is under stress and tends to defeat conservation voltage reduction.

The input impedance of a power converter usually is constant power, but it can be modified to be resistive. In Figure 13, I show that negative resistance tends to follow a constant power curve. Positive resistance, ideally, is a slope — a line going through the origin, increasing at a linear rate. Light bulbs are not linear. Many loads aren't, but that's the definition.

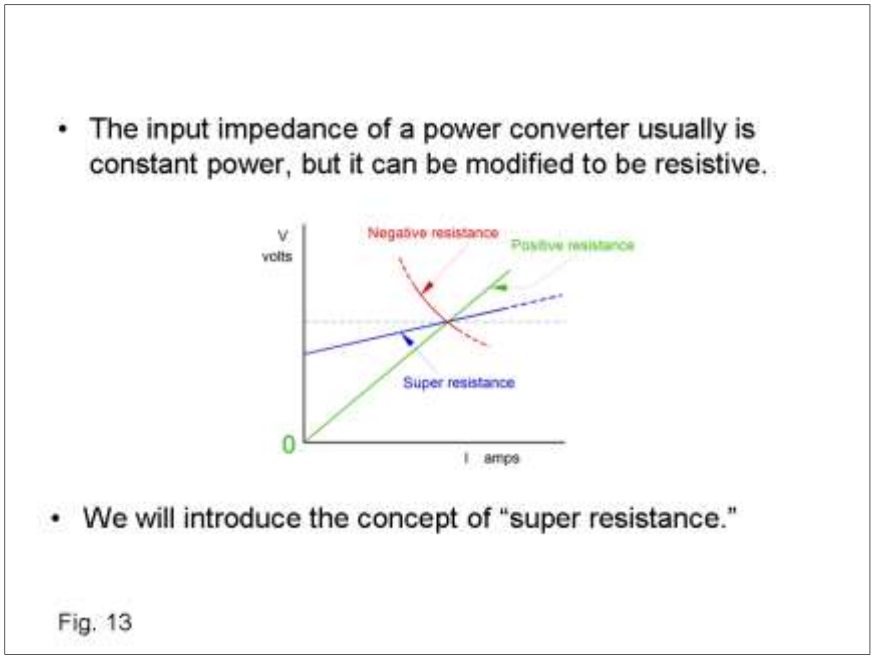

Fig. 13

With negative dynamic impedance, when the voltage goes down, the current goes up.

Figure 14 shows a bit more about the dynamic resistance as the slope of the resistance curves at the operating point. A true voltage source has zero resistance. It produces the voltage regardless of the load. A current source provides a current, regardless of the voltage. The super-resistor has a lower slope.

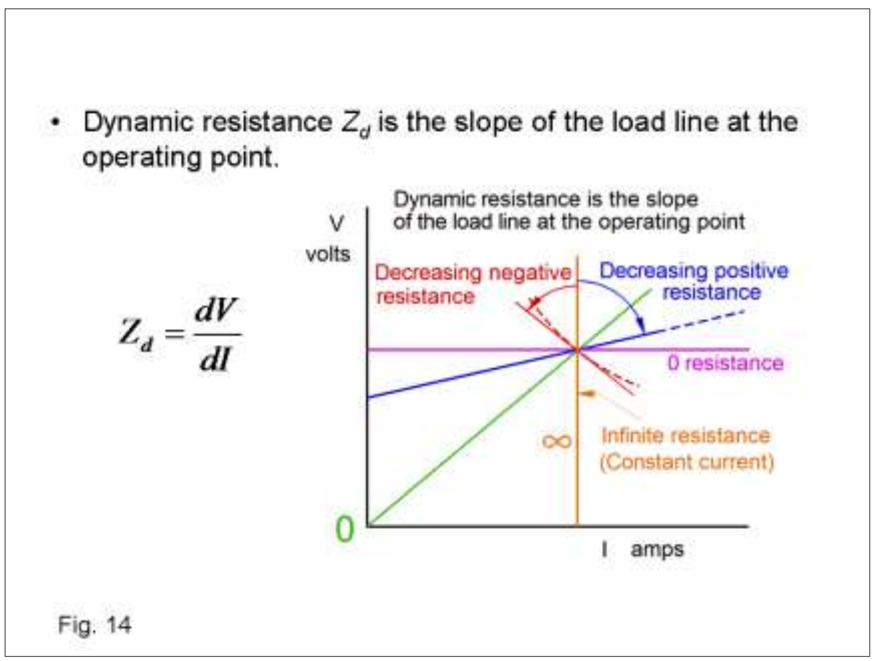

Fig. 14

"Constant power" does not mean that the input power does not change. It means that the input power does not change in response to input voltage.

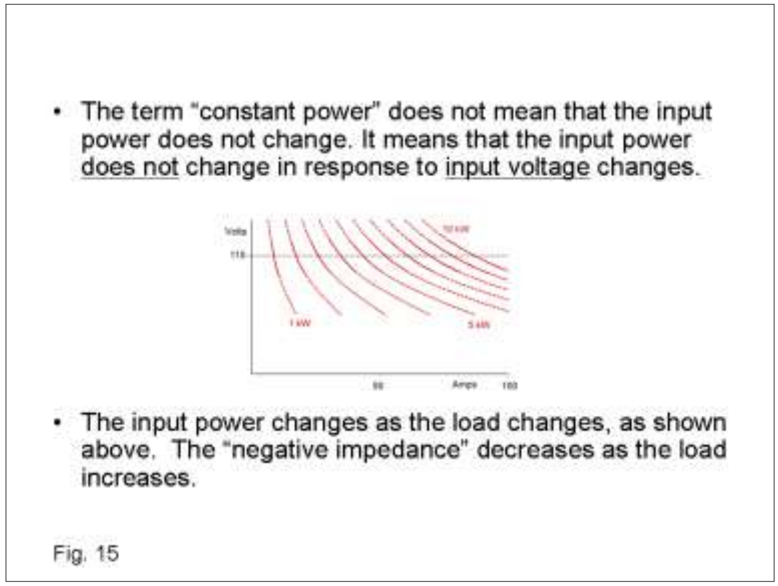

Fig. 15

The term constant power does not mean that the input power does not change. It means that the input power does not change in response to input voltage. Input power does change as the load changes, as shown above. The value of the negative impedance decreases as the load increases.

There are anecdotal stories about data centers that operate fine at low load; but once they increase the power, they are unstable. This is probably the reason.

Figure 16 shows again, graphically, constant power load. The voltage goes down and current goes up, and vice versa. The output power is steady and conservation voltage reduction is defeated.

Using "super resistance," the CVRf can be > 2.

With the resistor, the input current and the input voltage track proportionately. However, the output, the input power, and the output power differ, as shown in the lower curve. You need energy storage to supply that difference.

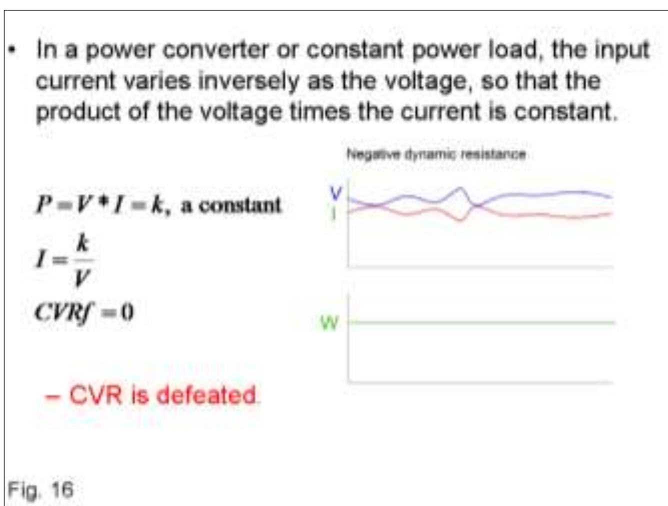

Fig. 16

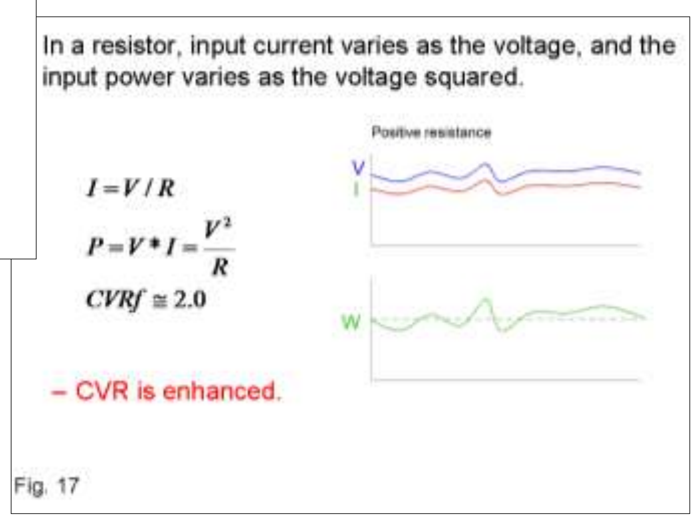

Fig. 17

Using "super resistance" makes the change in current greater than proportional to the voltage, so the CVRf factor can be greater than 2. Dynamic resistance is lower and the load line does not go through the origin. This is not something the power converter people normally would do, but it is something they might be willing to do if you paid them to do it. A subsidy in the rate for improved CVRf may be good policy.

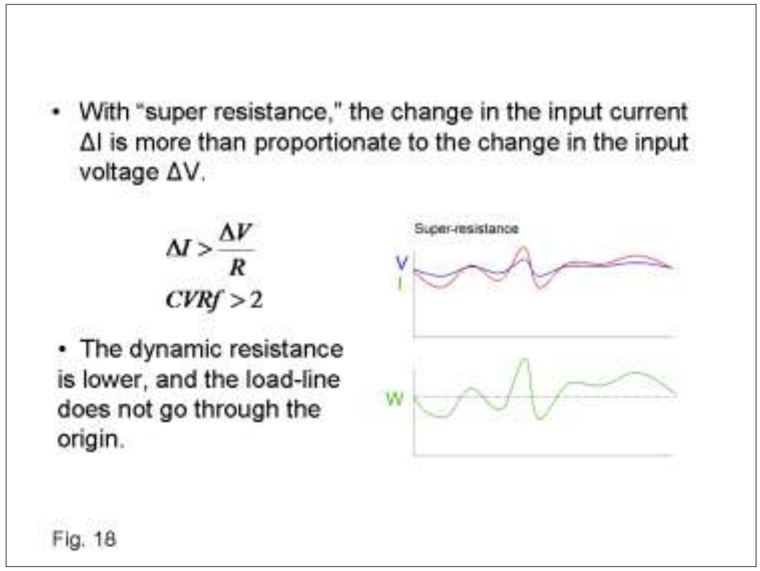

Fig. 18

A lot of the impetus for this workshop was to look for system-wide optimization. I felt that there were probably things power supplies could do that would be helpful to the utility, but they did not know enough to ask. Power converters can do a lot of things for the utility that they don't know enough to offer. If we get these two groups together and discuss problems and abilities, maybe we'll find some solutions.

A subsidized rate for improved CVRf may be good policy.

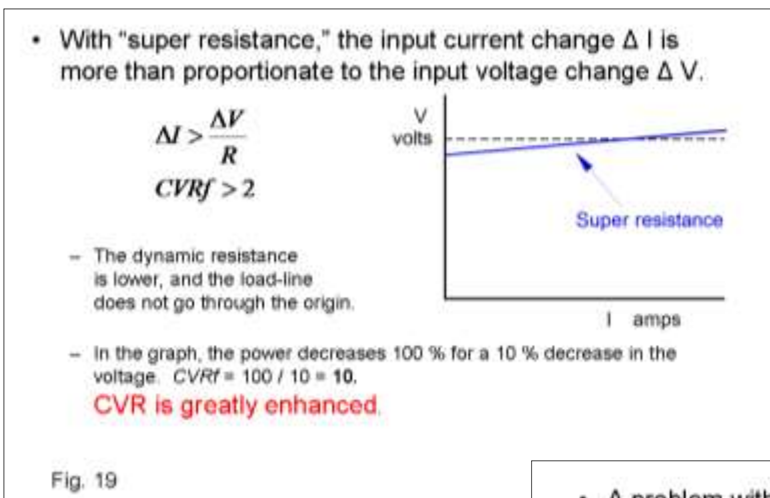

Fig. 19

Figure 19 shows a graph of super resistance. In this case, we have a load that decreases the power 100% with a 10% drop in voltage. The conservation voltage reduction is greatly enhanced. This could offset many other constant-power loads. It might be a suitable response for air-conditioning. If the voltage gets down to a 10% drop, it darned well better be disconnected.

In Figure 19, the current has a very high peak current with a voltage surge. You might want to modify a super resistance to be a regular, normal resistance for a voltage surge, with a break in its characteristics, as shown in Figure 20.

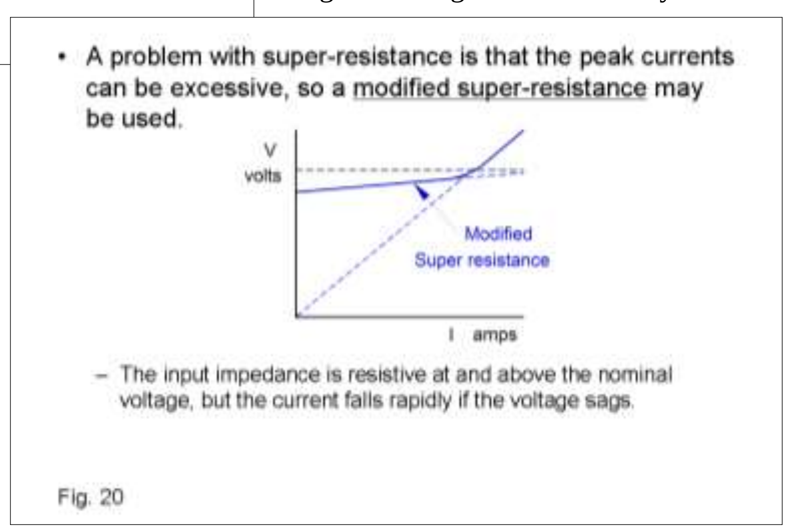

Fig. 20

This is as much to remind people that any algorithm that can be put into the current mode control can be implemented.

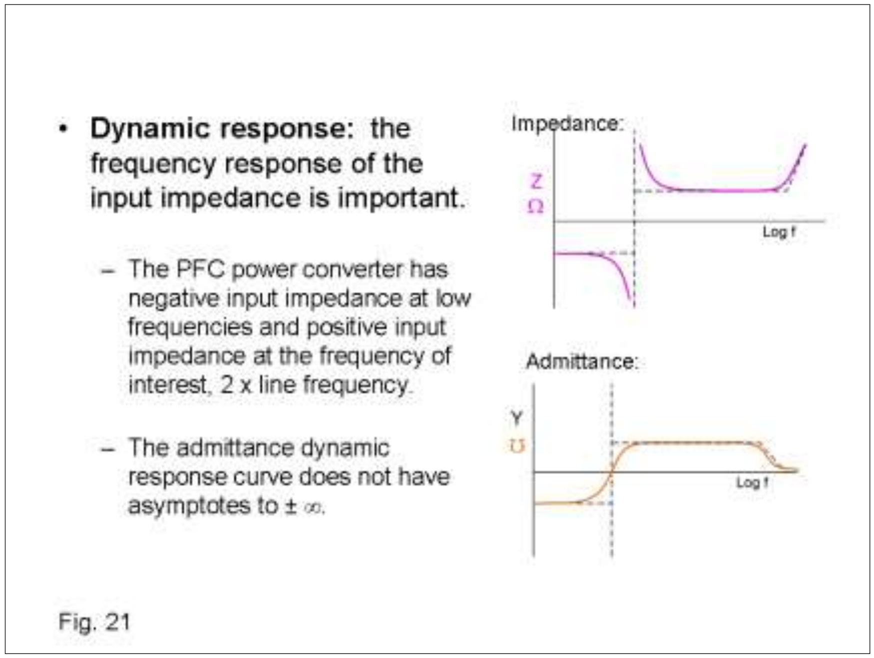

Fig. 21

Dynamic response, the frequency response of the input impedance, is very important. As shown in Figure 21, a PFC converter has negative impedance at low frequencies, as it must. The current goes down as the voltage goes up and vice versa. However, at the frequency of interest, which is generally two times line frequency, it has a positive resistance. In transitioning from negative to positive resistance, the input impedance cannot go through zero or you have a short circuit. So it goes through negative and positive infinity. For that reason, I have a slight preference for admittance curves. However, I think most people are used to impedance curves, so I'm not sure which will gain traction.

The dynamic response of the input impedance is very important. A PFC power converter has negative input impedance at very low frequency, as it must for energy balance, but it looks like positive resistance at two times line frequency.

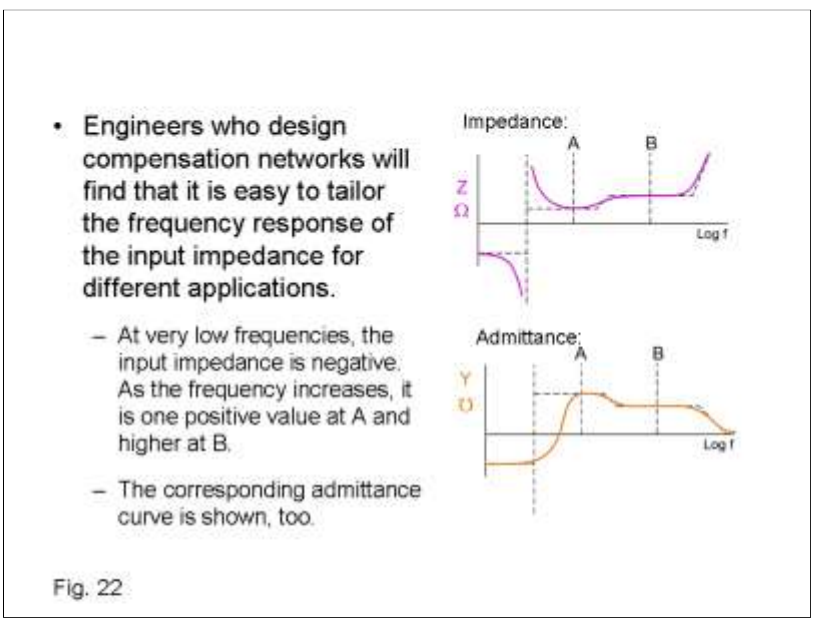

Fig. 22

Anyone who handles feedback and has to design a compensation network for a certain frequency response profile can apply that to the dynamic characteristics of an input current controller. This is an example where, maybe at Point B, we're doing power factor control. Maybe at Point A, we're going to offer a much lower impedance so we can help damp out the system resonance that might be the grid resonance or maybe substation resonance or something. Anyway, the point is to tailor the frequency response for different applications.

To maintain a positive input admittance to 0 Hz (dc), there must be a supplemental power source, such as an auxiliary generator.

To maintain positive admittance to dc, as shown in Figure 23, there has to be a second power source, such as a very large battery or a standby generator. You cannot have positive admittance, positive resistance, to zero with just capacitors unless they're unusually large.

- Load shedding:
  - To maintain positive admittance to dc, there has to be a second source of energy, such as a very large battery (UPS) or a standby generator.

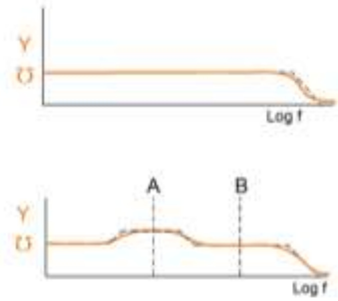

Fig. 23

Anyone who handles feedback and has to design a compensation network for a certain frequency response profile can apply that to the dynamic characteristics of an input current controller.

Air conditioners with variable-speed motor controls may have a CVRf of 0 (constant power). It is easy to convert an air conditioner or space heater to have resistive or super-resistive characteristics; just vary the thermostat set point as a function of the input voltage.

- It is particularly easy to make the input impedance of heaters and chillers look resistive or super-resistive: Just vary the temperature set point as a function of the voltage.

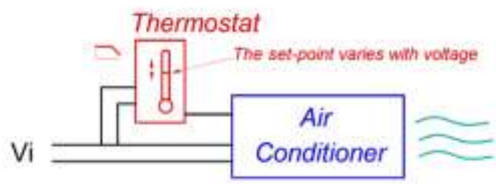

- With an air-conditioner, if the voltage goes down, set the thermostat higher. That will reduce its input current.

Fig. 24

It's particularly easy to make the input impedance of heaters and chillers super resistant. Tie the voltage to the thermostat and alter the set point as the voltage. This could either be a steady-state or a dynamic response… perhaps a dynamic response with a 15-second roll off.

With the air conditioner, a voltage increase could cause the thermostat set-point to go higher, which reduces input current. If the voltage goes down 10%, shut it off entirely, as shown in the curve in Figure 19 or Figure 20.

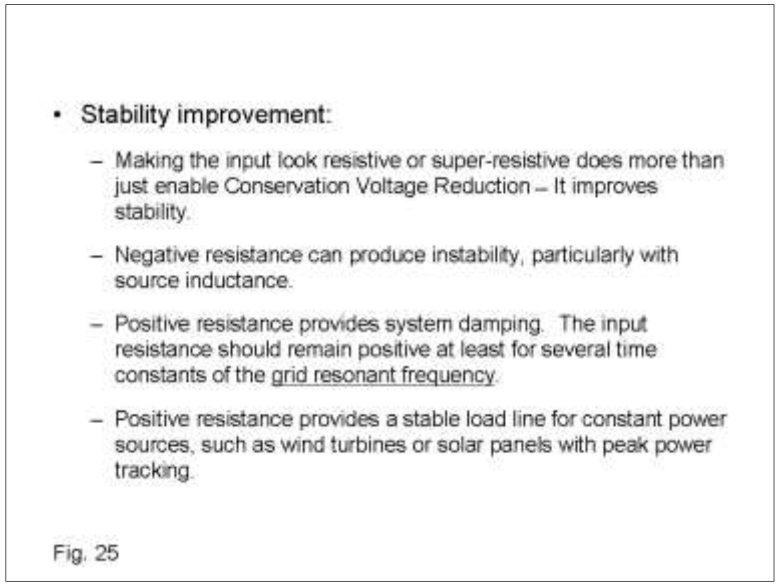

Fig. 25

Stability improvement. Making the input load resistive does more than just enable conservation voltage reduction; it improves stability. Anyone who studied electrical engineering in college probably learned about negative resistance oscillators. A lot has been written about it, so I won't rehash it now.

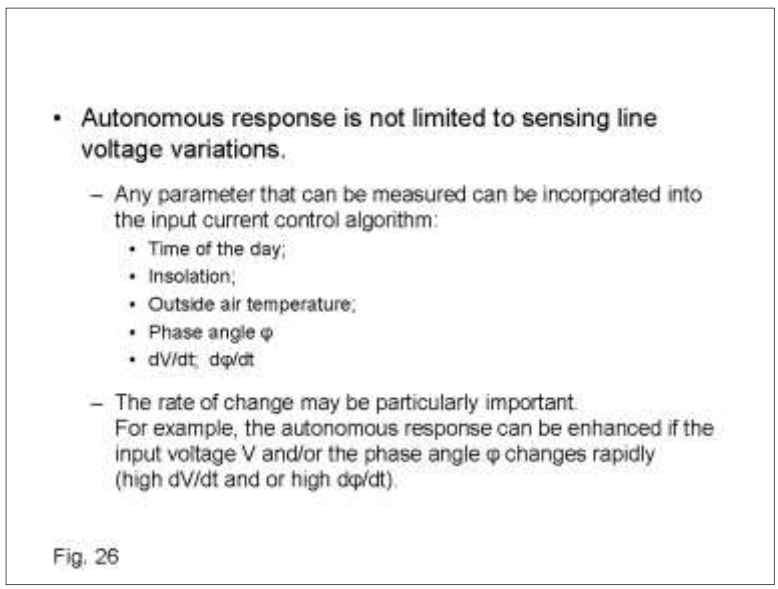

Fig. 26

Autonomous response is not limited to sensing voltage variations.

Any parameter can be measured and incorporated into the algorithm, such as:

- Time of day
- Amount of sunlight,
- Outside air temperature
- Phase angle of the voltage
- Change of the voltage
- Rate of change of the voltage, dV/dt
- Rate of change of the phase angle of the voltage, $d\varphi$.

The rate of change of the voltage, dV/dt, may be particularly important for dynamic response.

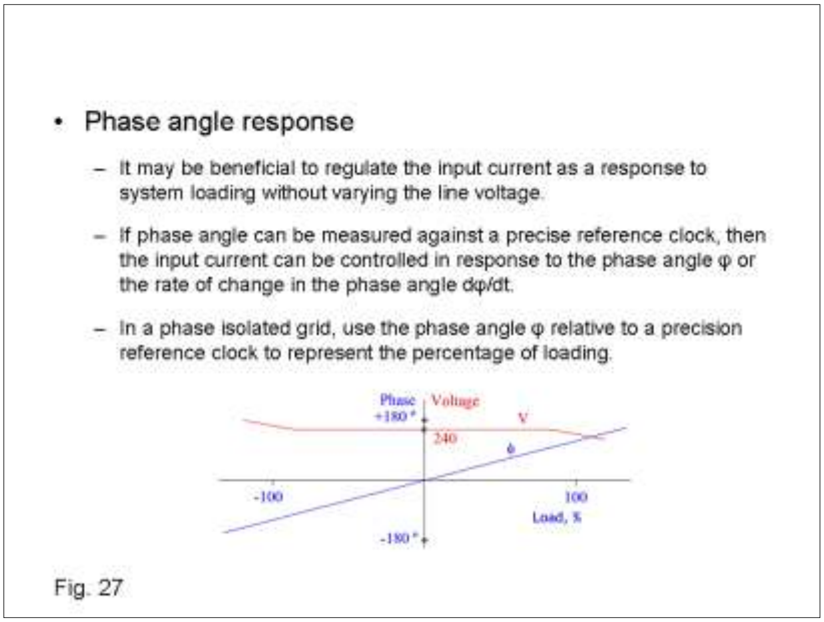

Fig. 27

Figure 27 shows is something I thought of as I was developing these slides in the past couple of weeks. I was trying to think of some way, using autonomous response, of getting information to the load about how much the grid is loaded. By this, I mean the local grid. Voltage is probably not a good indicator because, if the voltage droops, the grid is going to turn on the regulators and boost it back up. You might not know that the grid is loaded to the 90% capacity.

Phase-angle response may be a good method of conveying the status of grid loading.

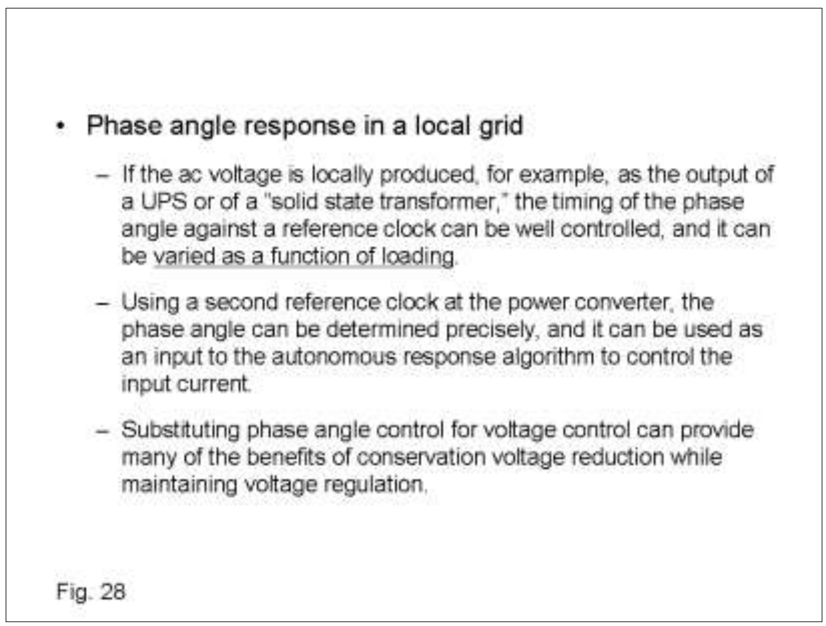

Fig. 28

Some systems measure line frequency as an indication of grid loading. Just as with the input voltage, it is likely that the frequency regulation will be corrected, so that a grid in near overload conditions may have nominal frequency. This becomes even more likely as community storage systems play a larger role in frequency regulation.

However, it occurred to me that on any system with phase isolation, such as the output of a solid-state transformer or a UPS, there is complete discretion about where to set the phase of the voltage. It can be referenced to a precision clock. A precision reference at the load can measure the phase angle of the voltage, referenced to the master clock. In this way, the controller at the load can tell very easily the percentage of the grid loading taking place.

[Added] If the grid becomes synchronous, locked to master oscillators, the frequency will no longer change with grid loading, but the phase-angle relative to the master oscillators will be a good indication of grid loading and virtually unhackable.

Phase-angle measured against a precision clock can be the principle way of conveying information about grid loading to the controller of a load. It is virtually unhackable.

Using phase angle measured against a precision clock can be the principle way of conveying information about the loading through a very large percentage of load change. This allows you to leave the voltage completely flat in the middle. The load still has information about the grid loading. Detecting this, some non-critical functions might start reducing their power at 50% loading. Life-support equipment might ride it all the way to the end, with variations in the middle.

The phase angle response is in a local grid. If the ac voltage is locally produced, for example, as the output of a UPS or a solid-state transformer, the timing of the phase angle against a reference clock can be well controlled and it can be varied as a function of loading. Using a second reference clock at the power converter, the phase angle can be determined precisely and used as an input to determine its response — to control the input current. Substituting phase angle for voltage control can provide many benefits in conservation voltage reduction while maintaining voltage regulation.

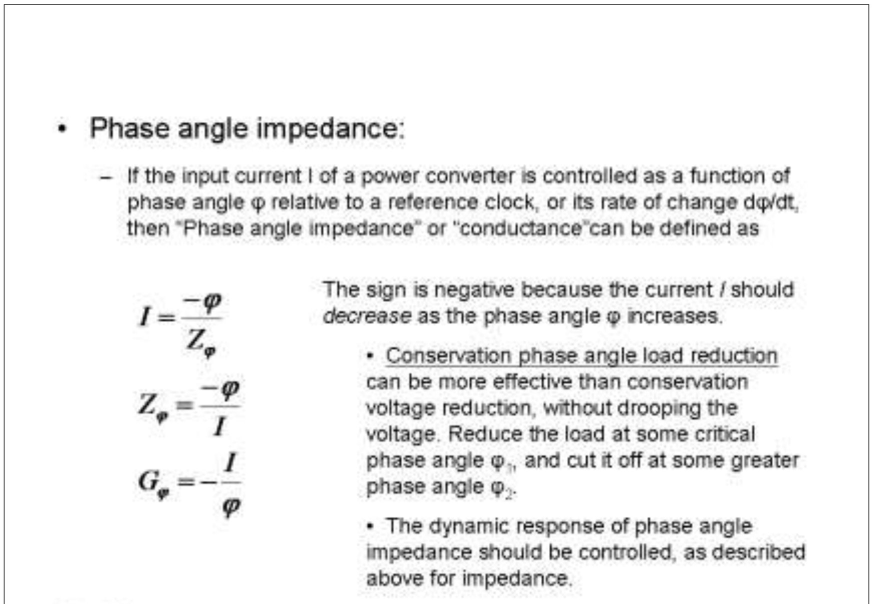

Fig. 29

I'm not a super analyst, so I don't really know if this is true; but I'm going to say that anything you can do with voltage, you can do with phase angle with such an algorithm. This is as much an illustration of what you can put in an algorithm — any parameter you can dream up, you can put into a control algorithm without limitation. Whether this is practical or not, I'm not going to judge.

Figure 29 defines phase-angle impedance. It's the control of the input current using the phase angle, not the voltage. I don't think you're going to start dimming the lights just because the phase angle got a little negative. But it will probably have some kind of step response -- ignore it up to some angle -- and then slope it at a steeper angle. You'll eventually get a constant current.

I think all of the relationships determined by Ohm's law can also apply to this control algorithm.

Thank you very much.

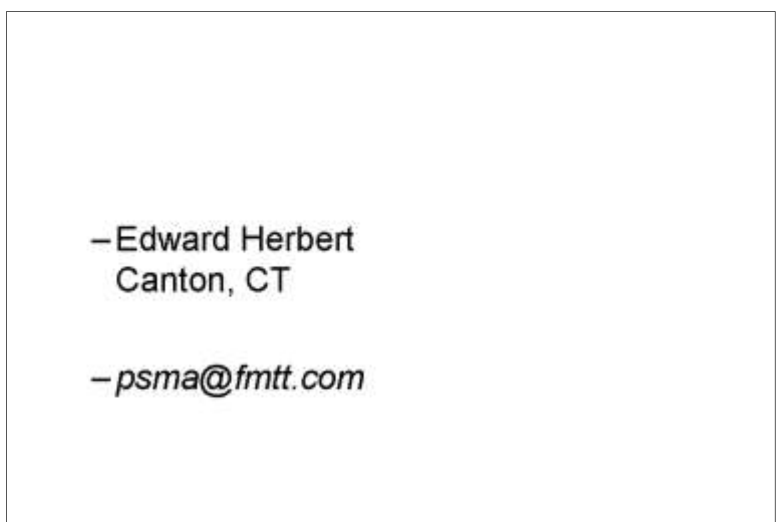

Fig. 30

Section VII
Third Panel Discussion

Third Panel Discussion — Jian Sun, Ed Herbert, Satish Rajagopalan, Greg Smedley

DR. JIAN SUN:

I just want to make a quick comment. Ed showed a lot of insight. I said in my presentation that impedance is a really unique concept. There's a lot of utilities, but there's also more likely to be misunderstandings. A lot of these things we cover today, imaginary power behavior, area-constant power behavior; define the impedance as being the ratio between the voltage and current average over cycle. It is no longer impedance. So you say that's a ratio between volt and current, but it's actually not impedance.

A really important characteristic of impedance is you put two of those in series, right? Total impedance is a sum of the two. You apply those rules and look at the quality that you're calculating. It's no longer impedance. I just want to put that out as a caution.

The dc system is very clear, whatever makes sense. When we talk about ac and area-constant power, it is a really tricky concept. I just want to add that; we can talk more afterward.

ED HERBERT:

If you are talking about the power that's consumed by a load, you can define your "impedance" or "resistance" (or give it a new name) as the power divided by the voltage. [Note: It should be the voltage squared divided by the power] That was how I used it. Maybe saying "real resistance" would be clearer than saying "impedance".

QUESTION:

Concerning storage, one of the observations I'm having is the grid is speeding up. I look at what's happening with bulk management and I'm concerned that some of the smaller storage systems we have, when you're coming off an inverter, the response time of the inverter can virtually add to the system. It can be very fast. If a fault occurs, especially a localized fault, is there a lot study going on in looking at that situation? Bulk management and bulk loading systems?

DR. SATISH RAJAGOPALAN:

I would say not right now. Nobody's even giving a thought to them. It's like you have these systems that are very fast because of pumps; bulk management is a challenging issue.

At some point I'm pretty sure we'll have a look at it; but right now, all of the focus is on just trying to get these storage systems into the field and seeing if they even work or they are suitable for an application. Pretty much all these are headed in that direction right now.

DR. GREG SMEDLEY:

I just have a quick question for that question. What aspect of the fault are you concerned about?

QUESTION:

A softball could probably cause a lot of oscillation issues. Especially if you're not able to have a tighter-controlled system that it could detect.

DR. GREG SMEDLEY:

I think what you're referring to is power converters typically phase launched to the grid with a voltage output. This can be avoided if the power converters actually provide pure current.

QUESTION:

When we talk about autonomous response, are we concerned that… The idea is very good that you have the distributed control. The more control variables you have, the better it is. We can resolve the problem by distributing it among all of these control variables now.

But is there a need for, like in your example, about the reflex? Now it's not the movement that is controlled by this. Yes, it didn't go to the brain, but it went to a central location or central control that controlled the hand. So I move a finger on my hand, I don't want every finger to go in that one direction. So what does that make it?

DR. GREG SMEDLEY:

There need to be rules developed that could be used for what the voltage profile should be on a given feeder. Then that voltage profile can be broadcast to the sub-control nodes and they would control out to that set point. You might have certain rules. Right now, there are rules for how do you set a cap bang or a load tap changer.

QUESTION:

But we only run into problems sometimes with a conflict in control. We have to be very careful that the same point is not being controlled by different…

DR. GREG SMEDLEY:

Okay. That is true. Good point. When you have distributed controllers, all operating autonomously, you need to make sure there is enough separation between them to ensure you don't have this kind of problem.

ED HERBERT:

Part of the question is going to be answered by Conor Quinn in his presentation. I have seen his slides, so I had a preview.

On the phase angle control I was proposing; that would have one input parameter, and that is the loading. That basically assumes you have a single source like a solid-state transformer or a UPS driving a small grid. I think there are ways of adapting it to a larger grid with multiple feeds, but I didn't get there.

In any system, there can only be one voltage source. Everything else has to be a current source or a resistor. If they're all current sources with resistive load lines, you won't have stability issues. I'll just put this out as a statement without trying to defend it: You can have different voltage sources as long as they have different frequency domains. By that, I mean you can have a voltage source at a solid-state transformer with a very slow response and you can have a very fast response at a storage capacitor. I think they can sort themselves out.

You have to be a little careful so the storage capacitor can handle the high-frequency ripple and the solid-state transformer would not pass it through to the grid. The solid-state transformer would adjust the load on the grid more slowly, so as not to shock the grid and bring the power balance back before the storage capacitor was depleted.

QUESTION:

I have a question about storage. What do your roadmaps show for the adoption of some of these storage technologies? Are we looking at 5 to 10 years before some of this stuff becomes possible?

DR. SATISH RAJAGOPALAN:

Well, it depends on the actual application of storage. If I look at the bulk level, it's a little bit more clear than the distribution of the edge of the grid scale just because the load level is much more defined in terms of the market.

There are proposed rule changes that would allow energy storage into bulk markets or those kinds of applications we see coming in a much near timeframe than within this next 3 to 5 years.

The distribution applications, they are fast emerging. There a lot of technical challenges, a lot of financial challenges. I would say at least 5 years before we see the wide-scale deployment of storage. We are still in the experimental phase right now.

QUESTION:

I want to comment. We have done the research for current advisement for smart grid applications and any integration, reliability, and control. I'm involved with one of them, supporting the university producing exactly the modular skeletal power converter transmission application.

I just want to pinpoint two barriers that the whole team, EPRI, the whole wider GE group -- we actually had our annual review two weeks ago, right after the GE summits in DC. They touch on distribution control. I want to point out that Professor Lau from Caltech was also doing exactly what the middle speaker was talking about the distributor, the ideal world, where you actually lower all the control from the control house.

DR. GREG SMEDLEY:

Right, so you have a distributed control.

QUESTION:

So there are different kinds of research projects that our group has supported. The one common theme among all the technology teams that we have at the GE meeting was when you try to push it through from demonstration to deployment for a bulk commercialization power; the utility industry that we're facing, just to ask for information, communication, and knowledge is vastly different from, say, communication from IT or all of the other industries.

I just want to make the point to whoever is here, you want to ramp up your power supply, you want to increase your voltage, I welcome you all. It's a pretty different space from what you guys are used to…different structure, regulation structure.

I think a lot of paralegals are trying to get together to write book to address that. In terms of the pace of improvement, everyone was just sitting there in the DC room and we were all… we just don't know the right way to go about that.

DR. GREG SMEDLEY:

In response to that, basically, I gave a presentation at a panel of venture capital investors and they said that when someone comes in to pitch an idea to them that goes into an automobile, they usually say to the person, "You're probably going to have a long beard that'll be colored gray before you see that product actually used in the marketplace."

They said that if someone comes in saying that have an idea for something to be sold to the utilities, they basically don't even talk to them. They said that utilities were invented to make the automobile industry look good.

Jokes aside, the key here is that when working towards utility-scale applications, there needs to be a high degree of patience and a willingness to work together to make it happen. There are other countries in the world where the grids are under a great deal more stress than they are in the United States. In some cases, those utilities may be more enthusiastic to try new technologies.

QUESTION:

I have a question primarily for Greg. I really like your concept of using the autonomous response to stabilize the grid voltage. That's very nice. I'm trying to figure out if I can do that with a micro inverter. Granted, the system doesn't see a single micro inverter, but maybe it will see 10,000 of them and maybe they could make a difference.

The key is I can make the phase angle anything I want of the current I choose out onto the grid, looking at the ac voltage waveform. The question is: can I do anything useful without having a control signal from the grid telling me what phase angle is needed?

DR. GREG SMEDLEY:

That's the whole point, basically. If you listen to what the grid or the system has to say, it will tell you everything you need to know. So when the grid is stressed, you will see a change in the voltage. You will see a change in the frequency.

DUSTY BECKER:

You have to become a grid talker.

DR. GREG SMEDLEY:

Grid whisperer!

QUESTION:

Just to add a comment on top of that. I was at a workshop at Solar Power International a couple of years ago and a converter company said, "Well, we can do all that." The look of horror from the utility people in the audience was priceless.

DR. GREG SMEDLEY:

There's an example. This has already been done. There's a company in the Northeast; they deployed, I believe, 200,000 node points.

JONATHAN POLLET:

Petro Solar.

DR. GREG SMEDLEY:

Yes, Petro Solar. That was about 200 watts or 200 VARs or 0.2 k-bars per system maximum if you had no PV going on. It adds up. 200,000 nodes is a pretty respectable amount of megaVARs.

They mounted it on the utility poles. That was their idea for distributing all of the wiring all over New York and New Jersey on the utility pole. That's how those nodes were created. My understanding is that many of those PV panels ended up as very nice sails when a big wind came through.

QUESTION:

If I'm not mistaken, they have a central control system.

DR. GREG SMEDLEY:

I believe that's correct. They are not autonomous nodes, but they are control points, at least, out there that can be activated. Actuators, we would call them.

Section VIII, Session 4
The Future

Changes to End-User Equipment

DUSTY BECKER:

Our next speaker, Conor Quinn, is Director of Technical Marketing at Emerson Network Power. He has product line responsibility for low-power standard products, in addition to driving technology roadmaps and other initiatives in energy-sensitive applications. He is a regular contributor to the specification and roadmap activities of industry groups, including Power Sources Manufacturers Association, Power Management Bus, and PCI Industrial Computers Manufacturers Group.

Conor Quinn currently serves on the PMBus Board of Directors. His experience in the power electronics industry spans over 20 years in design, management, and marketing roles. Dr. Quinn holds a BE in Electrical Engineering from University College Cork in Ireland and a Ph.D. in Engineering from the University of Minnesota.

DR. CONOR QUINN:

Thanks, everyone, for your time today. Well, I read this question "Am I smart enough for the smart grid?" and I said, "No! So I better sign up for your workshop!"

Then Ed and Dusty saw my name on the list and they had had someone pull out, so they asked me to speak. I wasn't smart enough to get out of that either. So here I am.

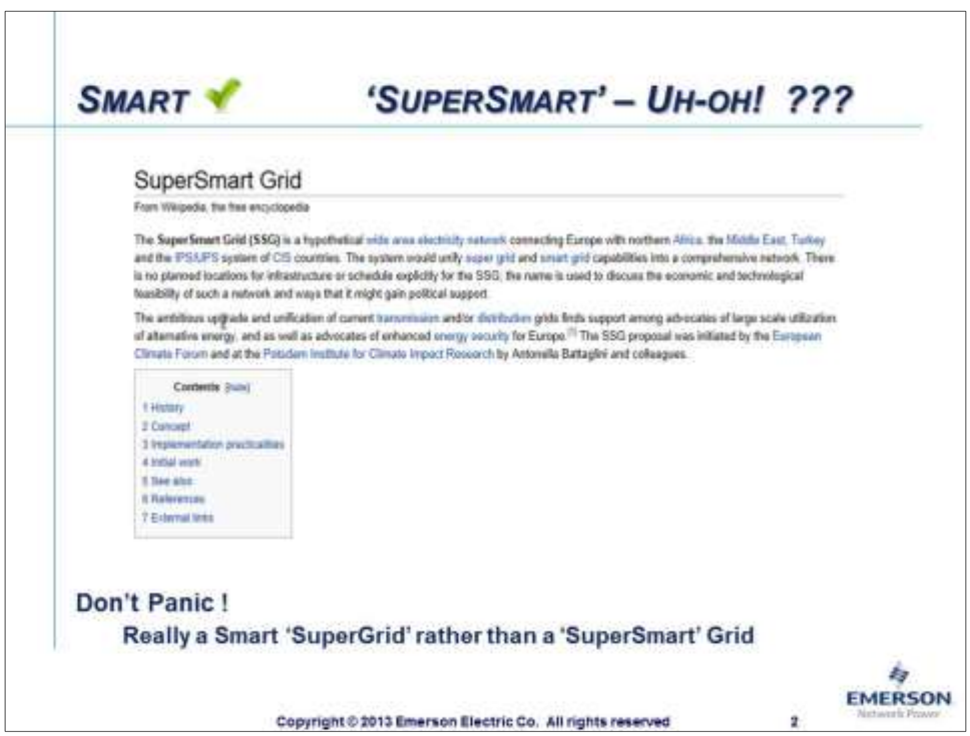

I said I wasn't smart enough for the smart grid. Then I got really worried when I started reading about super smart grids. It turns out; it's really a smart super grid, not a super smart grid.

So I stopped panicking… briefly.

I'm giving this talk from the perspective of a power supply designer.

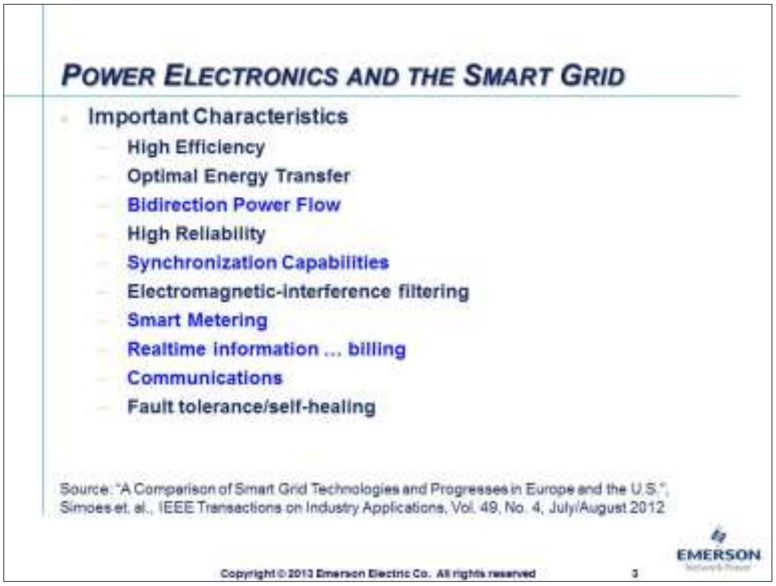

Here is a list of important characteristics for power electronics devices. Just bear in mind that I'm giving this talk from the perspective of a power supply designer for a traditional embedded power supply company. I'll try to stretch it out a little bit where I can. This list was put together by someone who was looking at the larger power electronic systems for the smart grid, but a lot of characteristics are very similar to what we're all fighting every day.

High-efficiency optimal energy transfer. Maybe bi-directional power flow. Some of us are not as experienced in that area as others. High reliability; we get that all the time; it's a given in our industry. Maybe synchronization is a little out there for some of us. EMI is a constant battle. Then the whole area of smart metering real-time information, building communications.

Five, ten years ago, the power supply industry was an analog world. Today it's a digital world.

Five or ten years ago, at least, the power supply industry was a very analog world. Today it's a very digital world. Almost every power supply we create above a certain power level is digital today... not only for communications, but for closing the loop.

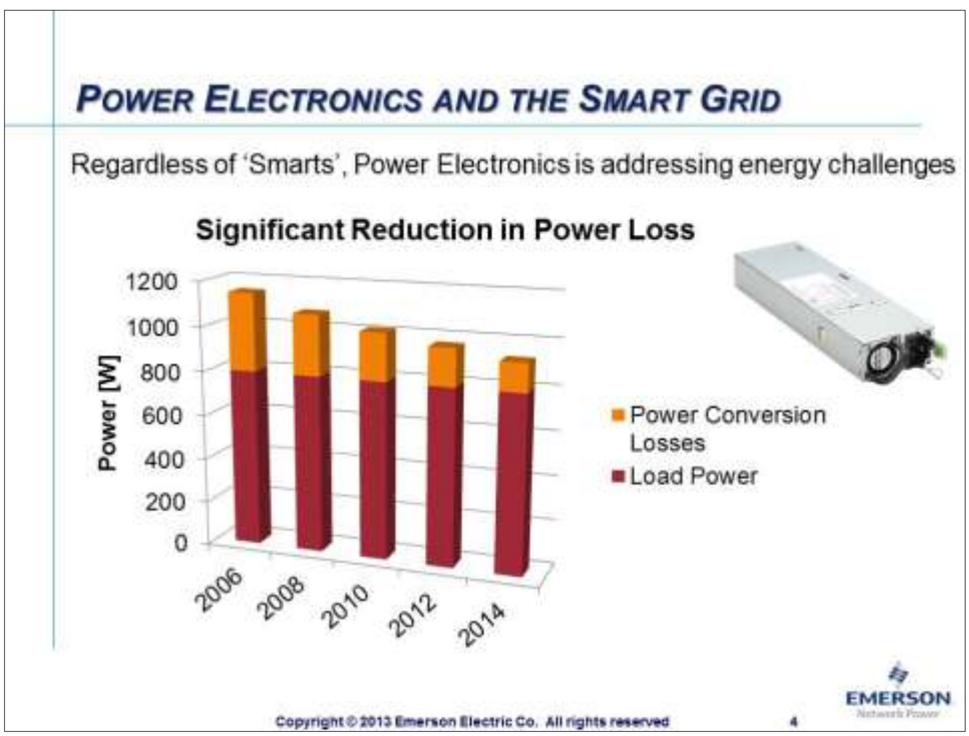

This is a world that we are becoming quite comfortable in and, in some cases, leading. One of your other speakers this morning was just talking about the advances in digital. It's now -- the processing power is now cost-effective enough that we can use it to do things in power supply we couldn't five, ten years ago.

Regardless of the smarts, I want to stand up here and give some kudos to the power supply industry and our suppliers, our management guys, power semiconductor guys. We've made huge progress trying to reduce or mitigate some of the stresses on the grid. You look at a typical 800-watt server application, for example; we've reduced, probably by a factor of two-thirds, the amount of losses that are due to power conversion. I'm talking about from the ac down to processors. So ac-dc and dc-dc. I'm not including UPSs; the number might go up a little there.

We've reduced, by a factor of two-thirds, the amount of losses that are due to power conversion.

So we've already made significant contributions. We are efficient. The comment from Satish was that our converters need to be more and more efficient. With that comes density and a reliability trade-off in some cases, but lots of good progress there already.

Our bosses only get excited when they see large numbers, so I better put a few numbers up here. We see numbers ranging from hundreds of thousands for demand-response appliances. Some of this data is from Darnell.

Sub-metering: We're going from 25 million units to 100 million units in about five years. All are predictions, obviously, so those are from the market analysts.

How quickly will some of this stuff become cost competitive?

DC microgrid power electronics. Now we're talking billions, according to the analysts. The potential is there. I don't know if we can achieve this, which is why I was asking questions about some of the roadmaps… how quickly will some of this stuff become cost competitive?

There's a big number here for data centers and telecom. As you can see, there are some pretty large numbers here in commercial industrial buildings. Those in the power line communications business depict really large numbers there too.

Here's a list of a some of the opportunities for the power electronics world. Again, I was trying to put a little bit perspective on this from the power management, power semiconductor, power supply people here. These are the areas where we can potentially play in some of it. We're playing in those power levels today and, in some of them, we probably have to move up a level or two.

Again, just a little bit perspective. I'm really talking from this group of members of PSMA, the traditional power supply manufacturer embedded people. There are really not too many people doing inverters, adjustable speed drives, that kind of a thing. I think we have some side projects with the active component people, who are our suppliers and possibly suppliers to the end equipment.

This third of the business; maybe some university academic people; somebody can sell them. They're already way ahead of us and know a lot of what I'm going to bring up as potential issues and challenges.

There's a huge gap in communications. Not just in how we communicate electronically, but how we talk to each other.

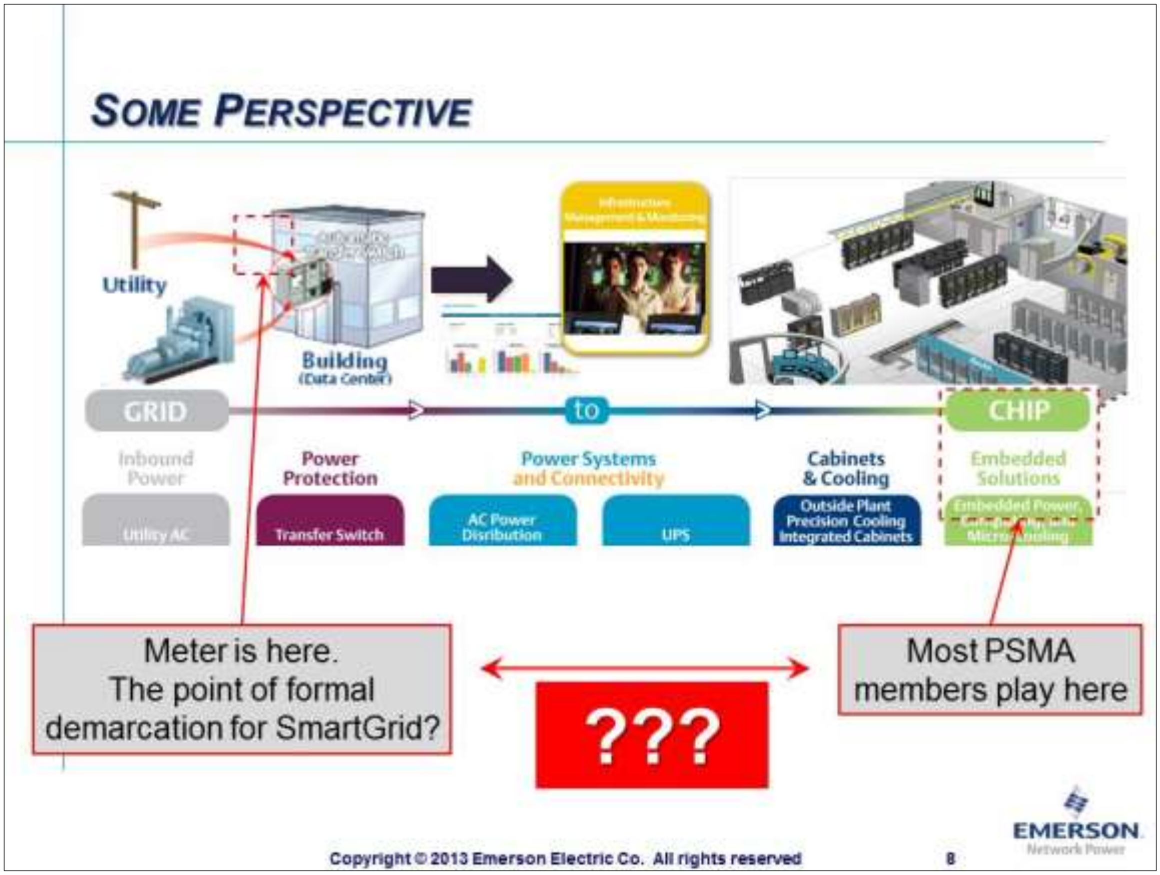

This is an Emerson diagram of different Emerson businesses laying all the way along here on the power path. The group I'm in and that group that I was pointing out, that part of the PSMA membership, would argue that the vast majority is played out down here. We don't play much up here, maybe with a few exceptions. Metering is here and, I think, that's where the utility people and us primarily come together. It seems to me, there's still a huge gap in communications. Not just in how we communicate electronically, but how we talk to each other. Hopefully this workshop is advancing those discussions.

I'm from the perspective of the power supply industry. I've talked to a few colleagues. I've talked to a few competitors. What do you know about the smart grid? This was probably the most common answer: "It's not a dumb grid."

A lot of us really don't know the issues.

A lot of us really don't know all the issues; we don't know what we're expected to solve. As the last speaker said, "Maybe there are things we know that you would like to know." So here we are. Of course my boss, being the skeptic and cynic that he is, said it sounds like a great way to get some money.

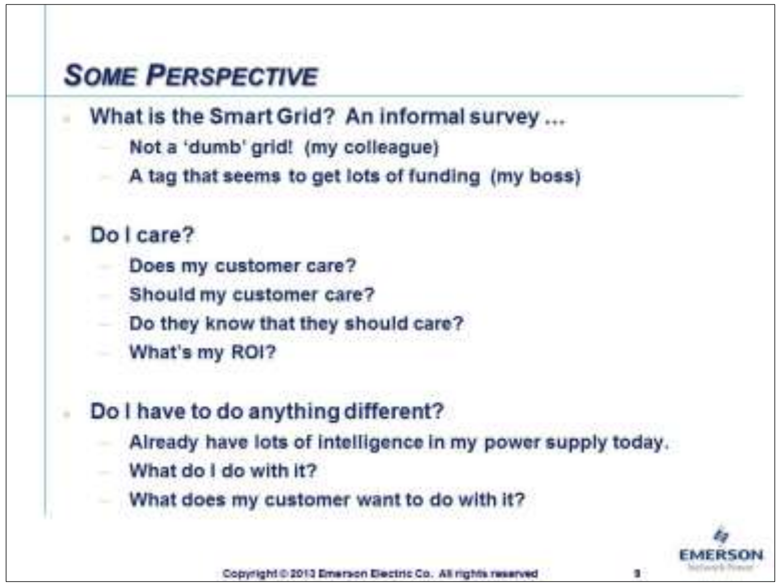

The "smart grid" tag seems to be applied to a lot of things these days to get funding. Ultimately, do I care? Does my customer care? Should my customer care? Do they know they should care? I don't have a lot of customers coming to me and saying, "We have this big challenge and this wonderful opportunity in the smart grid." Ultimately, what's my ROI?

So again, what do I need to do differently?

I came to this workshop hoping to learn a lot about this. So apologies that some of the slides, some of my thoughts here, are ending up on the page.

I already have lots of intelligence in my power supply.

I already have lots of intelligence in my power supply. What do I do with it? What does my customer want to do with it? What would the utility like for my customer to do with it?

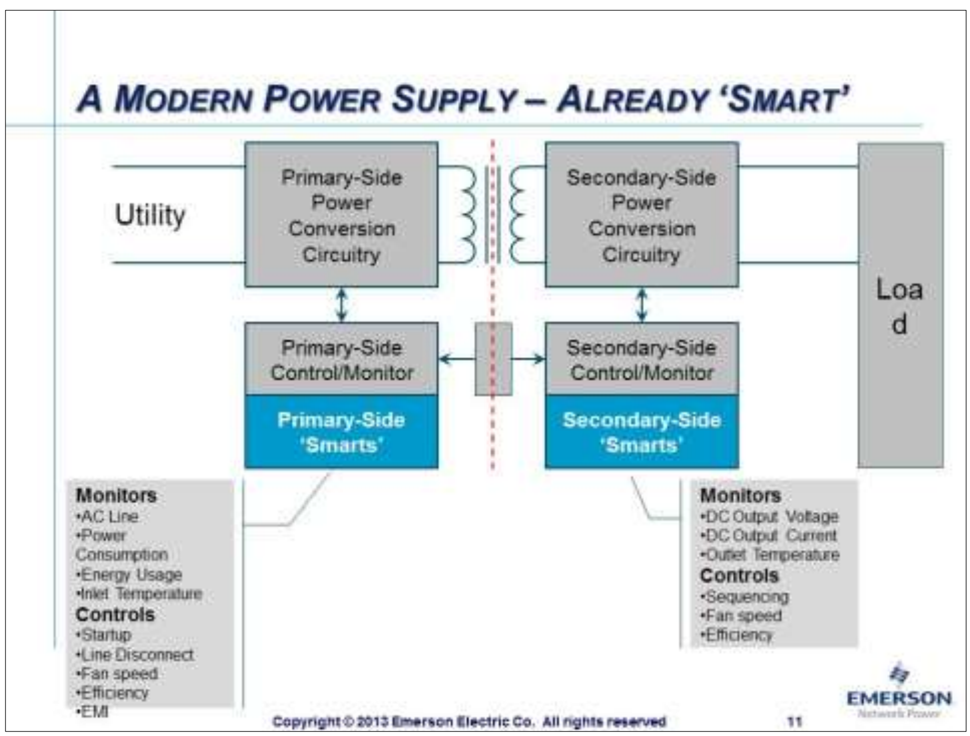

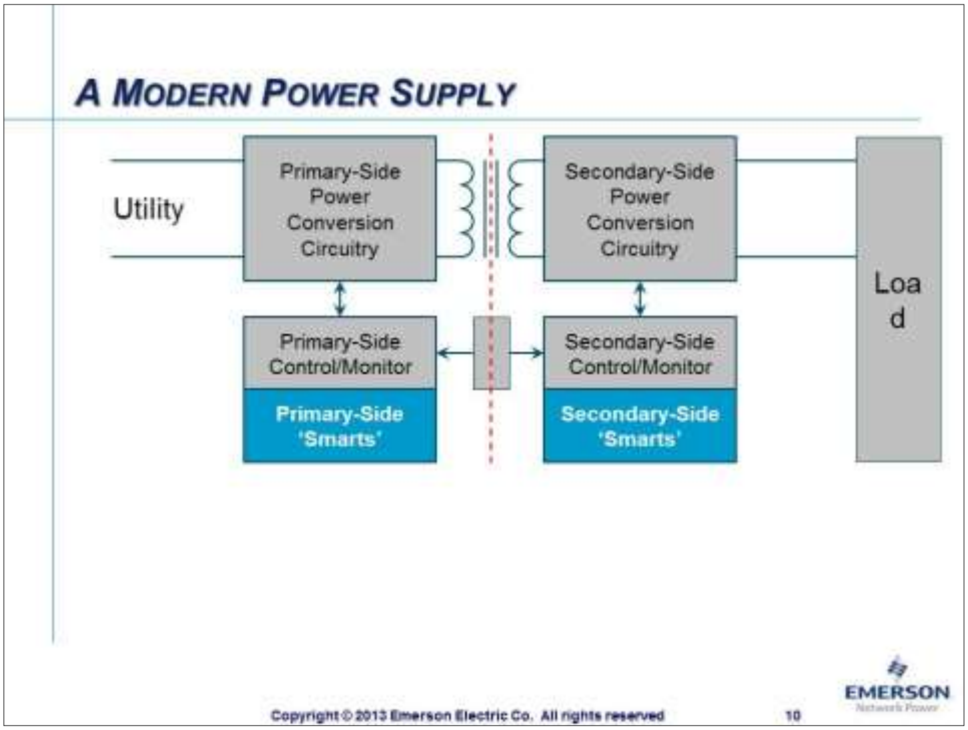

Some of the slides just walk through the traditional modern power supply for some of you who may not be as familiar. We've got the power conversion circuit, typically consisting of some primary side and secondary side. I think we use different nomenclature than the utility sector. Secondary and primary mean something different, as was pointed out this morning.

We've got some controls monitoring going on within the power conversion. Five, ten years ago, most of that was analog. Today most of that, at least above a certain power level, is digital. So now we have a lot of smarts that previously may have been added on.

We've got processing power. We can implement bus interfaces.

Now it's all there. We've got processing power. We can implement bus interfaces. All that kind of stuff is available to us and it doesn't add a whole lot of cost.

What are the efficiency improvements that we talked about? They've been done, for the most part, without adding to the price of a power supply.

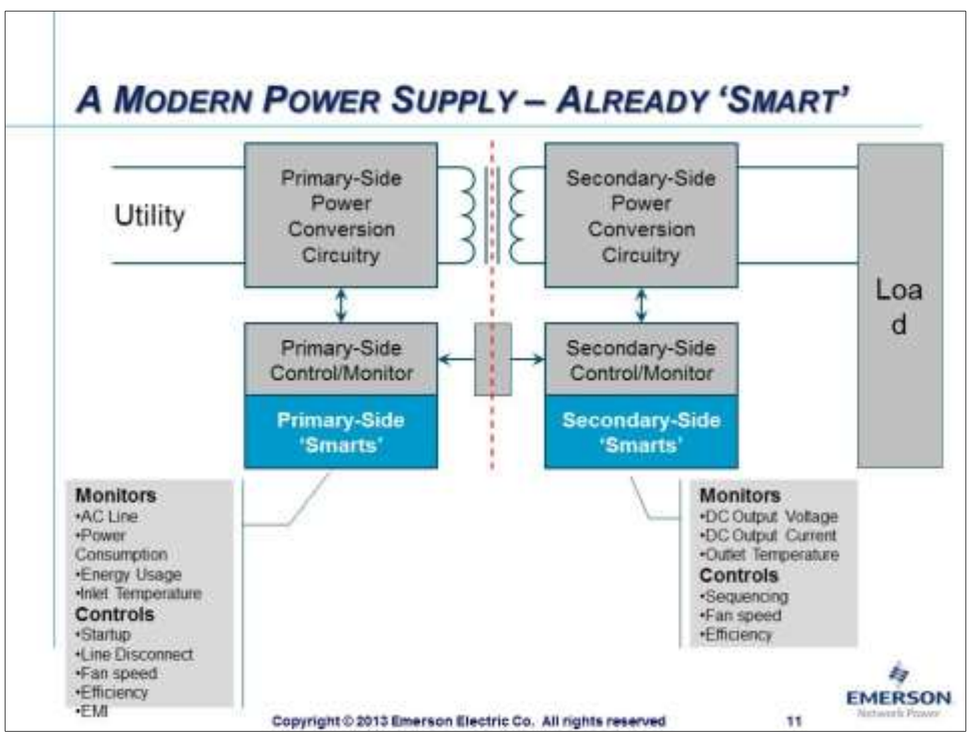

These are some of the kind of things we can do. We can and do monitor the ac line. We monitor power consumption and energy usage. Whereas five or ten years ago, we might have been measuring things to 5 or 10% accuracy because we couldn't afford anything better; now that may be down to a 2% accuracy. Maybe even down to 1% accuracy. In some cases, not enough for smart metering, but it's still probably good enough to report some information or do things that you might like us to do.

We're using smarts for startup.

We're measuring temperatures, primary side up, secondary side inlet, and outlet air temperature. We're using these smarts for startup; maybe for line disconnect, in certain applications. We can control fan speed. We're using it for efficiency today, so our power supplies are internally dynamic. Today, they're operating under different conditions: low line versus high line, heavy load versus light load. That's how we use some of the efficiency gains we've achieved over the last five years.

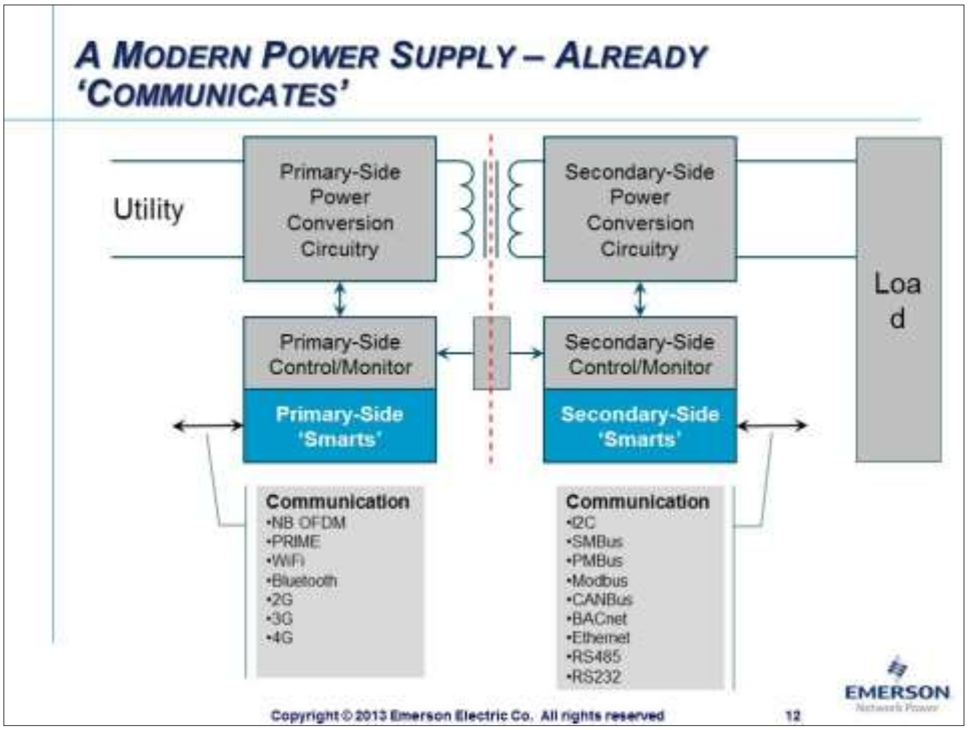

Now we can monitor and report what our output loads are drawing, how much that Intel® processor pulling from the power supply, things like that. There's a lot of capability in here.

Some of our customers use it. A lot of our customers say, "Oh, that's really nice. I'd like to use it. Let me think about it," and never come back to us. There are a lot of customers out there who can use it. I suspect a lot of customers in other parts of the ecosystem could use that information if they knew it was available.

I suspect a lot of customers could use that information if they knew it existed.

All right. We talked about some of the stuff we're monitoring. We're also doing a lot of reporting. Apologies if I'm mixing up physical layers and protocols and things like that, but it's an alphabet soup of types of protocols and physical layers that we're implementing today. All of these: I^2C Cord, SMbus, Modbus, CANbus, BACnet for energy management, Ethernet; all of these could be on the primary side. We could communicate back to the power line with our targeted frequency division multiplexing, the prime standards. Or we can, and do, wireless communication.

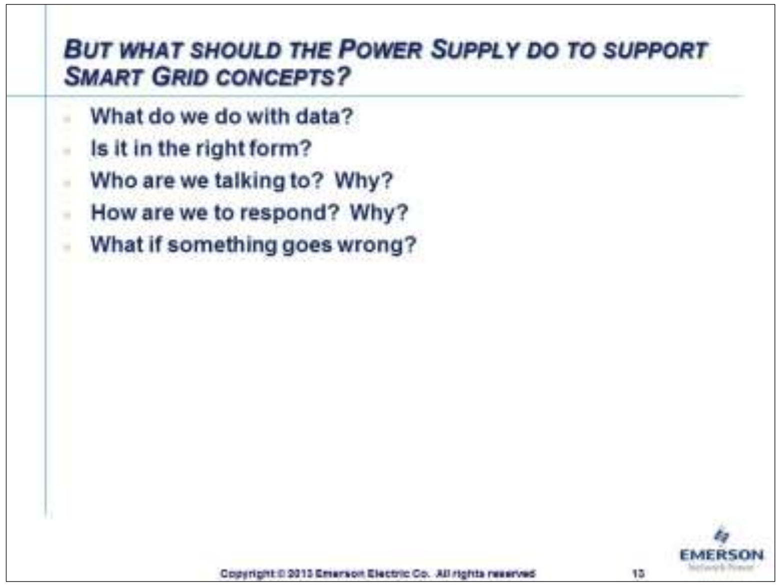

All of this stuff, we can do today. And we are doing it. Some of it; like 2G, 3G, 4G; may cost a little more than a lot of our customers want to pay, but the capabilities are there.

What if something goes wrong?

What do we do with the data? Is it in the right form? Who are we talking to and why? How are we to respond? What if something goes wrong? We talked about autonomous earlier; there are probably a lot of other things that we can do.

We can talk to the load. We do talk to the load today. So there are things we can do there. We can talk to a hub, so maybe a home area network. Rather than talking to the load, we're talking to a device that pulls a lot of other devices together and, together, they make some decisions.

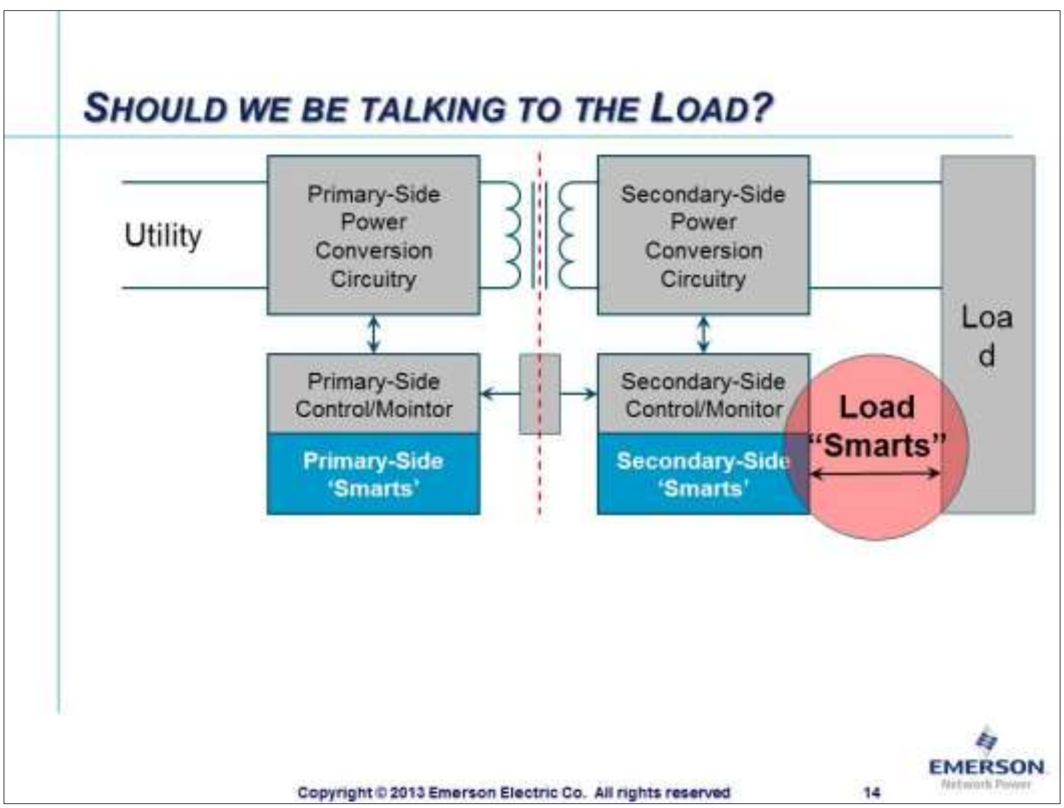

The question here is what are we doing with this data?

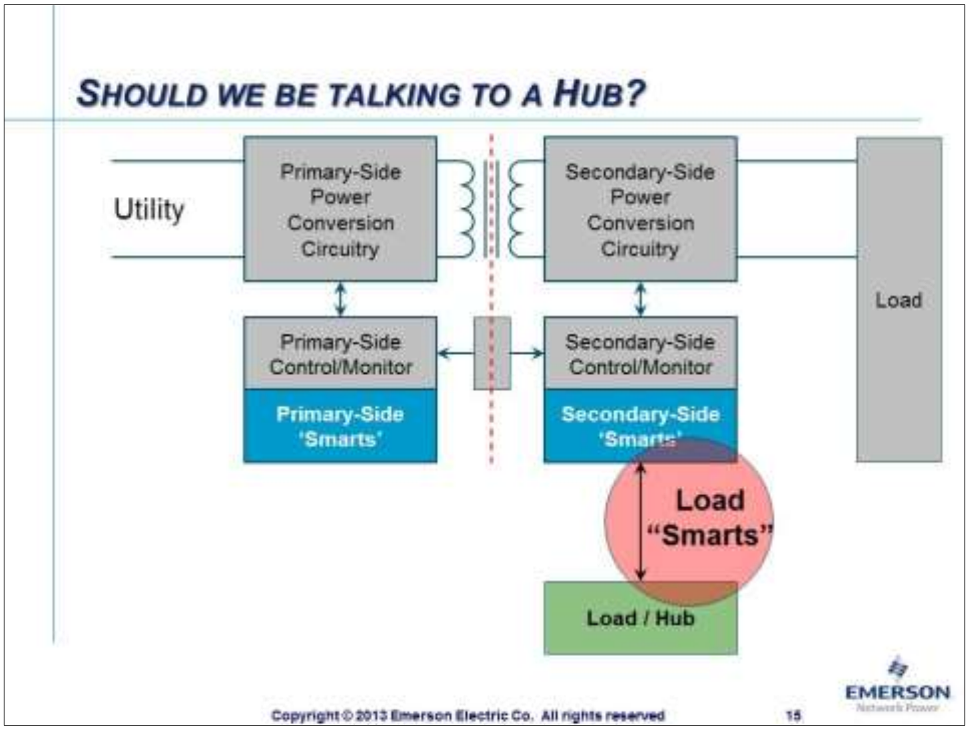

We could be doing that on the secondary side or the primary side; still a hub. We could, through some sort of coupling mechanism, implement power line communications. Again, the question here is: what are we doing with this data? Where is it going? Who's listening? Who wants to see this data?

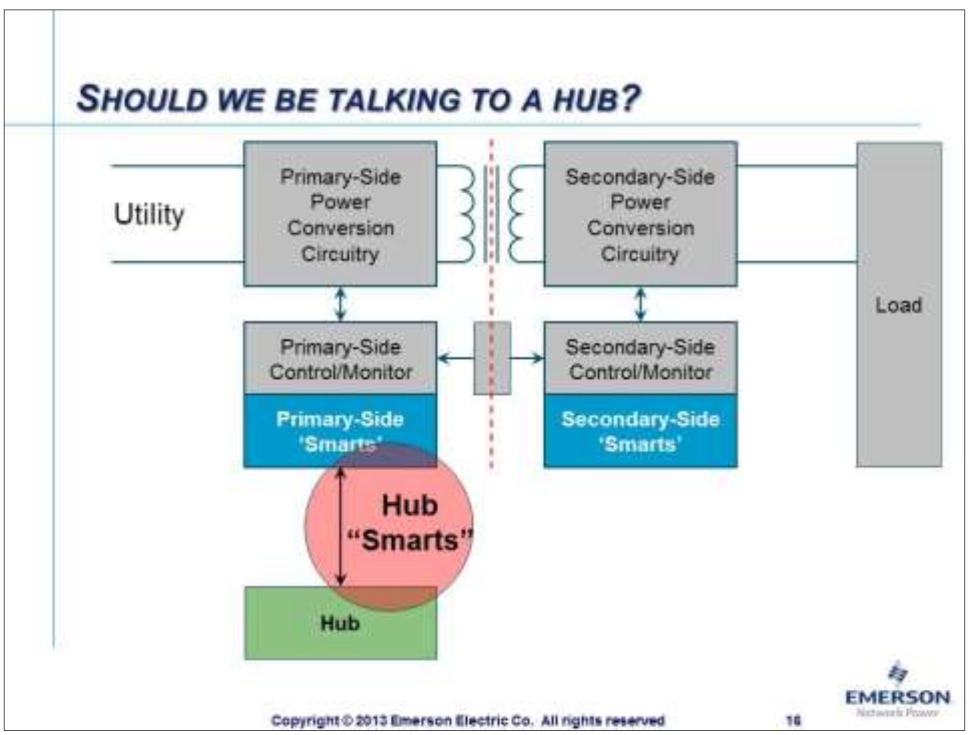

The point of all these slides is that it is available and we do it today. I think we can make a lot more use of it, but we, as an industry, need a lot of guidance about what to do with that.

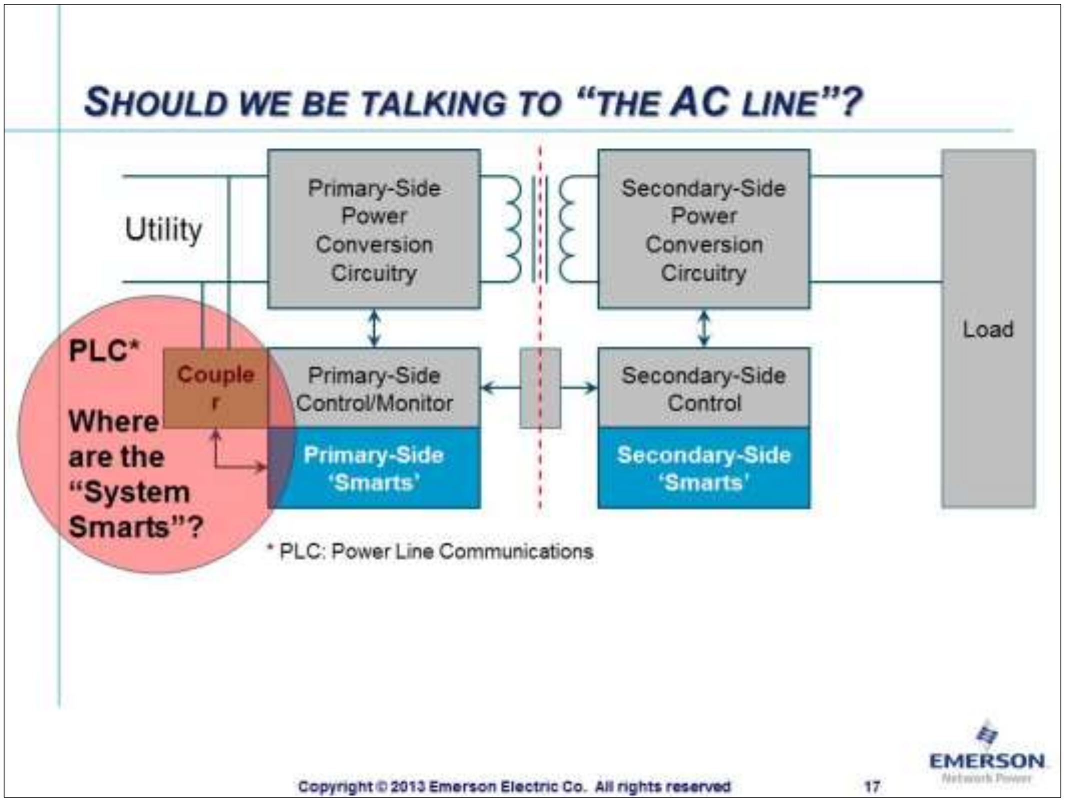

Here are some examples to give you an idea of what some of these applications are starting to do. I think there's a lot more that could be done. I'll try to touch on how the power suppliers are using them, but, obviously, I have some questions too.

Let's start to look at some of the home and residential stuff. I don't do any of this at home today, but I just started to dig in a little bit. What an alphabet soup this is too!

> *How many are just going to mess things up for the rest of us?*

How many of these play together? How many of these compete with each of other? How many are just going to mess things up for the rest of us? What do we already know? What do we need to do?

Two of those press releases, I think, were dated the same day. So there are different things going on in the power-net, obviously.

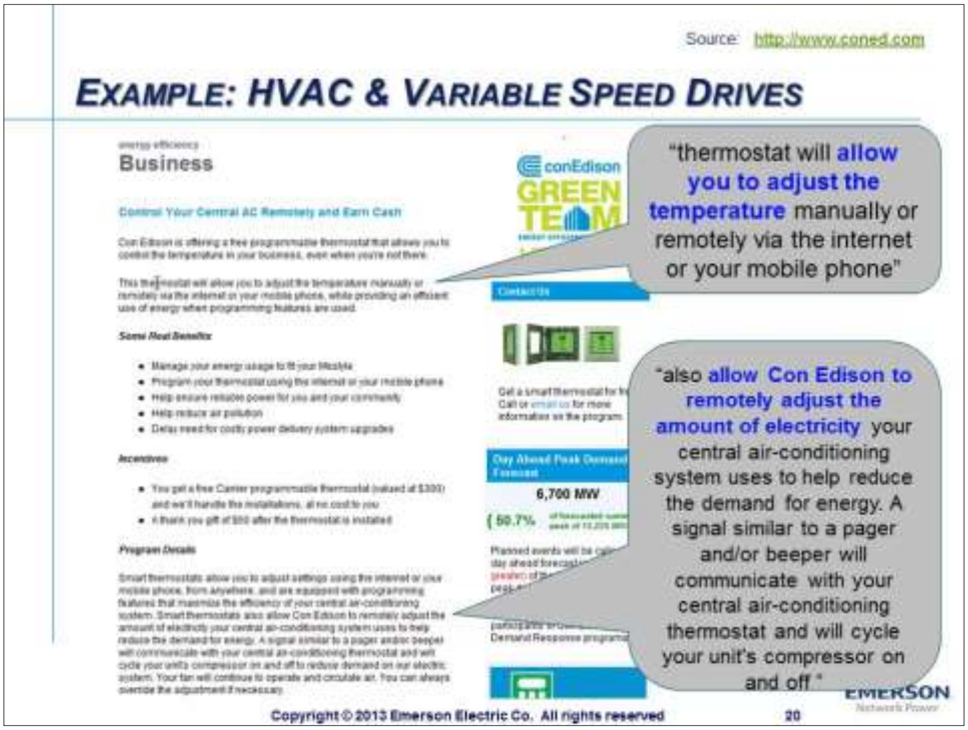

This one was mentioned earlier, so I won't spend a whole lot of time on it… but we do a little bit of adjustable speed-drive business in our group, powering HV/AC equipment. Con Edison has their program and other utilities have their programs where they want to talk to you or your HV system.

The reason they're doing all this, is they really want to get access to control the thermostat.

They sell this to the user as a thermostat that will allow you to adjust the temperature manually or remotely. So they give you a smart thermostat, they give you an app, and you can do all of this remotely. Save yourself some time. Of course, the reason they're doing all of this is because they really want to get access too, so it allows them to do all this.

Just an example of what's being done today with smart, sometimes very smart, thermostats (sometimes with medium smarts on it). A thermostat can communicate with the furnace or air handler controller or the air conditioner. They, in turn, communicate with the adjustable speed drives that are mounted on the motors or mounted in proximity to the motors. So a lot of smarts are creeping into areas where power electronics are part of the solution.

LED lighting is very efficient, you can have a lot of control over it.

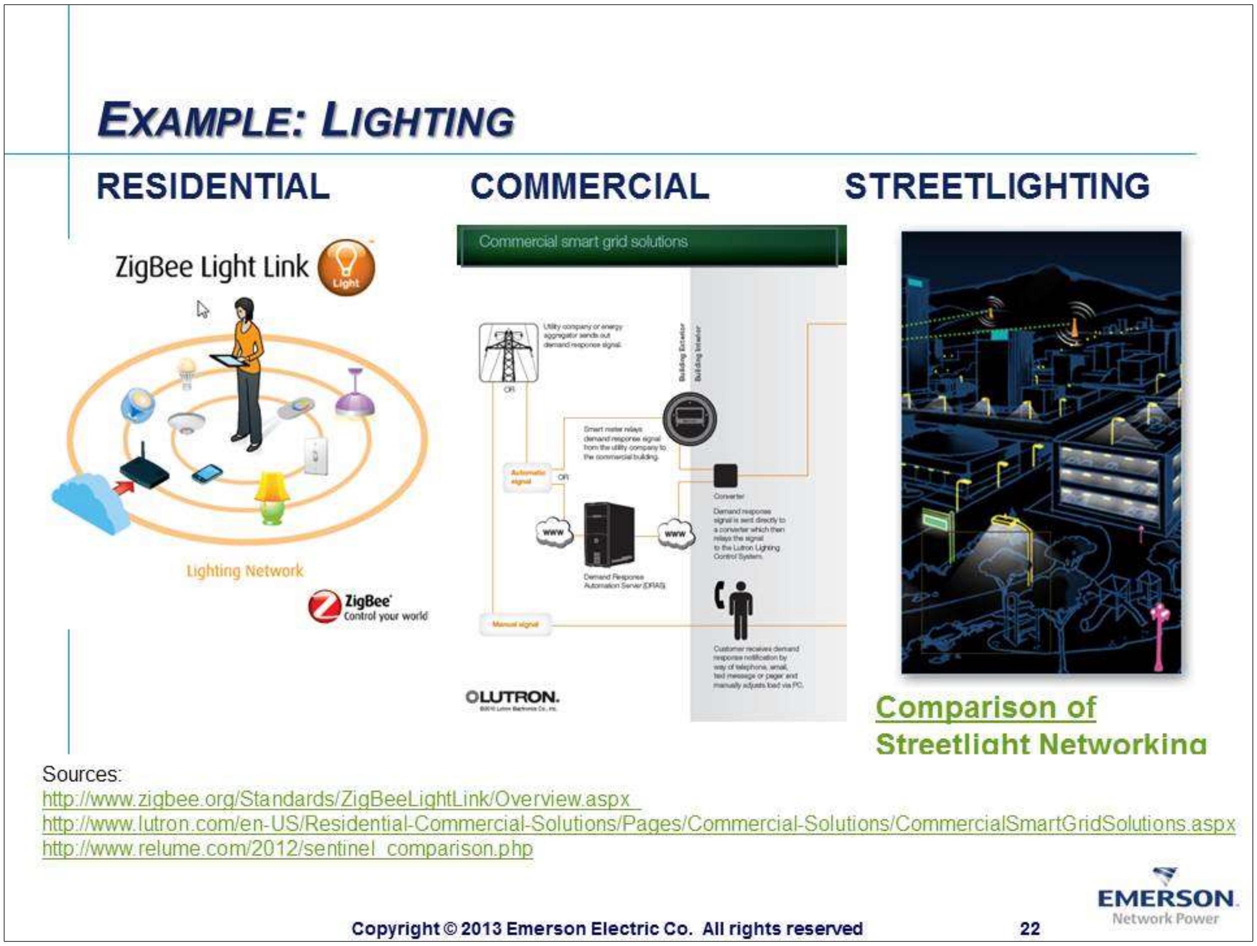

Lighting. A lot of stuff is going on in lighting. LED lighting is very efficient; you have a lot of control over it; but it costs a lot of money today. It's expensive to buy an LED light bulb, as some of you probably know. It gets a little more cost competitive when replace whole fixtures. They're having to sell a lot of the value based on smarts and things other than just efficiency.

So there's a lot of work going on in this area on various types of networking: a ZigBee Light link[2] at the residential level; a system from Lutron link[3] for commercial level where you've got servers and demand-response units.

Then you go to the street lighting, where LEDs can be sold on reliability if you can find power electronics with the same reliability. There's a good link here[4], but this compares some of the different streetlight networking systems that could be compatible with the smart grid; the networks ranges from radio repeater systems to power line communications.

---

[2] http://www.zigbee.org/Standards/ZigBeeLightLink/Overview.aspx
[3] http://www.lutron.com/en-US/Residential-Commercial-Solutions/Pages/Commercial-Solutions/CommercialSmartGridSolutions.aspx
[4] http://www.relume.com/2012/sentinel_comparison.php

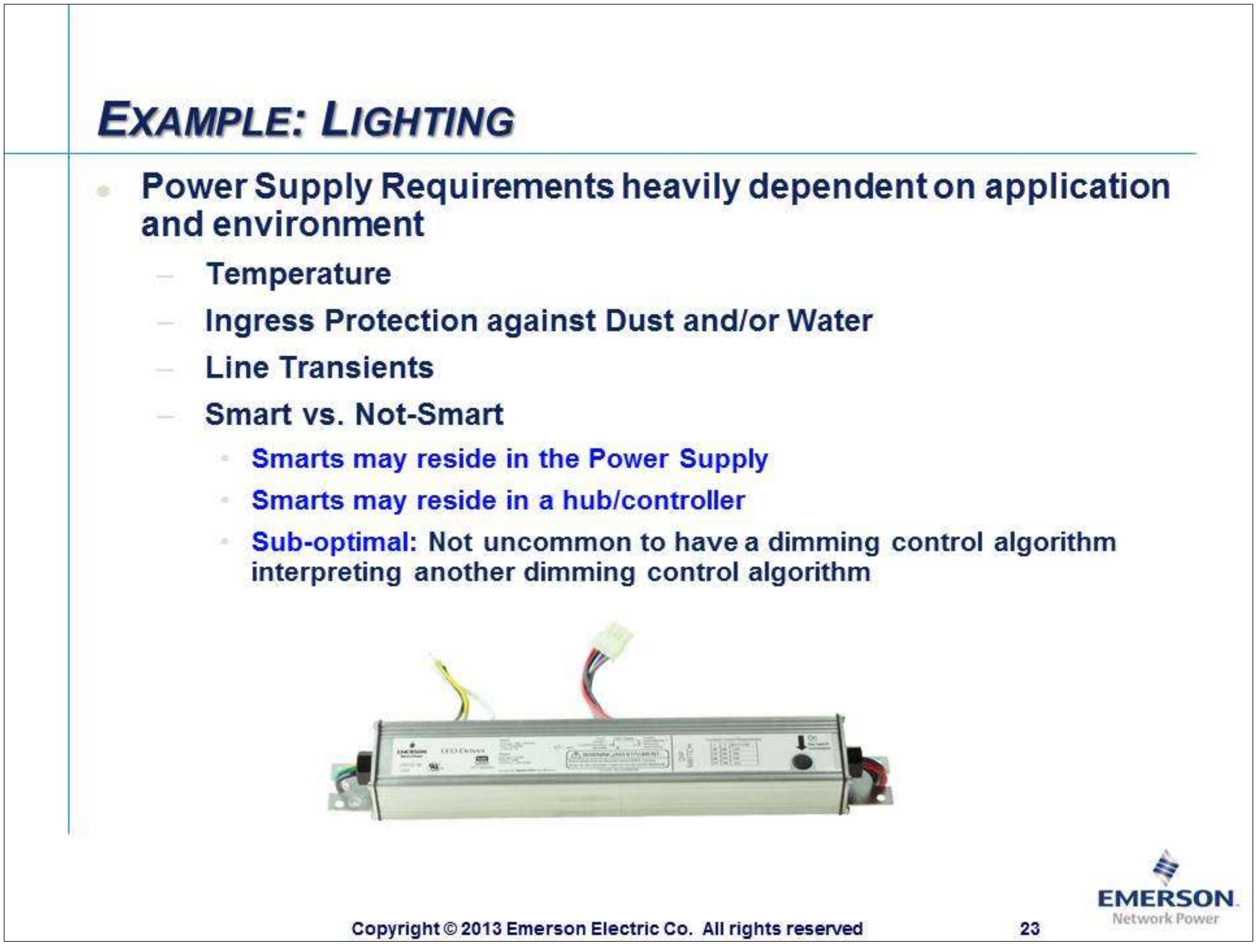

In terms of the power supplies for lighting -- what they look like, what they feel like -- is very heavily dependent on the application and environment. There is temperature, ingress protection, line transients, smart versus not smart, and all of that kind of stuff to work through.

Some smarts reside in power supplies, but a lot of smarts reside in hubs or controllers externally.

What you see a lot of is some smarts residing in power supplies, but there's also a lot of smarts residing in hubs or controllers externally. I would argue they are very suboptimal here. The smarts are in multiple layers, but the smarts are not always used at those multiple layers. We've got products that can do a lot of fancy dimming, but it turns out that the user side just gives us a simple analog surge, a 100-volt dimming, which end up taking out all of the smarts. There's duplication of cost and a duplication of effort there.

Solar is another area where people are trying to differentiate as price competitive/cost competitive issues come in. So they're trying to do more and more with that inverter.

In addition to inverting and putting real power onto the grid, these inverters can do reactive power provisioning.

I hear there's an inverter from Fronius[5] that is 10 kilowatt, I think. They've now pushed some of this same technology down to lower power levels. As well as just inverting and putting real power back onto the grid, they can do reactive power provisioning. They're advocating the advantages of their three-phase balance circuits and grid back up. I forgot what that's all about. I think it's a storage unit or something.

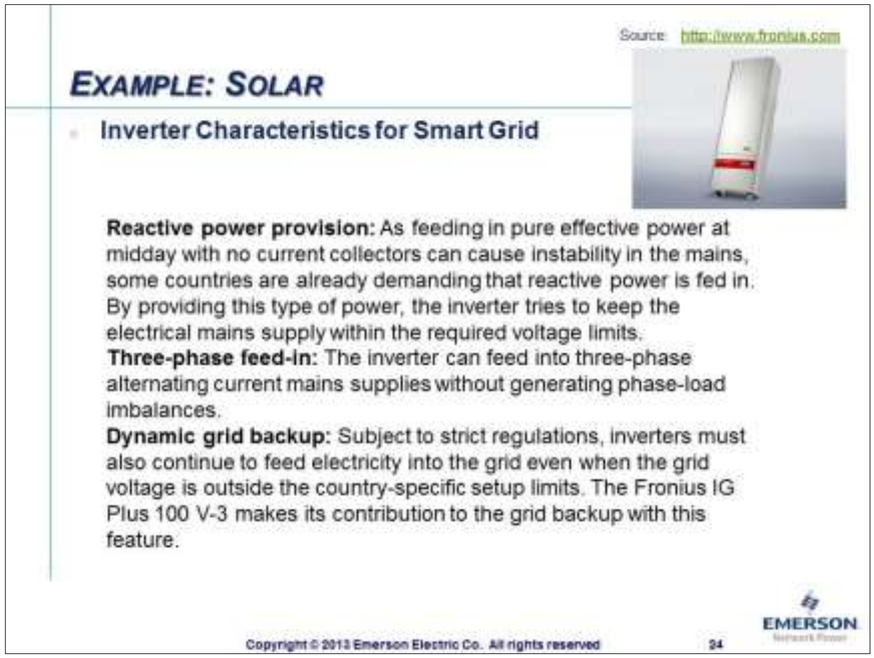

Then you move up a level to the utility-scale converters and storage[6]. More standards or more profiles are being implemented on top of or within standards at the facilitation and integration of inverters and storage in this case too.

[5] http://www.fronius.com/
[6] http://www.dnp.org/default.aspx

Electric vehicles have been mentioned by a couple of people this morning. The embedded power industry is starting to get more and more involved. In some cases, in the area of chargers, and, in some cases, in the area of dc-dc converters onboard the vehicle. It's definitely of interest to our industry. There's different levels of charging.

I think that the key point here is the rate of proliferation is probably the biggest concern: how that can be handled? You can buy EV chargers now on Amazon. I learned that yesterday.

Level	Application	Voltage	Amperage	Associated Standards	Description
1	Residential	120	<30	SAE J1772	Same power source used for standard home appliances in the US. Provides ac energy to vehicles; often portable devices. Standard speed charging
2	Residential and Commercial	208 – 240	30 – 40	SAE J1772, IEC 62196, IEC 60309 16 A	Can use the same power source as larger home appliances in the US (i.e. dryer) to provide an ac energy supply. Moderate speed charging
3	Commercial	400+	>40	CHAdeMO	Dc energy supply for networks and commercial use. High speed charging

As an example of autonomous behavior, Battelle and AeroVironment have some frequency-monitoring technology that they believe will help stabilize the grid.

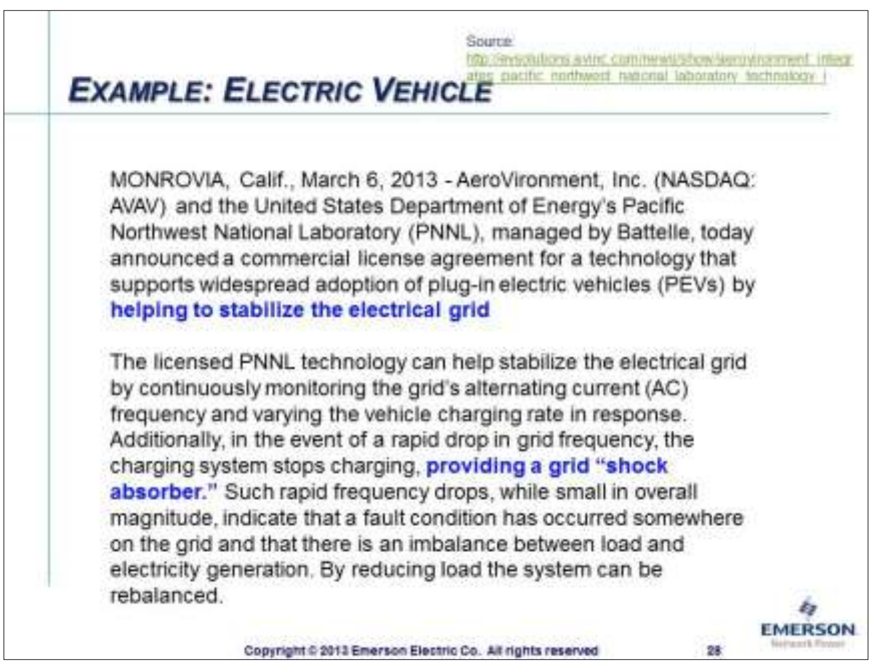

Because I was pulling some stuff together, this is another one I came across. I sent this off to Ed as a possible example of some autonomous behavior. Battelle and AeroVironment have entered into an agreement: they've got some technology that they believe will help stabilize the grid. In this case, they're doing it by monitoring what happens to the frequency on the grid when you add these kinds of loads.

Data centers. Two of the top three issues are of concern to data-center people, the facility people, and the networking people; tie in with this conversation. The whole business of monitoring and managing, what's going on there? Not only their IT capabilities; but the power, power usage, availability, and up-time.

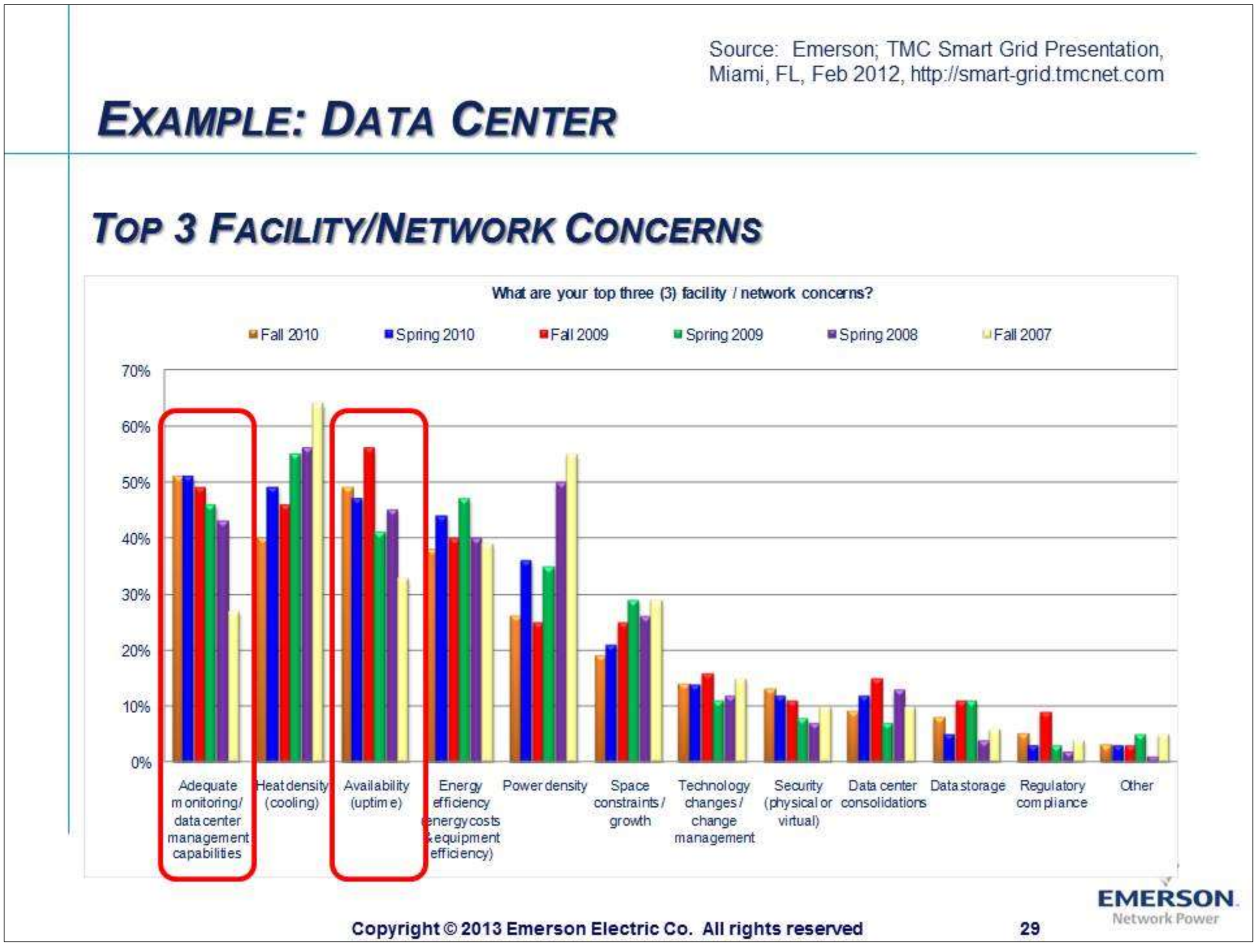

What are your top three? Two of the top three turn out to be very much tied to the smart grid.

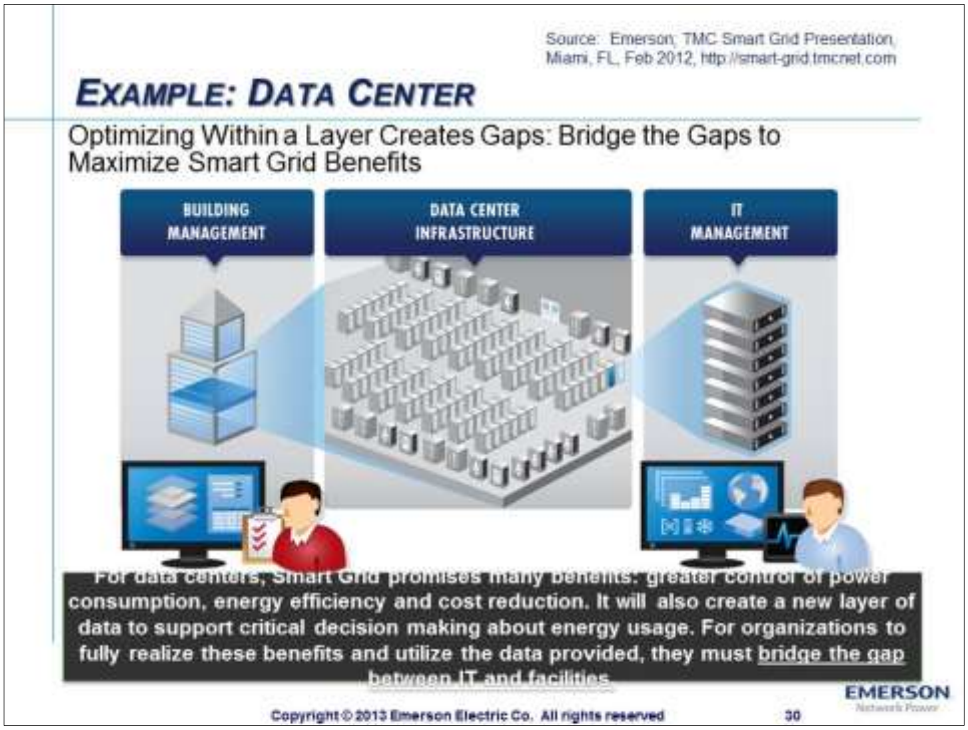

Getting back to optimization and avoiding sub-optimization: there are a lot of layers here. Building management people have concerns at data centers. Infrastructure people want to know how to tie all the servers together and power them. The IT guys just want bits in and bits out. This is all heavily computerized today.

How do you respond when they are demand-response commands?

So how do you tie it all together? How do you pull it in together? How do you communicate with the utility? How do you respond when there are demand-response issues, demand-response commands, things like that? There's a lot of work going on in this area today, memory being a big part of it. Tying all of these layers together and ultimately being able to respond to issues of power and power quality, things like that, are on our agenda.

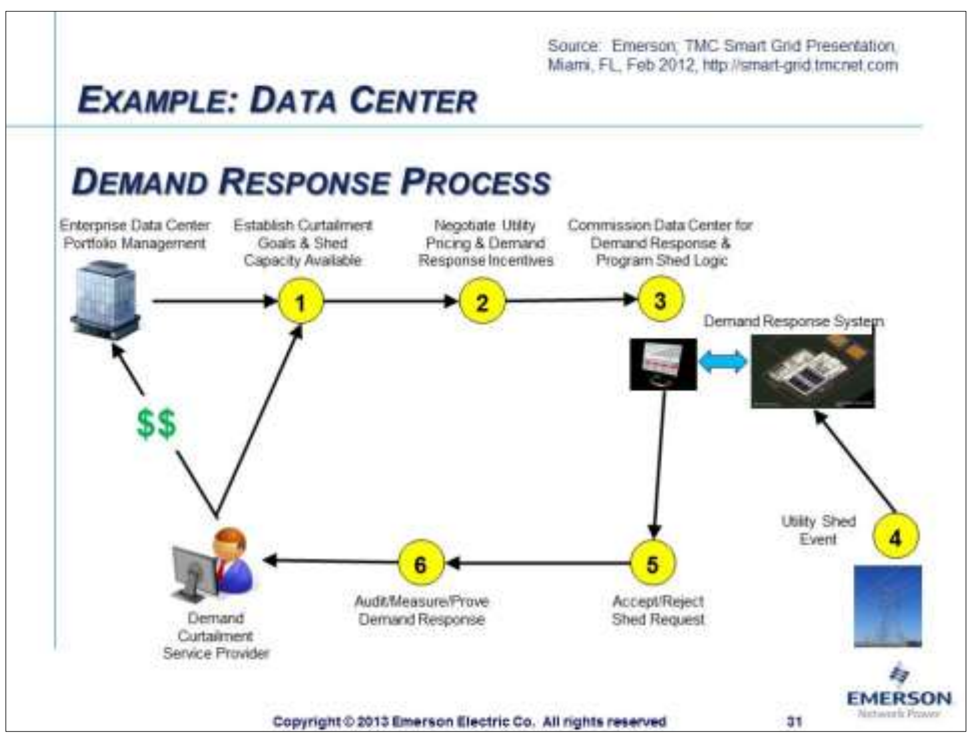

Here's an example of utility shed event and man's response. You've got some goals, tail-end goals, and shed capacity already predefined. You've got someone sitting there, I guess watching the dollars, making sure it's okay if I turn this server off and not that server.

Microgrid was mentioned earlier. 380 volts. Some of us think that this is high voltage. Some of us are moving into 480 volt ac and now we're starting to see 800 volts. That's getting really high. A lot of processors work with less than 1 volt. That's a huge range.

380 volts... Some of us think that this is high voltage.

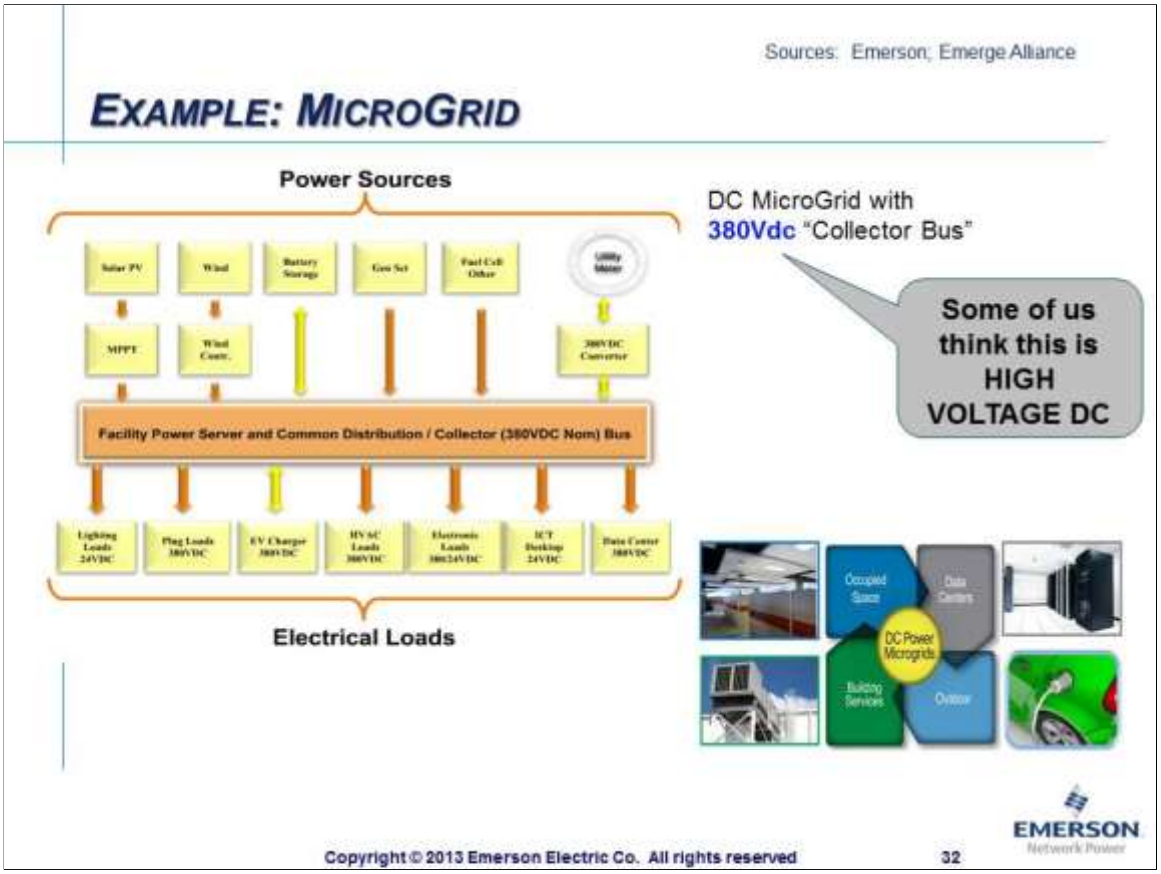

The microgrid was talked about this morning, so here's just an example of a dc-type microgrid where everything ultimately gets tied together at this 380-volt bus that can facilitate storage. Some people argue that 380-volts dc is inherently more efficient than some other systems. Some people argue the opposite. It can certainly help you tie in with photovoltaic systems and others. Regardless of whether or not you buy into the pure efficiency argument, the conversion efficiency argument; there are a lot of other reasons for looking at these kinds of systems.

Dennis talked about convergent lines this morning. So microgrids are being looked at for many different types of systems here, many different types of applications.

These slides are more of the same. The power electronics people. Virtually every single one of these boxes has power electronics in them, so there are a lot of opportunities if we knew what to do.

One other comment about this one and even the previous slide… All of these arrows are showing power flow. They're not showing communications and data flow. That's something that need to be addressed.

There's a lot of clarity needed around integration, communications, standards, etc.

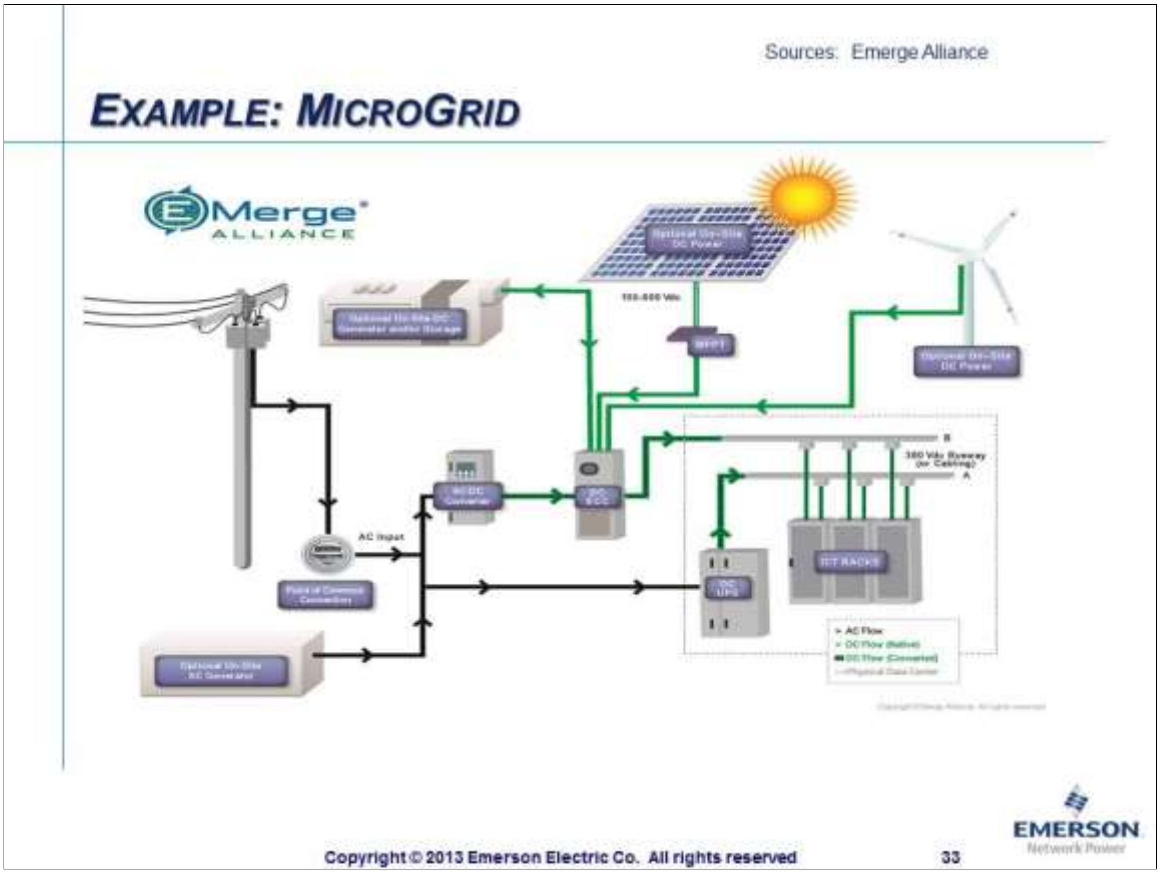

I think, in the summary, power supplies already are very efficient. We know how to do efficient power conversion today. They're already quite smart. But maybe we don't know what to do with them. So much is unclear for our industry regarding how we play in this space. There's a lot of clarity needed around integration communications, standards, etc. Greg, in his talk, mentioned the system value, and I think a number of other presenters did this morning.

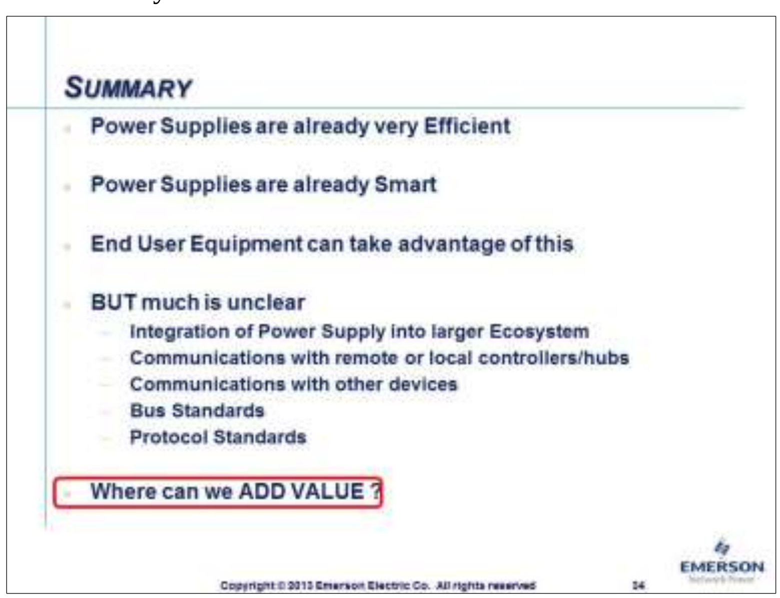

How do we tie all of this together? We're moving beyond just power conversion. Really, what we're asking, is where do we play and where can we add value?

Thank you.

System Optimization

DUSTY BECKER:

Keyue Smedley is our final presenter. She received her Ph.D. in electrical engineering from the California Institute of Technology. She is currently a professor in the Department of Electrical Engineering and Computer Science at the University of California at Irvine, the Director of the UCI Power Electronics Laboratory, and a cofounder of One-Cycle Control, Inc.

Dr. Smedley's research activities include high-efficiency dc-dc converters, high-fidelity class-D power amplifiers, single-phase and three-phase PFC rectifiers, active power filters, inverters, V/VAR control, energy storage system, and utility-scale fault current limiters. She is an inventor of One-Cycle Control and the Hexagram power converter.

Dr. Smedley's work has resulted in more than 150 technical publications, more than ten US/international patents, two start-up companies, and numerous commercial applications. Dr. Smedley is a recipient of UCI Innovation Award 2005. She was selected to be an IEEE Fellow in 2008 for her contributions in high-performance switching power conversion. Her work with One-Cycle Control, Inc., has won Department of the Army Achievement Award in the Pentagon in 2010.

DR. KEYUE SMEDLEY:

I am honored to have this opportunity to and share ideas with you and I am very excited about this workshop. I must compliment the organizing committee for their diligent work. They have crafted this special package with scope and depth that touches many critical and important areas for the smart grid of the future.

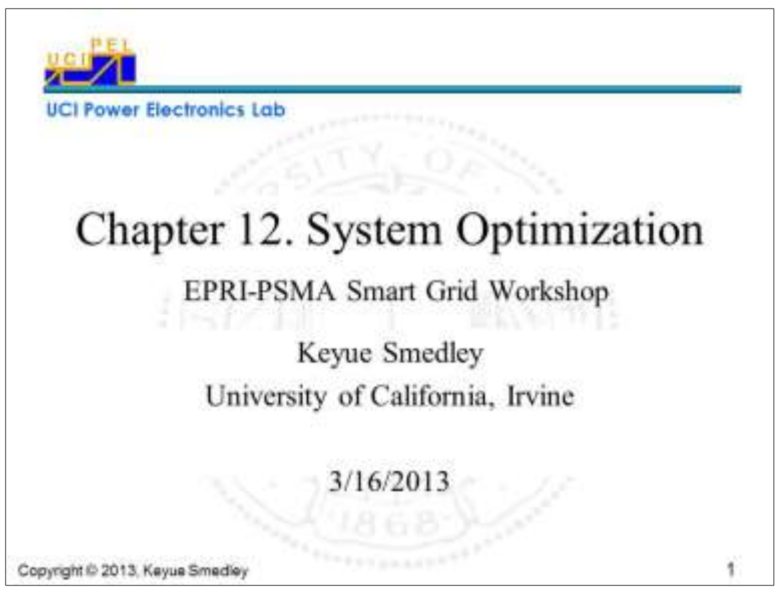

My topic today is system optimization… how to make the grid smarter.

I am from the University of California Irvine, which is about 30 ~ 40 miles from here. So this is a really nice venue for me: easy access.

Power electronics enable the smart grid!

Today I will discuss power grid challenges and opportunities, as well as a vision for a smart grid. Advanced power electronics enable a smart grid. I'll elaborate on congestion control, voltage control, and demand control. It is crucial to reduce power electronics costs for building a power electronics super highway.

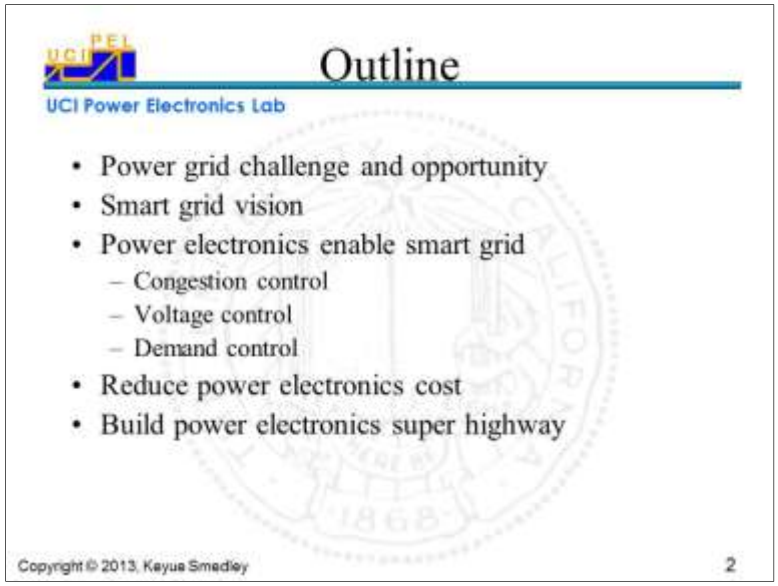

The U.S. power grid is one of the largest machines made by humankind. It has about 700,000 miles of transmission lines and about 5 million miles of distribution lines.

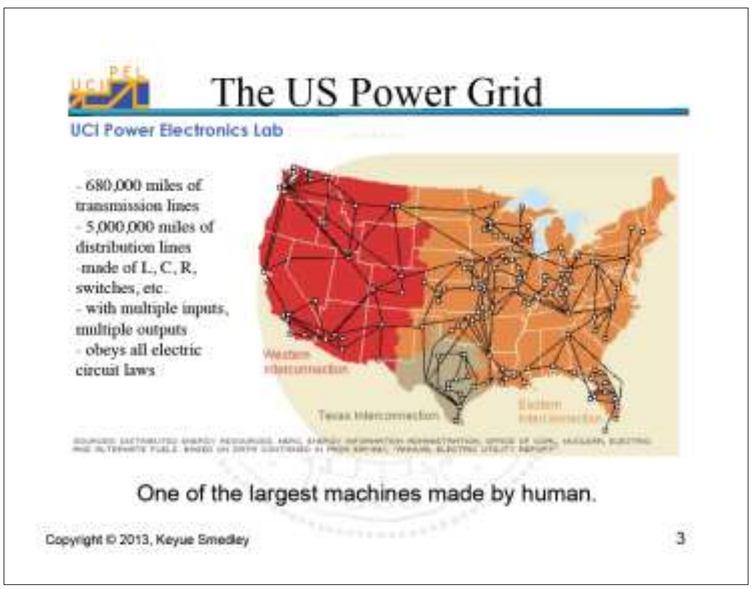

First, our power grid. The U.S. power grid is probably one of the largest machines made by humankind. It has about 700,000 miles of transmission lines and about 5 million miles of distribution. Don't be intimidated; it is really just a circuit with inductors, capacitors, resistors, switches, sources, and loads. From a system point of view, it is a multiple-input and multiple-output system with a lot of nodes. That is an understatement - millions and millions of nodes, if we consider each home is one. It is indeed huge.

The grid is really just a circuit with inductors, capacitors, resistors, switches, loads, and sources; and it obeys all of the laws of circuit physics.

From a circuit point of view, it obeys all of the circuit laws, such as Ohm's law and Kirchoff's law. It is manageable. But how to tackle the scale, the giantness, the complexity… that intimidates people.

The grid started from separate, individual, small circuits; then grew into this giant, interconnected machine over the course of a century.

Grid infrastructure updates are lagging the demand growth. The direct consequence is reduced operation margin.

In earlier days, when we build it, the grid had a lot of margin… thus it had stability. In recent years, however, congestion and blackouts have become more frequent. This is because the infrastructure updates are lagging the demand growth. The direct consequence is reduced operation margin, which pushes the grid to operate at its hairy edge. This also stresses our friends at Cal ISO; who are struggling to keep the grid out of trouble.

Tighter control of the system; can help keep the system operating in a safe region. A more accurate model and smarter control can extend the life of the present infrastructure for many more years to come. This requires a more accurate model and smarter control strategy.

The traditional energy supplies for the grid generate a huge amount of CO_2… directly contributing to environment issues.

We need to be responsible. We want to use more renewables, such as wind and solar.

We have all witnessed serious environmental changes. We need to be responsible. We want to connect as many renewables, like wind and solar, to the grid as possible and as quickly as possible. The problem is that the present grid can't take it. Traditionally, we had centralized controllable generation and predictable loads, so it was not too difficult to manage demand and response.

But with a lot of renewables in the grid, the situation is quite different. Dealing with the wind and solar power is like dealing with wild animals, as these energy sources are distributed and vary minute by minute, second by second. To harness their energy, the grid has to maneuver dynamically to smooth out the intermittency. The present system can tolerate a small percentage of this, but isn't equipped to handle a high penetration of renewables. We need to renovate the system.

When there are a lot of renewable sources on the grid, the situation is quite different. Dealing with the wind and solar power is like dealing with wild animals; they are widely distributed and vary minute by minute, second by second. They cause the grid voltage to fluctuate at the point of connection, which can be driven out of spec. To harness renewable energy, the grid needs a highly maneuverable and dynamic control capability to smooth out the intermittency. The present system is not equipped to handle a high penetration of renewable. This is a pressing issue to renovate the system.

The line loss is typically 8%. So line loss, alone, in the United States is about $26 billion a year.

Let's look at some challenges. First, line loss. There is typically about an 8% line loss from a generation station to end users. So this loss, alone, is about $26 billion a year in the United States.

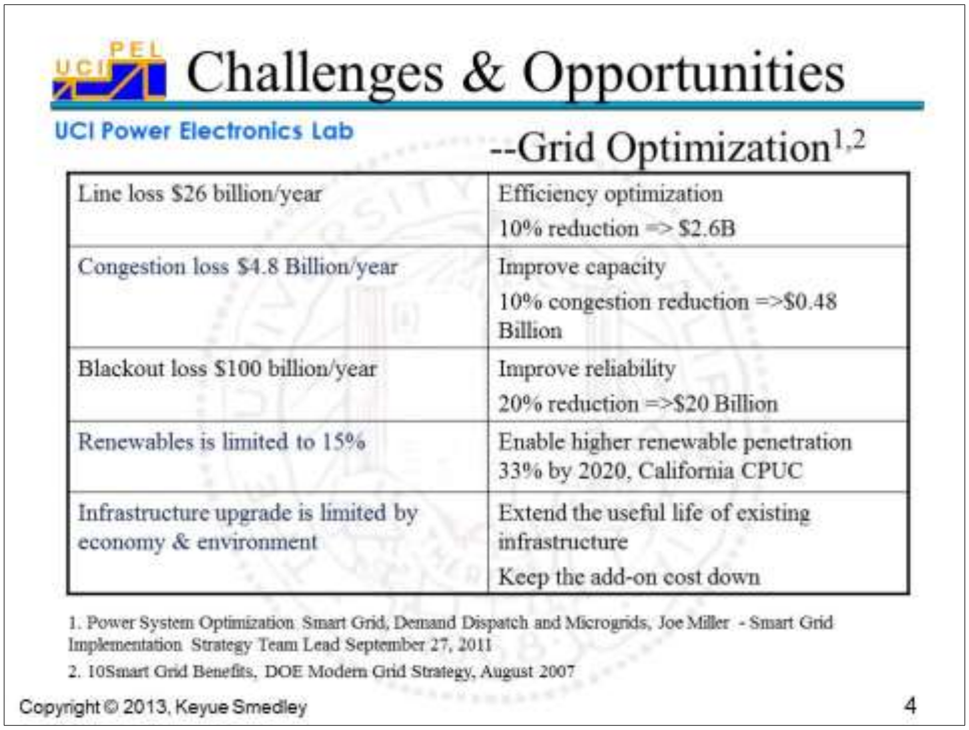

Line loss $26 billion/year	Efficiency optimization 10% reduction => $2.6B
Congestion loss $4.8 Billion/year	Improve capacity 10% congestion reduction =>$0.48 Billion
Blackout loss $100 billion/year	Improve reliability 20% reduction =>$20 Billion
Renewables is limited to 15%	Enable higher renewable penetration 33% by 2020, California CPUC
Infrastructure upgrade is limited by economy & environment	Extend the useful life of existing infrastructure Keep the add-on cost down

1. Power System Optimization Smart Grid, Demand Dispatch and Microgrids, Joe Miller - Smart Grid Implementation Strategy Team Lead September 27, 2011
2. 10Smart Grid Benefits, DOE Modern Grid Strategy, August 2007

Congestion loss. This is very visual. When you drive down the highway and you're stuck there... you've lost your productivity by just parking.

A similar thing happens to the power grid.

Imagine when one corridor is heavily loaded. It cannot take the massive wind power out of the wind farm and transport it to the end users. The only thing you can do is curtail it. But then the wind power owners might not be very happy, because that means loss. The loss due to congestion is estimated at about $4.8 billion per year and may actually be more.

Blackout loss is about $100 billion a year. That's huge!

Blackout loss. It is about $100 billion a year. That's a big number; a huge impact on the economy. As I mentioned before, the power system infrastructure is stressed. Due to the economic recession, we are not doing many upgrades because the cost is too high. Moreover, due to environmental concerns, we cannot afford to populate our land with transmission lines and power generation stations.

On the other hand, connecting more renewables to the grid puts more stresses on the already stressed system. When the grid is stressed and operates in the hairy edge of failing, we really do not have much room to maneuver. A small transient can tip the grid to instability and blackout. These are some of the challenges in front of us, but these are also opportunities to for us to optimize our power grid.

If we optimize the system for efficiency; 10% of that is $2.6 billion! If we optimize to reduce congestion, 10% is $0.5 billion. By optimizing the stability to reduce a 20% blackout, we can save $20 billion!

The CPUC of California mandated that we have 33% of renewables by 2020. However, utilities push back; they don't want more than 15%.

The CPUC of California mandated that we need to have 33% renewables by 2020. These are opportunities. It is important to extend the useful life of existing infrastructure and keep the add-on costs low for upgrades.

By looking at the landscape, we see tremendous opportunities in front of us on all layers: manufacturing, power electronics, research development… there is a chain of opportunities to optimize the system.

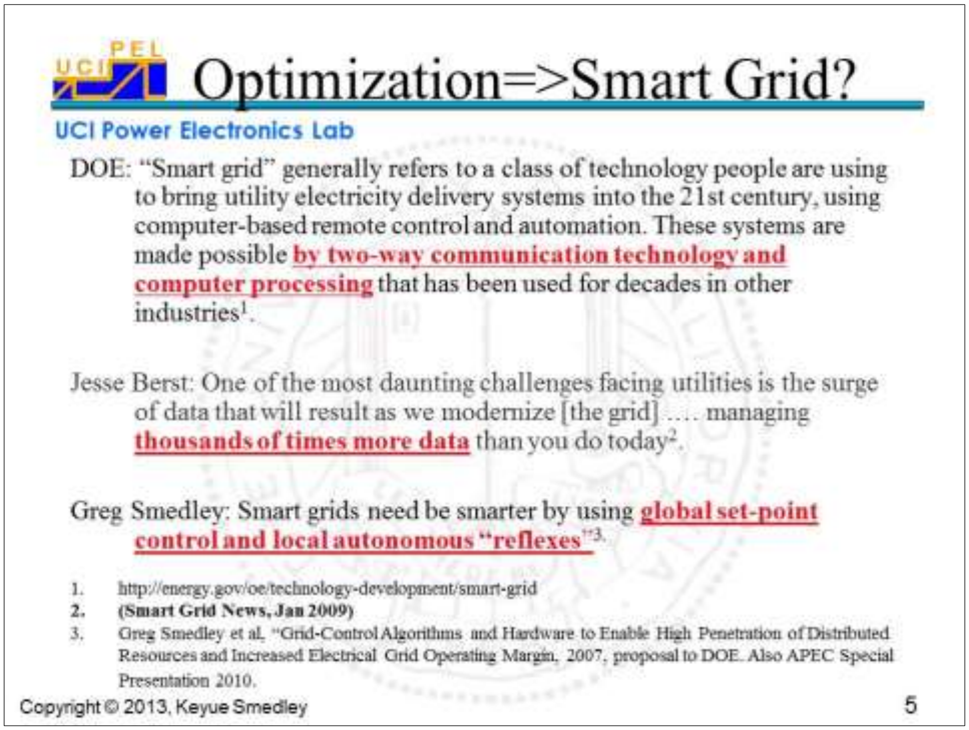

So how to optimize the grid… Right away people point to the word "smart." The media is talking about it because it's hot. Researchers are excited about it because there are funding opportunities there. Smart is a buzzword that catches attention.

Who knows the original intent… but people like smart phones, smart cars, smart cities, smart everything! It's time for change. The world is entering a smart age; everything needs to be smart… our grid too.

According to the DOE, two-way communication and computer processing make the grid smart.

What is a smart grid? DOE believes that two-way communication and computer processing make the grid smart. This idea excites IT professionals and IT industries: great opportunities to take a lot of data at all those nodes, send it to super computers to process, then control the nodes.

I have no doubt that the communication and computer will play an important role in the smart grid. However, one key element is missing in this picture: ultrafast control system actuators. It is important to figure out how to minimize the burden of the massive amount of data.

Let's read what Jasse said: "One of the most daunting challenges facing the utility is the surge of data that will result as we modernize; managing thousands of times more data than we do today."

I think "thousands of times." is an understatement. Large data handling can be a challenging issue and data processing speed is another challenging issue. The electrons traveling at the speed of the light are not going to wait for the data to be processed before they are out of control.

One of the most daunting challenges is the surge of data that will result as we modernize the grid, with thousands of times more data than we handle today.

Dr. Greg Smedley talked about the smart grid concept in 2007: We need set-point control and local autonomous reflexes. The global set-point control may be handled by two-way communication and computation, while the local autonomous reflexes must be accomplished by ultrafast real-time power electronics. Such a system will be smart without the burden of handling gazillions of data.

If we use a global set-point control and local autonomous reflexes, we can make the grid smarter without the need to handle gazillions of data.

First, large data communication and processing does not have the speed necessary for real-time control of electrons. Furthermore, our power grid is a dynamic control system. A control system needs actuators to execute the set-point commands. Power electronics can shoulder the responsibility of actuators due to power handling capability and fast dynamic speed.

I have a smart grid vision. It includes a control room, two-way communication, controllers, actuators, control variables sensing, disturbances sensing, local high-bandwidth feedback loops, optional feed-forward loops.

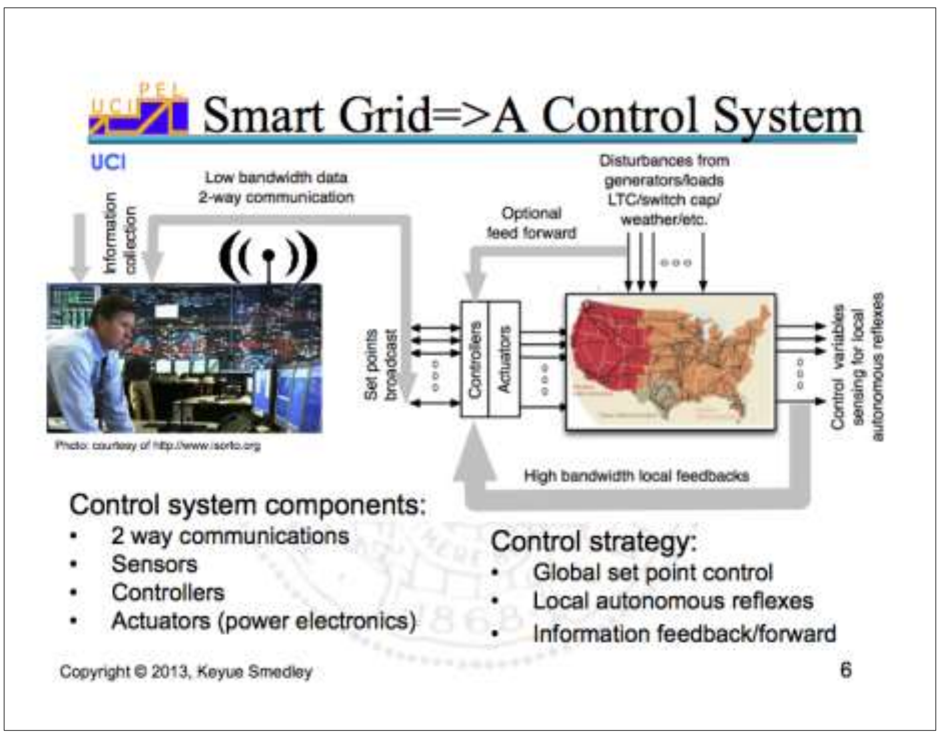

A control room collects grid information and other information (including weather, market, forecast), makes grid-control decisions, and broadcasts the decisions as set points to all or a selected cluster of nodes. The grid control decisions include of global voltage profile, demand profile, and the schedule. The set points are control references that provide commands for control variables at the local nodes. A typical node receives the set points, compares the sensed control variables with the set points, collects optional feed-forward information from disturbances, makes local reflex control commands, and executes the control commands by using the actuators.

The objective is for the nodes to follow the set-point references and to reject the disturbances from generators, loads, load tap change, switch capacitors switching in and out, and weather.

For fast autonomous reflexes, the feedback needs be fast and real time. The information traveling between the control room and the nodes does not need be real time or high bandwidth.

Ultra-fast actuators are an indispensable part of this control system. Why ultra-fast? Only ultra-fast actuators can handle the fast intermittency transient. Are the traditional, load tap change, switching capacitor banks fast enough? No, they are much slower than the intermittency transient and they have limited cycle times before end of life. Power electronics have the speed and accuracy.

This is a great opportunity for the Power Sources Manufacturing Association to promote power electronics technology as an enabler for the smart grid. We need it to connect renewables to the grid; we need it for energy storage; we need it for voltage control, frequency control, power control... It is really a powerful thing in our hands.

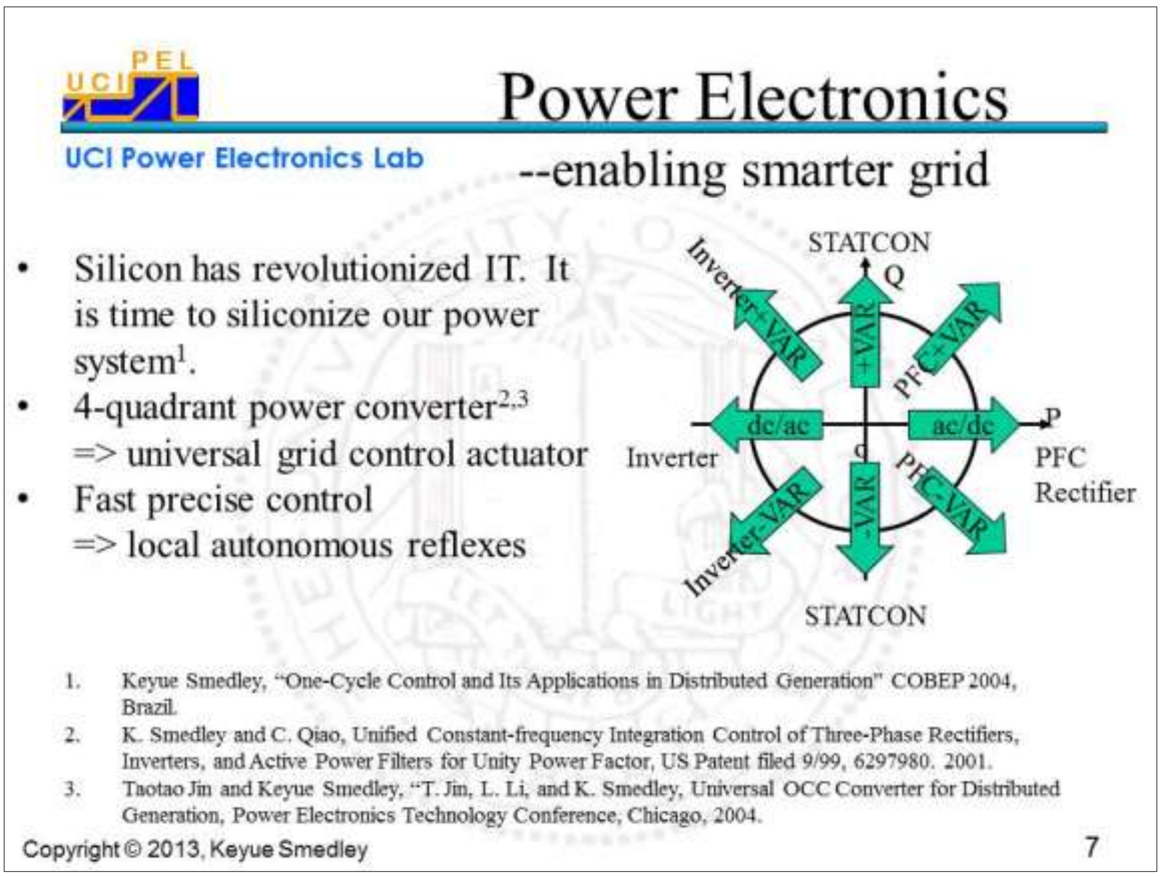

In the last century silicon has revolutionized IT. Now it is time to "siliconize" our power system. Before the arrival of Intel®… we had many different computer platforms: giant, slow, and expensive. It was hard to image that one day we could have personal computers. Now, the number of smart phones, with computation capability way beyond the old main frame computers, is going to outnumber our population.

What made the IT revolution successful? Intel came up with a chip; the microprocessor, which standardized the computer architecture with a software layer on top of that. This architecture can be applied to servers, desktops, laptops, smart phone. Standardization is the key! It dramatically improves the performance, reliability, productivity, and cost.

A four-quadrant power converter unifies all the grid actuator functions we need.

Similarly, we need to standardize power electronics. Right now, power electronics manufacturers are doing different things. Different power converters were created for different functions and are generally giant, slow, and expensive. A fast and precise four-quadrant power converter is what we need to unify all the functions for realizing all different local autonomous reflexes.

A four-quadrant power converter has P and Q reactive power dimensions. If the power flow is along the P direction, we have a PFC rectifier. If the power flow is along the -P direction, we have an inverter. If the power flow is along the vertical axis, we have Q control, that is a VAR supply or STATCOM. The four-quadrant control capability enables a power converter to operate at any point inside or along a circle in the PQ space, limited by the current rating. This kind power converter is universal.

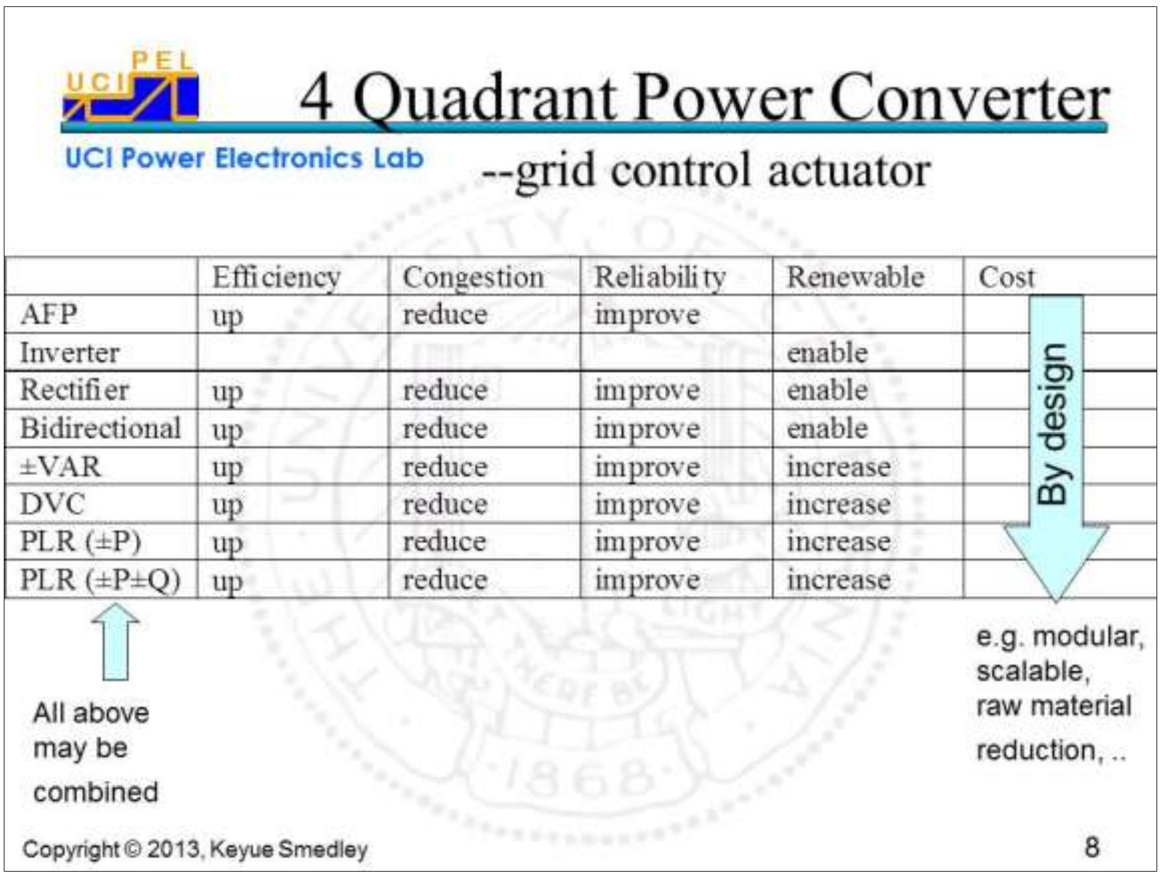

	Efficiency	Congestion	Reliability	Renewable	Cost
AFP	up	reduce	improve		
Inverter				enable	
Rectifier	up	reduce	improve	enable	
Bidirectional	up	reduce	improve	enable	
±VAR	up	reduce	improve	increase	
DVC	up	reduce	improve	increase	
PLR (±P)	up	reduce	improve	increase	
PLR (±P±Q)	up	reduce	improve	increase	

With a four-quadrant power converter, we can actually make a product, such as an active power filter, inverter, rectifier, bi-directional power converter (ac to dc, dc to ac), a VAR control (plus and minus Q), a dynamic VAR compensator, a peak load regulator with positive and negative power flow, or peak load regulators with plus or minus P and plus or minus Q, regenerative ac or dc loads.

With a standardized four-quadrant power converter, we can make an active power filter, an inverter, a rectifier, a bi-directional power converter, a VAR control, a dynamic VAR compensator, a peak load regulator, or regenerative ac or dc loads.

With this kind of capability, we are equipped to optimize efficiency, congestion, and reliability, as well as enable the expansion of renewables. In addition, modular, scalable, power converter design reduces cost further by high-volume production and optimized raw material usage.

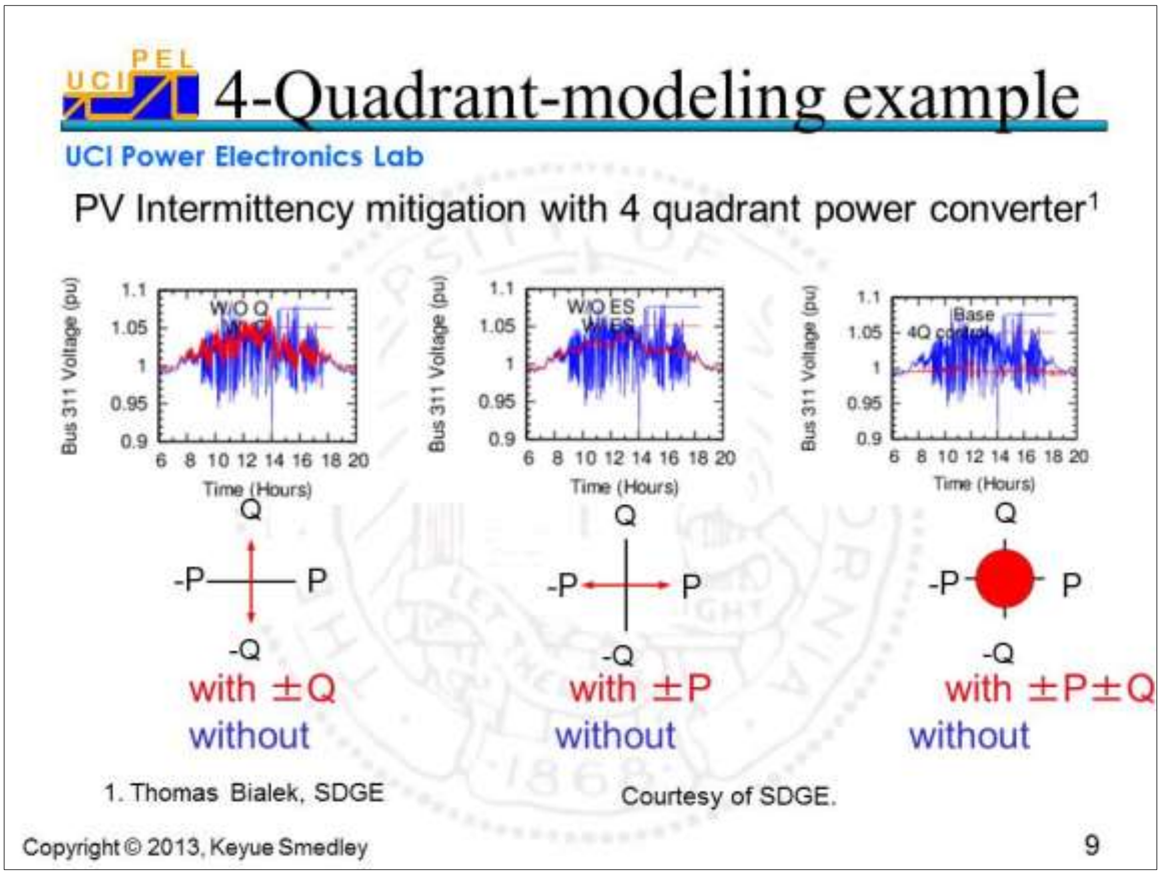

People in utilities have dreamed of fast four quadrant converter for grid control. Here is some simulation work from SDGE, aiming to solve renewable intermittency. The blue curve is solar power voltage, not constant, but spiky and all over the place, due to intermittency. By adding fast-plus and minus-Q to the system, the voltage changes from the blue curve to the red one, which is smoother (left plot). The impact from intermittency is greatly reduced.

By using some storage to facilitate plus and minus P control, the voltage profile is even smoother (center plot). With both plus and minus P and plus and minus Q, the voltage profile can be pretty much constant (right plot).

If you have some intermittent generators, such as solar or wind, and dynamic loads; a four-quadrant power converter can make them look like model citizens in the grid.

This is the powerful. If you have an intermittent generator, such as solar or wind, the four-quadrant power quadrant converter can actually make it look like a model citizen: tame, beautiful, and obedient. If you have very dynamic loads, such as an EV station with fast battery charging, the four-quadrant converter can make the load look like a model citizen: tame, beautiful, obedient.

Such an idea is no longer a dream. It has already become a reality. Ultra-fast and precise four-quadrant converters were demonstrated by One-Cycle Control (OCC), Inc. in 2006 (5 kVA) and in 2008 (40 kVA) under DOE sponsorship.

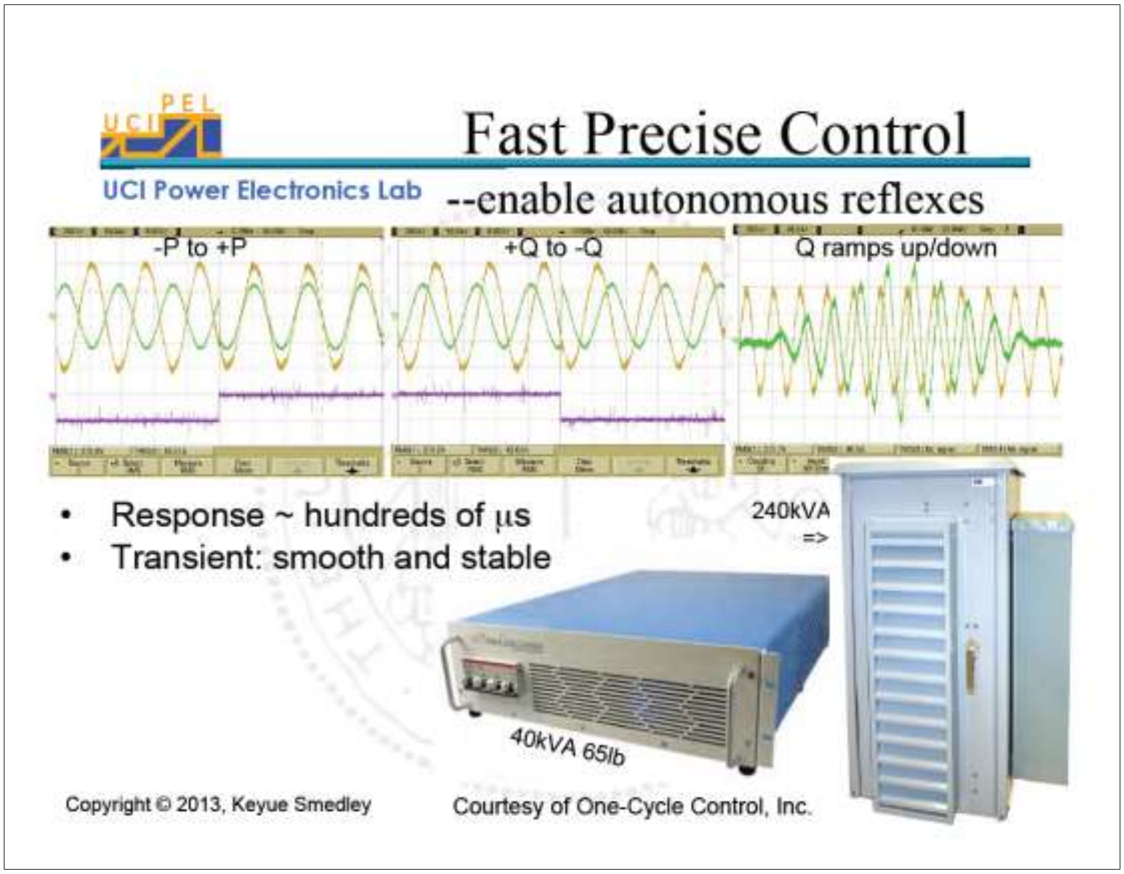

Here is an example of one-cycle control four-quadrant converter, three-phase, 40 kW, 65 pounds. It is about the size of an old desktop.

The power converter can reverse its active power or reactive power flow in a matter of a few hundred microseconds, from P to -P, from Q to -Q. In the plots, the golden curve represents phase-A voltage, while the green one, the phase-A current.

First look at the speed of power control in the left plot. The current changes from 180° out of phase (inverter mode) to in phase (rectifier mode), following the command (the purple curve) in about 100 microseconds.

Regarding reactive, I really appreciate the earlier speakers who explained very clearly what Q is. It makes my job easier. In the center plot, the current changes from 90° leading to 90° lagging, following the command (the purple curve) in about 100 microseconds. In the right plot, the Q value follows the ramping command precisely. The four-quadrant power converter is ultrafast! With this kind of speed, this kind functionality; it can execute the set-point command from a remote control room to achieve local autonomous reflexes for controlling voltage, power, frequency, power factor. The transients are smooth and stable.

Let's look at the modular and scalable features. Six power converters can be configured in parallel to handle 240 kVA, as shown in the rockmount unit. Modular design allows scalable from kVA to MVA.

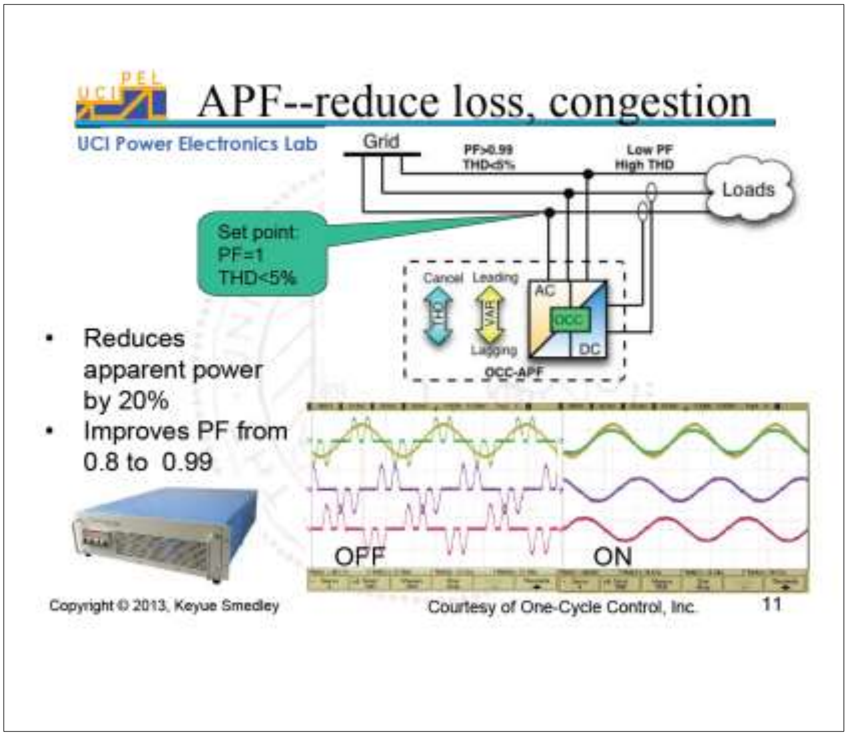

Let's use our quadrant converter to perform local autonomous reflexes for power factor correction. In this case, the four-quadrant converter follows a power factor set point command to maintain unity power factor at the point of connection. The four-quadrant power converter in this example acts as an active power filter in a pilot demonstration. It increases power factor to one and reduces total harmonic distortion to below 5%. As a result, it reduces the apparent power by approximately 20% across the network. It improves grid operation margin while reducing line losses and alleviating congestion.

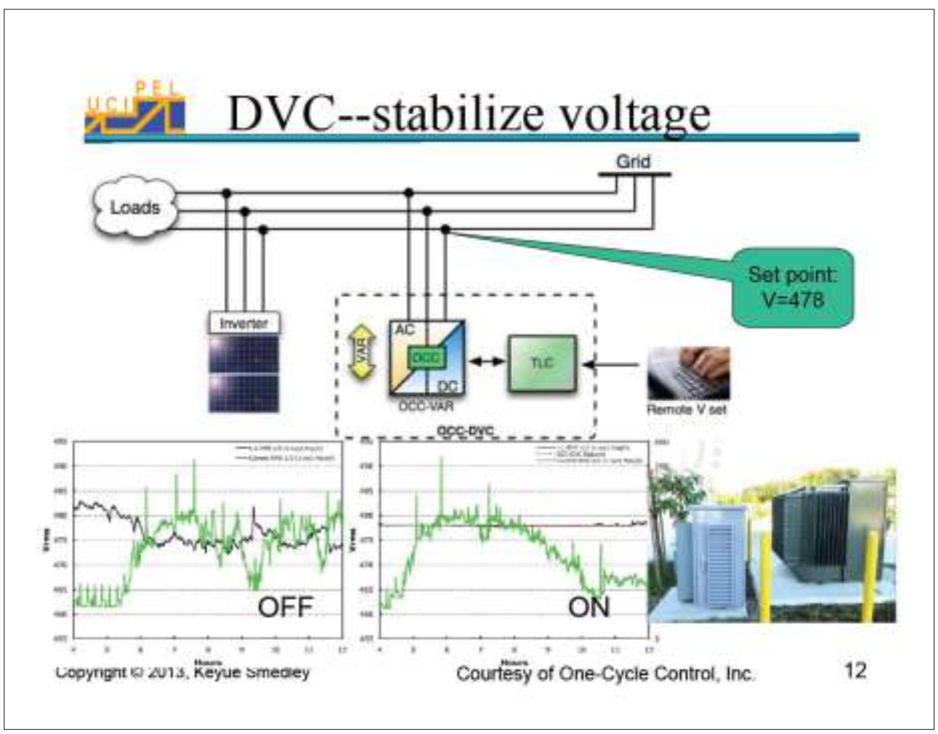

Now I am going to use a dynamic VAR compensator (DVC) for voltage stabilization. Command the set point of terminal voltage to be 478 V and BANG! The voltage is stabilized to 487 V. In the left plot, when the DVC is off, the voltage (black) fluctuates as the current (green) changes. In the right plot, when the DVC is on, the voltage is kept constant no matter how the current changes. The DVC controls the voltage to a straight line. When a cluster of nodes are controlled by such DVCs, the grid voltage profile can be stable and resilient to disturbances.

We have the technology to optimize the grid. With a wide deployment of dynamic VAR compensators, power filters, and energy storage; the congestion goes away and frees up 20% of the transmission capacity and the voltage is stabilized to prevent grid blackouts.

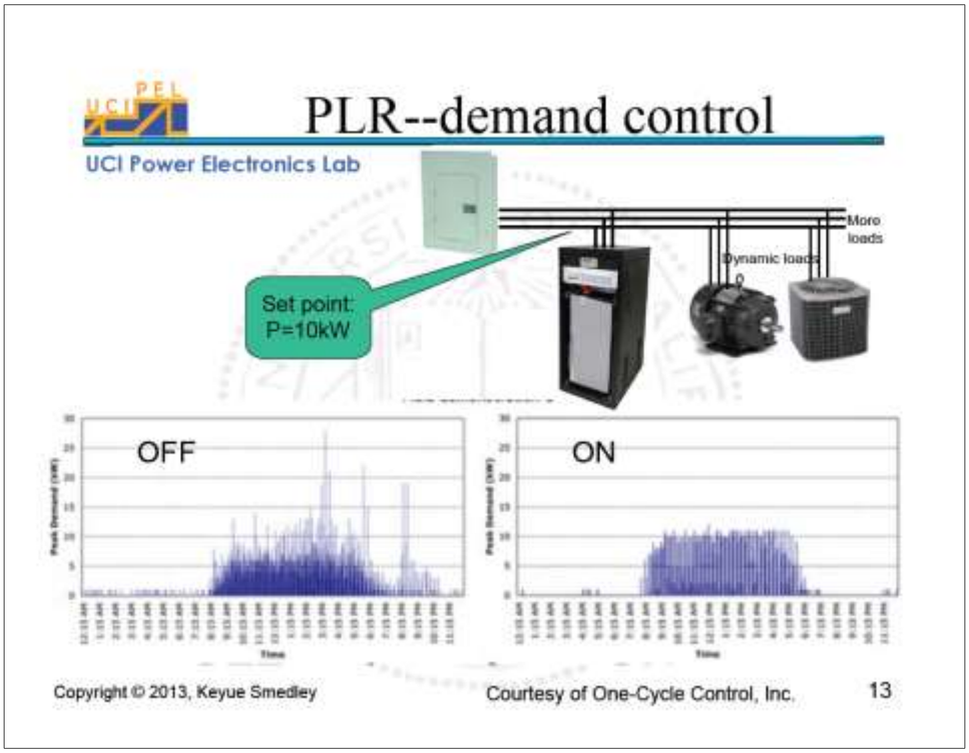

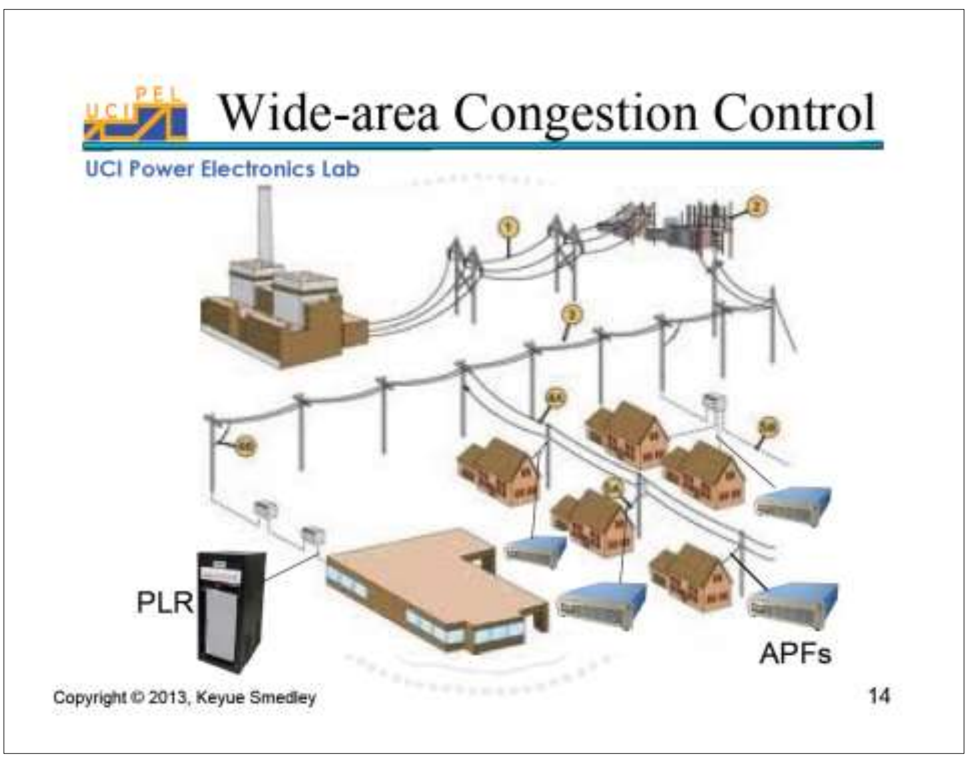

I'm also going to talk peak load regulation (PLR) using a battery pack and a four-quadrant power converter. This is also a pilot demonstration. Without the PLR (left plot), the power flow is very spiky. Utilities charge heavily for this. In California, peak load charge for industrial users can be ~30 – 40%. By using demand control with this OCC-PLR, we can shave the peak like a buzz haircut (right plot). It can cut demand charge.

We have the technology and tools to optimize the grid. Imagine if we deploy the four-quadrant power converters, like active power filters and energy storage PLR, in a wide area. We can do away with congestion and free up ~20% of the grid capacity.

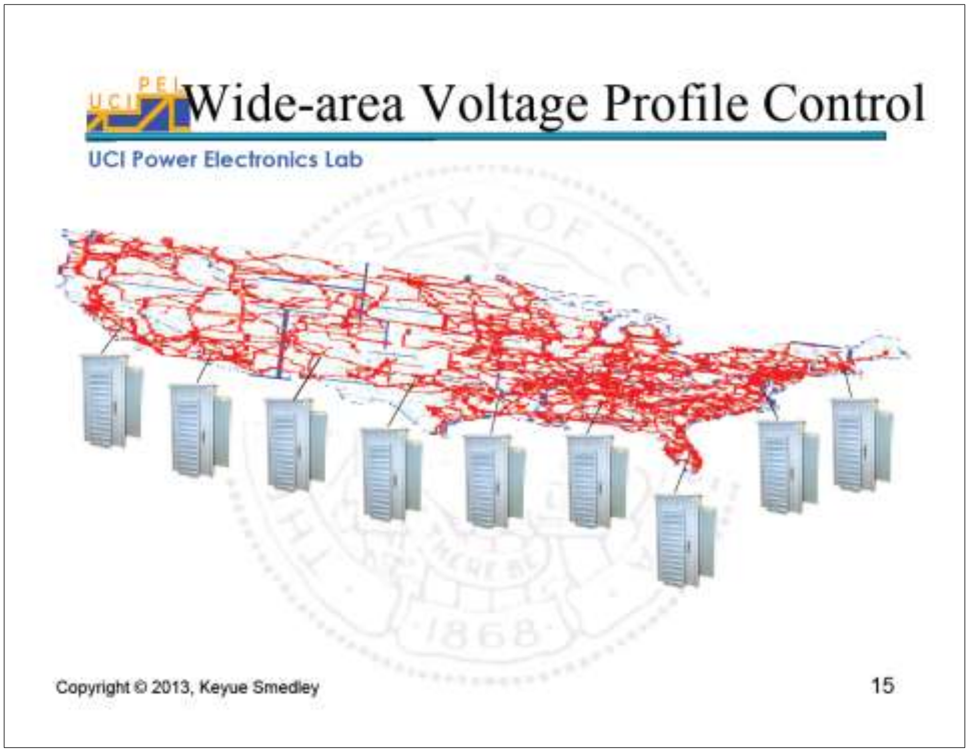

It is possible to deploy DVCs over a wide area for voltage profile control by broadcasting the set point to the nodes to maintain the voltage profile within specifications regardless of disturbances.

If we use weather condition data to predict power production from wind and solar, together with load forecast, we can schedule the voltage profile accordingly, dynamically, as a function of time and geographic location. It does not have to be a flat surface; instead, it can follow a moving set point of control as needed to optimize grid functions such as efficiency and reliability. We have the ability to control this giant machine to be a well behaved adult, instead of an angsty teenager.

We can control this giant machine to be a well-behaved adult instead of an angsty teenager.

I have shown examples of what power electronics can do. An earlier speaker expressed his concern on power electronics costs. Let me address this issue.

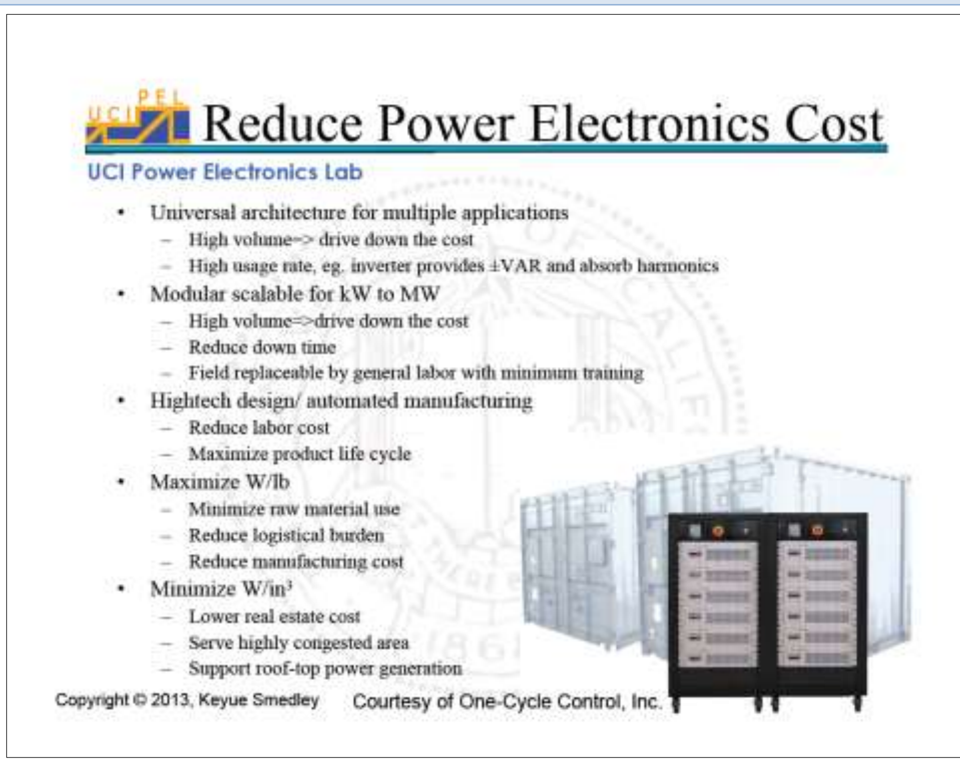

A universal architecture for multiple applications means the same power converter, same design, and same product can do the voltage control, energy smoothing, rectifying, active pulse filtering, or inverting… then the volume can drive the price down.

First, a standard architecture is the key for cost reduction. The same power converter can be configured to serve multiple applications; such as voltage control, power smoothing, rectification, active pulse filtering,… or to be an inverter. The combined production volume can drive down the cost dramatically.

The second measure to tackle cost is a modular, scalable design to cover kilowatt to megawatt applications. This further pushes volume up; pushing production, shipping, and maintenance cost down.

Furthermore, design automation, maximizing power per pound, and maximizing power per cubic inch reduce the cost as well.

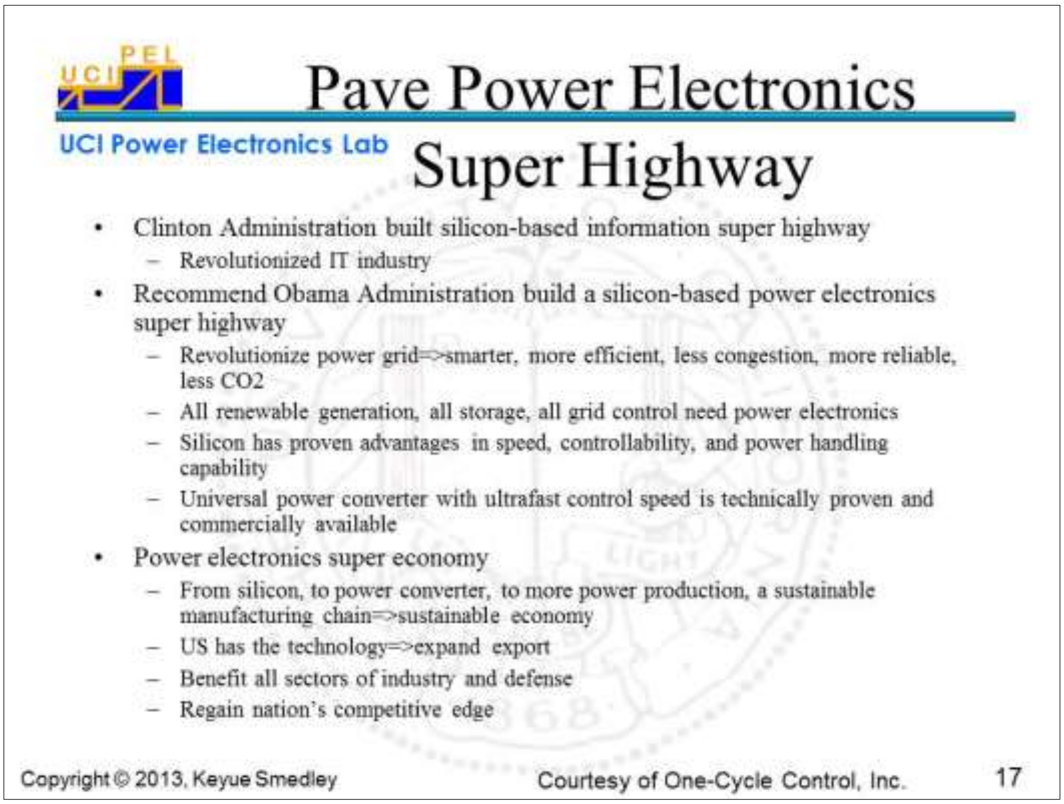

We have a huge opportunity before us, considering how big the power grid is.

The Clinton administration put investment into the silicon-based information super highway that revolutionized the IT industry. I recommend that the Obama administration build a silicon-based power electronics super-highway. Power electronics is an enabling technology for a stronger grid and more renewables. We have the need; we have the technology; we can create the jobs; we can change the energy landscape.

I recommend the Obama administration build a silicon-based power electronics super highway.

Investment in a power electronics super highway is vital for optimizing our present grid, facilitating high penetration of renewables, and creating sustainable industries.

The U.S. has the technology and this technology can benefit all sectors of our industries and our national defense. We can export power electronics products internationally to promote renewable power production, enhance stable electric power supply, and reduce the carbon footprint globally.

Thank you!

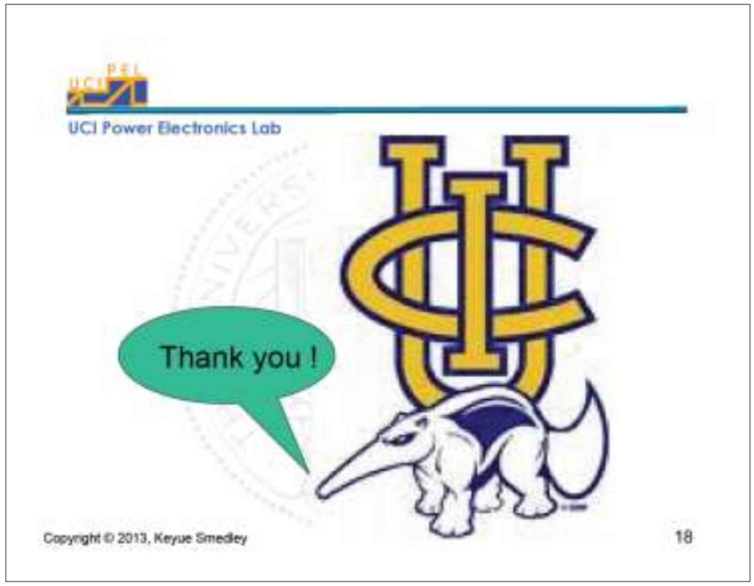

Section VIII
Summary

Summary

CLARK GELLENS

I'm going to be real brief because I have a feeling you're all going to leave if I'm not!

I have a very few slides here just to help me summarize.

This is an "executive" summary, I guess that's because I used to be an executive and I used to summarize.

So are you ready for the smart are grid? Did you learn a little something? I think this was a great session. Really wonderful speakers. I look forward to the summary of the material.

I think we made enough points about how complicated this is all going to be. In particular, at the edge of grid, with the interface being communication, sensors, and computational ability with the grid.

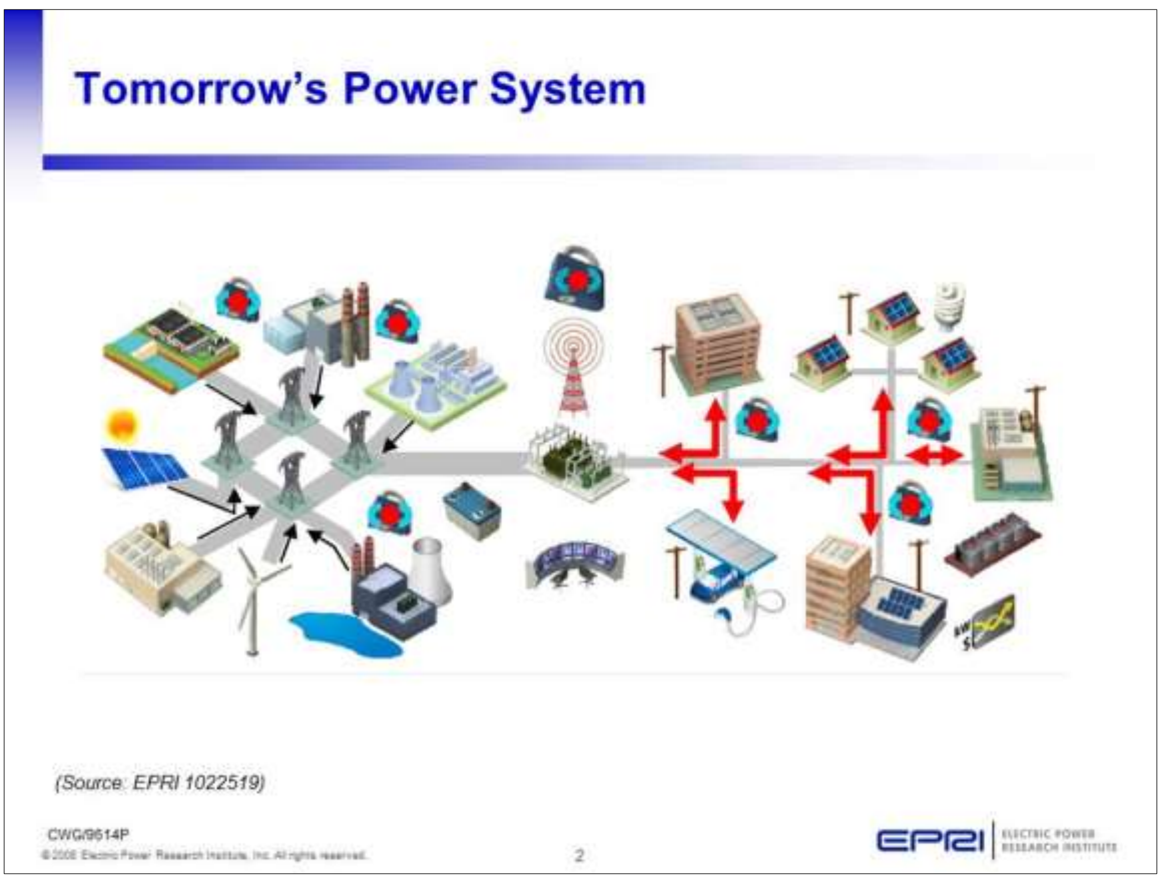

Phenomenal opportunities, driven, in part, by all this technology… there's a lot of opportunity for power electronics.

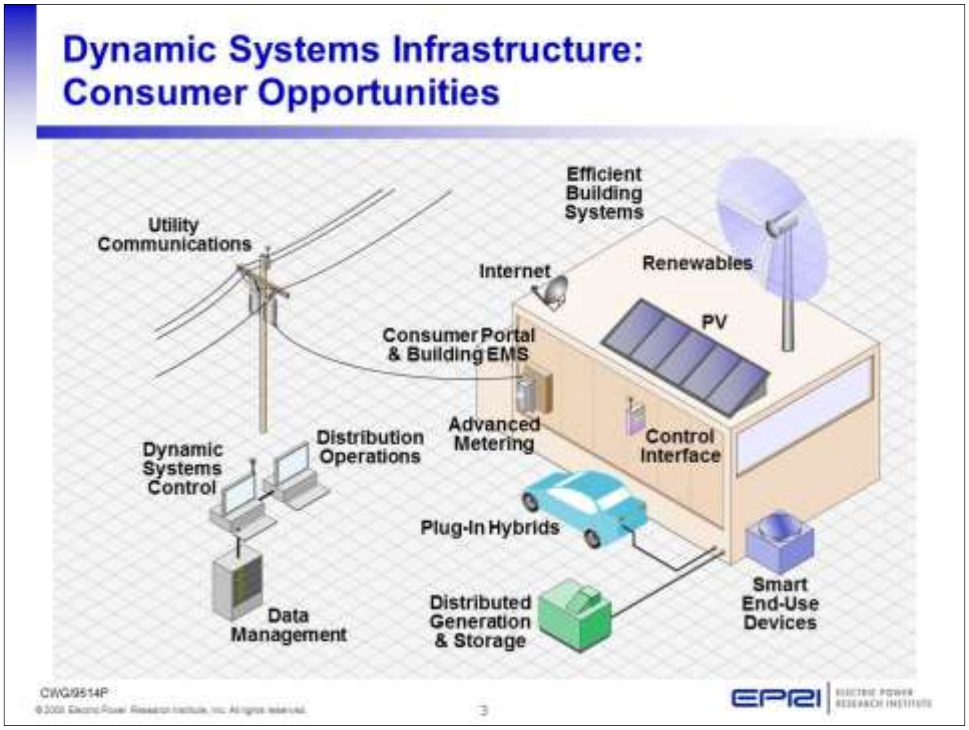

If I start just looking at these pieces of little red stars, you've got power electronics opportunities all over the place. Those of you who are in this business ought to be thinking about a greatly enhanced and expanded application.

Some of these we didn't even mention, some of these energy devices will need power electronics. We can store electrically, but I can store thermally as well and balance the two together. There are a lot of opportunities doing that.

If I move up in the power system, maybe to a distribution level or on a campus; I've got, again, a lot of opportunities for power electronics. The same with the bulk power system; that's more obvious where power electronics can play a key role. Several points have been made by previous speakers in that regard.

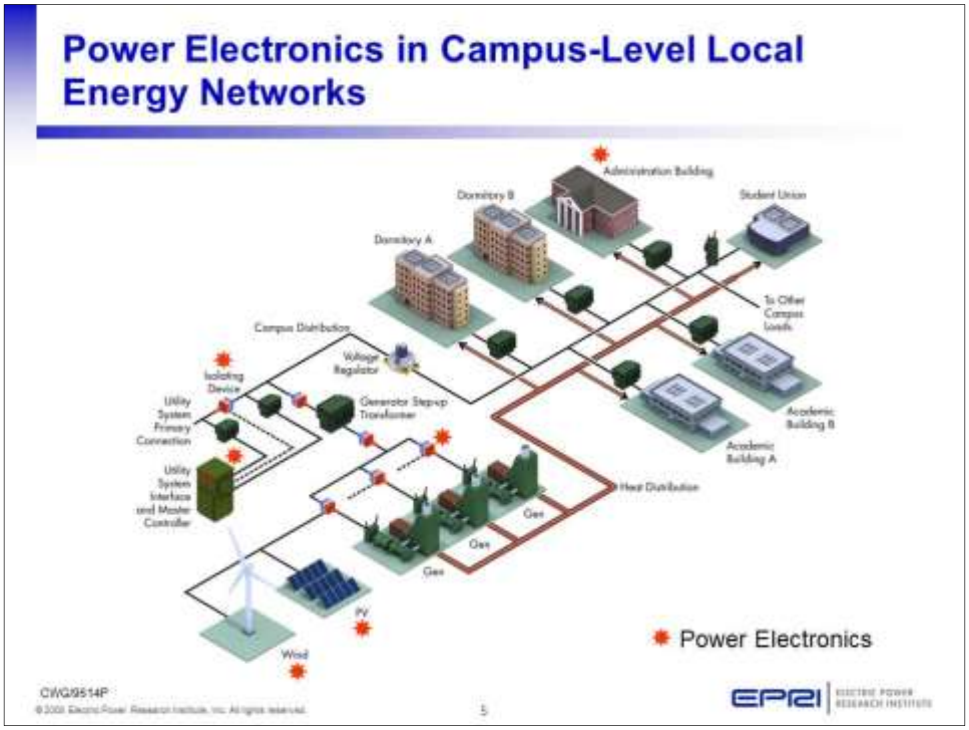

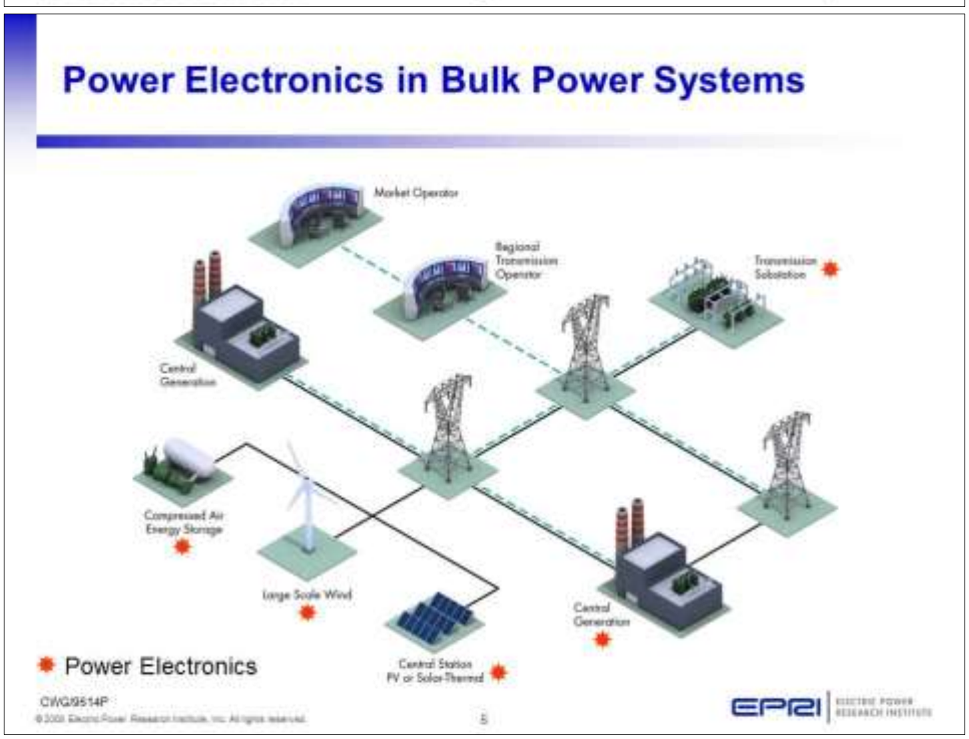

Add on top of that, all of this activity going on and we've only barely touched on that. Dennis Symanski did a great job in at least highlighting some of them: sensor data integration and analytics. It's not just smart grid, it's also smart generation. It's smart management. Diagnostics, managing the assets, is a huge area that is not well managed.

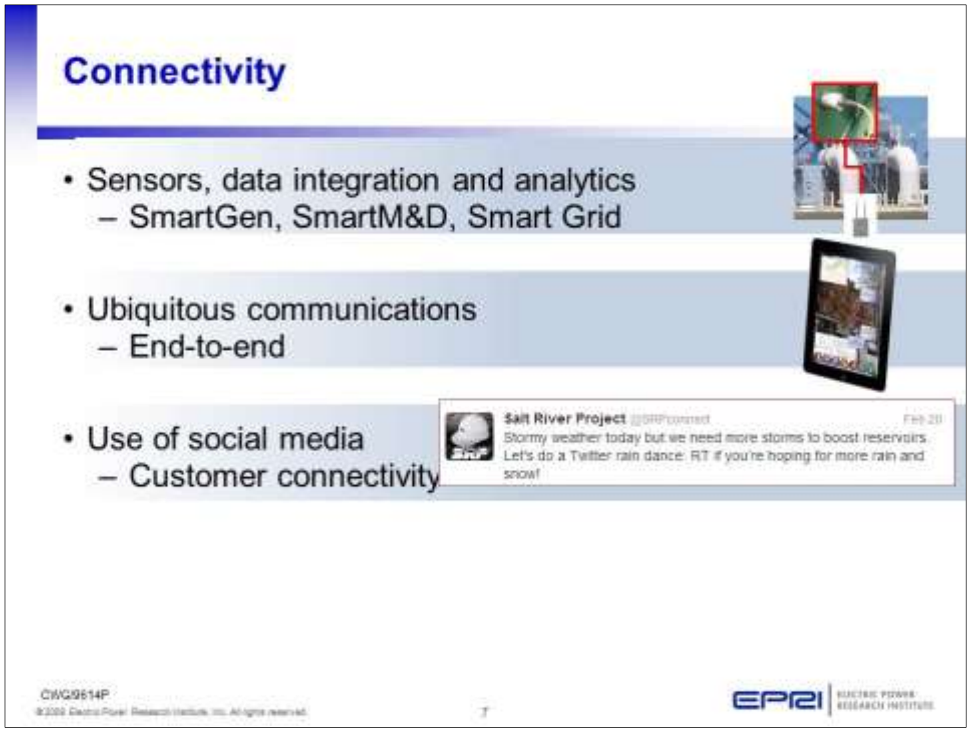

Managing power transformers' 40-year life was mentioned, but we didn't discuss that the average substation transformer in the U.S. is 42 years old. How do you fix that problem? You use a lot of this stuff that we've been talking about. This ubiquitous communication, end-to-end with the customer and with Cruiselift? This mutual assistance program that we use in our industry. You know, Sandy comes, and we've got 17,000 trucks that come in from other parts of the country to help out.

What do we want to do tomorrow? We want to hand them a tablet and, on that tablet, we'll have all of the circuit diagrams and all of the information they need. They'll have the GPS, they can take a picture, literally, of the street and end up with a list of materials, dispatch it, send it to control center, and have a truck come meet them with all of the equipment they need to rebuild that street. That's going to happen in short term.

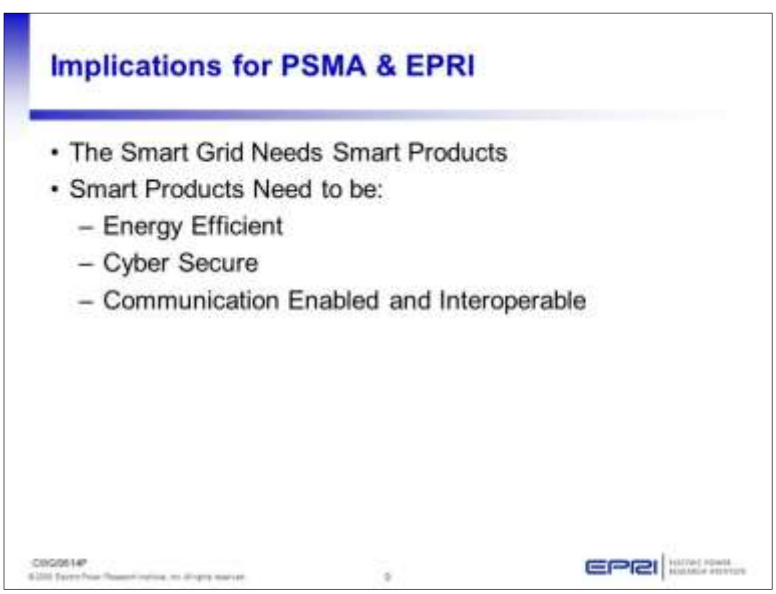

What I can't tell you, is how far this is going to extend with customers. I have no idea. I don't know whether Bill Gates is right when he said that everything tomorrow will be computer... even the walls in our buildings will be computers.

Think about how quickly this is all going to move! What does it mean for you and I in this business? I don't know… maybe even the use of social media for connectivity for the little stuff, like talking to the customer on behalf of the utility.

So what are the implications for you guys and maybe EPRI, since we have a joint workshop going on here? The smart grid needs smart products. Those smart products need to be energy efficient, cyber secure, communication enabled, interoperable, and probably a couple of other things. They're going to have to have that capability in the home and the office back to the utility: two way! They're going to have to enable energy management, demand/response, and all these other attributes we've referenced.

I wonder if it's time to establish a joint task force. The smart grid needs smart products. Maybe, together, we can outline, go over some guidelines, discuss what attributes those products should have…

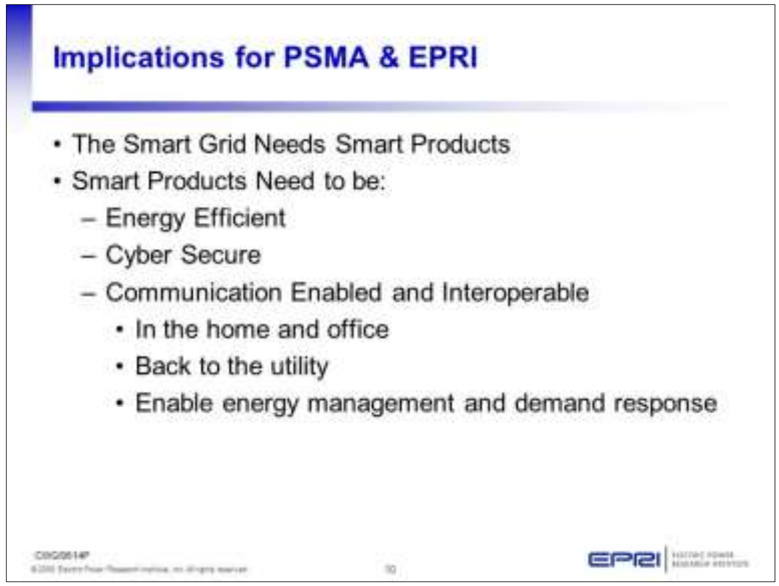

Panel Discussion

QUESTION

Is there any chance the demonstration you showed will happen in Northern California anywhere? Remember the UC Irvine plan of zero energy?

CLARK GELLENS

Well, that's out of the DOE funding and that's going to be public knowledge. It's just being implemented right now. They're just starting to put equipment out at the UC Irvine faculty housing and all around campus. Because it's DOE funded, it's all going to be public knowledge.

In Westminster, about 30 miles south of here, there is an advanced technology center. They've got a demo sensor there with a lot of the same technology. Everything from taking a look at the distribution circuits remotely to taking a look at smart converters. They're looking to get the four quadrants running.

QUESTION:

Anything in Sacramento?

CLARK GELLENS:

I'll send Dusty information on that. There are 15 different demonstrations and we've got some executive summaries, public summaries.

QUESTION:

I understand the need for closed meetings, but what about open workshops?

What about 3-phase power in residential areas that is affordable?

CLARK GELLENS

If everybody wants it, it will be more affordable.

We wouldn't have electric motors, but we'd have a lot more copper and a lot more distribution systems.

QUESTION:

Is it not more efficient than single phase?

CLARK GELLENS

It depends. You can't equate efficiency with cost.

I did an analysis of 25 KW demand. You might say that the overall use is efficient because you use less space. There are air conditioner motors and so on and so forth. Anything below that, no.

In Europe, they have a return path ground and that makes their power system cheaper and brings in stray voltage.

Section IX
Chairman's Remarks

PSMA Co-Chairman’s Remarks and Workshop Analysis

EDWARD HERBERT

Co-Chairman, EPRI-PSMA Workshop
Co-Chairman, PSMA Energy Efficiency Committee
Member, PSMA Board of Directors

The workshop brought together an amazing collection of talent. Careful review of the presentations reveals a lot of information. Analysis of the workshop material reveals much more. Remaining unanswered questions will be fodder for future work.

It is our belief, as power converter designers, that our equipment has capabilities that would be very useful to “The Grid,” if only the grid operators knew to ask. On the other hand, we do not know enough about “The Grid” to know what to offer. This is not helped by the hubris of engineers – we tend to believe that we have all the answers and do not need advice from others.

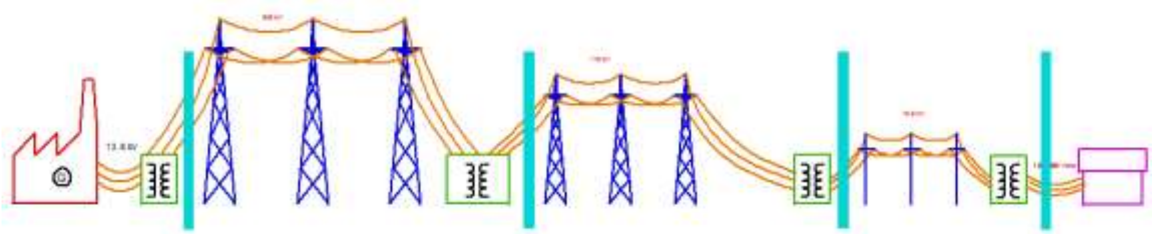

The intent of the workshop is to study the interface between “The Grid” and power converters, looking ahead to the “Smart Grid.” By applying engineering analysis to problems revealed in the workshop, we may have identified a number of opportunities for PSMA members to enhance their products and help to stabilize the grid.

The grid is generally divided as pictured above – bulk generation, bulk transmission, sub-transmission, distribution, and the end users. There are several functional divisions within the grid – it is remarkable how similar the boundaries are, generally as defined by the vertical lines in the sketch.

PSMA’s turf: Most of the equipment made by PSMA members is at the right end of the line. Presently, it is mostly loads, but increasingly it will include inverters for alternative energy and energy storage. As GaN and SiC components become available, “solid-state” compensators will be used more widely. All are marketing opportunities for PSMA members. We have relatively little interaction with the larger equipment on the left, above, but accessories and controls that enhance its operation and/or prolong its useful life are possible applications for PSMA members.

“The Grid’s” turf: There are many issues surrounding the operation and control of the grid. We can facilitate control and compensation by providing well-behaved active loads and distributed power sources. A concern of PSMA members is that we not interfere with the grid,” such as by being a conduit for hacking or injecting noise.

Problems and Solutions

We consider the various problems discussed in the workshop and try to find common solutions that can be implemented in controlled power sources and loads at the distribution level. A common solution may be improved impedance control. Power factor corrected power converters already implement impedance control to make the input "look resistive" and the control algorithm can be modified to optimize a "synthetic impedance."

Analysis suggests that controlled power sources and loads should have a synthetic series impedance of about 2 %R + 10 %X, with controlled modification for emergency conditions. Percent impedance %Z may be unfamiliar to PSMA members, but it is how the utilities characterize the impedance of transformers and it is a very useful tool.

Inverter failures. The extraordinary failure rate and poor deployment success for energy storage inverters suggests that stability issues are a serious problem between nearby power sources. Additional buffering between power sources is needed and a strong synthetic reactive source impedance may provide a solution that does not increase losses and enhances voltage regulation.

Voltage, frequency, and load lines. There is a very strong and linear relationship between voltage and phase shift when adding vectors, which is what ac compensation and power control is all about. The effective voltage, Vr, across an impedance between two voltage sources is $Vr \approx \sqrt{\Delta V^2 + (V_a * \theta)^2}$; where θ is the phase angle in radians. This has several implications. One is that phase shift between voltage sources is just as critical as voltage differences, but it is harder to control. However, once the phase shift control is quantified and managed, it becomes a powerful tool for managing power.

A key is to have strong load-lines responsive to voltage and phase shift, and that is determined by the series impedance of power sources and loads. In the distribution system, it is desired to have tight voltage control, thus the relative small 2 %R. Phase difference is transparent to most legacy equipment, so the much higher 10 %X is practical.

Compensation issues. The compensation for varying loads requires a thorough analysis to understand the problem. Current compensation is best done at the load and a reactive source impedance provides a solution. It is important to limit di/dt.

Resistors and induction motors. The grid is presently stabilized in part by resistive loads and induction motors. Both are being phased out. Simulating the input characteristics of induction motors in loads provides both voltage and frequency stability.

"Negative resistance." The negative resistance (more correctly, negative dynamic resistance) of "constant power" loads means that current is not a suitable parameter for voltage control. All controlled loads have negative resistance at lower frequencies with an impedance inversion at time constant of the control. Power and VARs are preferred for control.

Why have a grid? Why an ac grid?

The grid exists and will remain because the best sources of power are not located where power is consumed. Value is a consideration and "cleanness" and "greenness" will become increasingly important. If power is less expensive, cleaner, and/or greener at a different location, we need "The Grid" to transport it from source to consumer.

The "war of the currents" rages still, with many proponents of a dc grid. However, the ease of shifting voltage in transformers remains a compelling reason for ac.

Another factor, less recognized, is that the ac grid allows power flow between nodes of nearly equal voltage by controlling the phase difference. Proponents of dc power distribution often cite the problems of synchronizing different sources of power. When differences in phase angle are quantified and controlled, they become powerful tools for power management that are not available with dc distribution.

Energy storage and power noise

If "zero net energy" buildings become the rule, energy storage and power noise control will become increasingly important. Solar panels are the most likely energy source for private use in populated areas, and they produce intermittent power having variations due to cloud movement, daily cycles, and seasonal variations. The grid must be able to handle it.

If power production for the whole day is concentrated between 9:00 am and 3:00 pm, peak production will reach 400% of average consumption. This will require vast amounts of energy storage.

Power noise. Power noise is unwanted power perturbation. A power source that produces power that surges and dips generates power noise. A load that takes power intermittently generates power noise. At low frequencies (hours, days), this is the realm of power management. Power perturbations of seconds or minutes are power noise. Power noise can be cancelled using energy storage, absorbing momentary excess power and returning it when needed.

Virtual energy storage. Electrical energy storage, such as batteries and capacitors, is very expensive. Fortunately, there are large reserves of "virtual energy storage" that can be utilized by improving power management. Usually, energy storage is thought to be like a battery or capacitor – you put electrical energy in and later get it back out in kind. However, it is just as effective to modulate a power source or a load to absorb the variations. As an example, a home owner might make ice to use later for air conditioning. It matters little, minute by minute, how fast the ice is made, so he can vary his consumption sympathetically to cancel the power variations.

Controls that implement virtual storage are a great opportunity for PSMA members.

Other Discussions

Modeling the "Smart Grid." Engineers like to model things and we take a crack at it. A model of a power supply with remote sense may have a lot to teach us about the smart grid and demand/response through data links. It can be analyzed for loop response and stability, and failure modes can be identified.

Computer failure, data interruption and corruption, sabotage and hacking. The excellent presentation by Jonathan Pollet on cyber-security scared the daylight out of us. The smart grid model is used to identify nodes that may be subject to hacking and to discuss computer failure, data interruption, and hacking vulnerabilities.

Designing for the "Smart Grid." One objective of the workshop is to provide guidance to PSMA members as they adapt their products for the smart grid. We present some ideas.

Metrics. Before we can control the grid, we have to know how to quantify it and measure it. We need metrics for:

- Frequency
- Voltage
- Impedance
- Current
- Power
- Load-lines
- Cleanness and Greenness

Compensation. Compensation is very important and VARs are poorly understood. We attempt an explanation through SPICE models and vector diagrams.

Power Management. We explore the dynamics of energy storage, particularly "virtual energy storage," and the optimization of the load-line as a fundamental power management tool. We also discuss briefly how to introduce market influence to power management.

The economies of "free." We explore the effect of "free" power during times of surplus production. Anything that stores "free" power is worthwhile, even if very inefficient. We also explore "leaky bucket" storage. It is worthwhile putting "free" power into a leaky bucket and trying to use it before it is gone.

Islands and emergency response. When something goes wrong, we want to maintain as much service as possible wherever we can and restore full service as quickly as we can.

Modeling the "Smart Grid"

We are engineers, and engineers like to model things. We chose a power converter with remote sense for a simple model.

The smart grid has a central computer with remote sense over a data link. Part of the regulation is through a shunt regulator, also remotely controlled through a data link. There may be feed-forward, also through a data link. There may also be additional local power sources at the load, shown below, as a solar panel. The solar panel may have maximum peak power tracking, so it may be resistant to control.

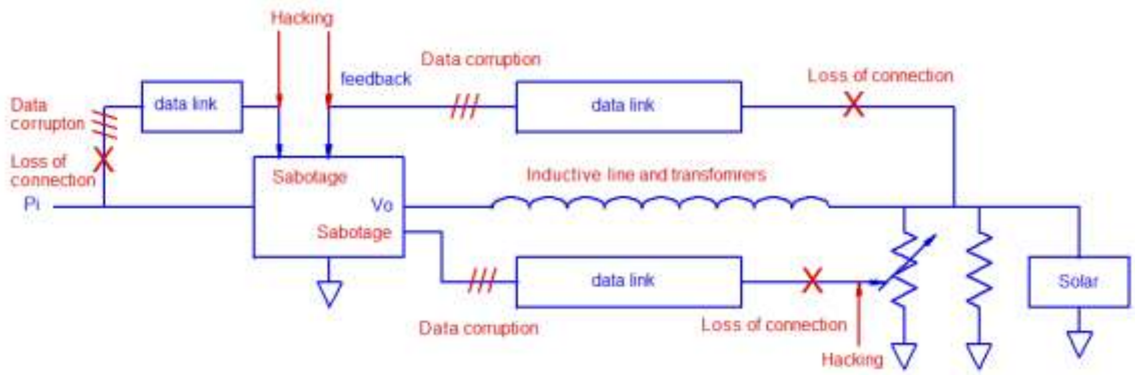

A power supply designer knows how to write all the equations, do a stability analysis, Nyquist plots, Bode diagrams, design the feedback compensation, and determine the frequency response and phase margins to make it stable.

Anyone who has designed a power converter with remote sense knows what happens if the remote sense line is broken or picks up noise. The transport lag through the data link adds uncertainty.

A lost sense lead is always a problem with remote sense. A lost control lead for a remote shunt regulation stage probably is comparable. Hacking is analogous to "noise" in a control system. We show nodes where hacking "noise" might be injected.

Computers and Data Links Fail

Computers and data links fail. For the Grid" to be reliable, it must function well without the computers and data links. Computer and data link failures are "contingencies."

Hacking and Sabotage

Hacking and sabotage can be lumped with other computer and data link failures. The model shows several nodes where hacked data can be introduced. We need to categorize hacking and sabotage so that priorities can be established.

Primary hacking and sabotage. Some forms of hacking and sabotage have such serious consequences that they may disrupt the operation of the grid with implications for homeland security. These are a first priority and must be identified. Changes must be made to eliminate the possibility of serious disruption, or worse, consequential damage, such as from high over-voltage.

Secondary hacking or sabotage. If hacking or sabotage has a small effect or affects a small part of the grid," it is unlikely to cause disruption. If hacking changed the calibration of some function by a percent or two, it is unlikely to cause disruption. If it affects only a small number of small customers, it is unlikely to cause disruption.

Defenses. The best defense against hacking and sabotage is to make them ineffective by limiting the damage that they can do. If the worst-case hacking or sabotage can affect the grid by only a few percent, or only at a very slow rate, it would no longer be any fun and the hackers would stop trying.

- **Parameter adjustment**. A central computer running a calibration algorithm may be able to make certain functions much more accurate, and the improved accuracy would carry over to the default mode. A percent or two of adjustment may make a significant difference to performance, for example, a voltage calibration, but it would never destabilize the grid.
- **Scheduling**. Power delivery profiles can be transmitted by a central computer and data link to distributed power sources and loads. Power changes should be slow enough to be detected and corrected before any disruption can occur. Secure data and verification by echo and third-party validation may make scheduling as secure as banking transactions.

  If Jonathan Pollet is correct, the ability to disconnect all customers simultaneously is a very serious cyber-security issue. There is no compelling reason to instantaneously disconnect customers who have not paid their bills. It would suffice to give a day's notice, followed by a 5% per hour decline in power allocation. To prevent the possibility of large-scale simultaneous disconnections, the one-day delay can be a random 20 to 28 hour delay. Both the delay and the ramp down of the power allocation can be hardwired in the smart meter algorithm so they cannot be changed by hacking. Most errant commands to disconnect would be caught before major disruption occurred and some customers might pay their bills during the warning period.
- Others functions may include marketing and arbitrage. The key is to transmit the data, verify it rigorously before it affects any output, and have the outputs change state slowly enough so that errant behavior is detectable with ample time for corrective action.

Designing for the Smart Grid

The smart grid offers great opportunities for PSMA members, but as an evolving technology, there are significant risks. Some will become experts in handling those risks and, to the extent that they do so successfully, they may have a market advantage.

For those who want to minimize the risk, we offer some advice.

Expect changes and design your power converters for easy upgrade, preferably field upgrade. The central power circuits with their heavy wiring and heat sinks are least likely to be changed. They can be optimized and chassis mounted.

The control circuits are more likely to change, so they should be on plug-in modules. If they contain firmware, consideration of hacking must be a primary priority.

The highest risk is in a data link, especially if it allows remote control or modification of the control algorithms. Power converter designers are accustomed to providing data and control interfaces, but the smart grid raises concerns to a whole new level. Can it be hacked? What happens if it is? Could it even be a homeland security issue, as Jonathan Pollet suggested, if a wide deployment of products could simultaneously make large step-changes in demand?

What is the potential liability if the power converter is hacked and causes substantial down-time or damage? Could easily hacked equipment be subject to recall?

Preferably, the hardware that interfaces with the smart grid is external to your equipment, interfacing through standard power converter bus interfaces, with all command implementation and security being the responsibility of someone else. If it is internal, it should be on a replaceable module and the equipment should work well in a satisfactory default mode with the module removed.

It may be possible to include robust security. Experience with smart meters shows that very often the security measures are not used. Including security measures may not be a defense if customers are known to ignore them or it they are non-standard or overly complex, which would discourage use.

Mainly, we should ensure that there is a robust default mode.

Metrics

Before the grid can be controlled and monetized, the parameters that affect the grid must be defined:

- Frequency
- Voltage
- Impedance
- Current
- Power
- Load-lines
- Cleanness and Greenness

Frequency

Understanding the role of frequency, transients, and the frequency domain of various functions and equipment is essential.

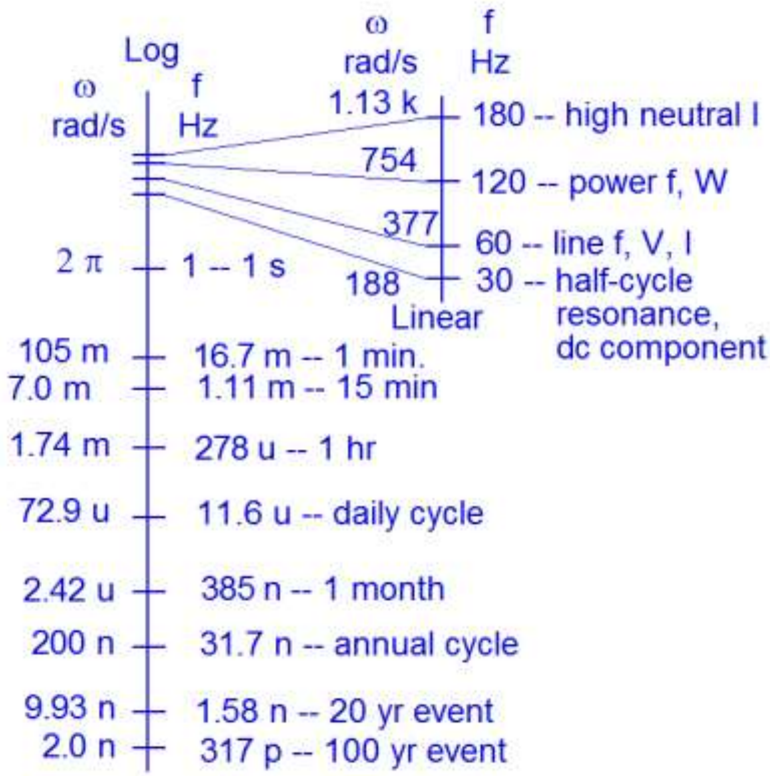

Compensation and power quality operate at or close to line frequency, 5 Hz and higher.

Regulation services are from about 4 MHz to 5 Hz.

Energy storage that compensates for the daily variation in solar panel output operates at about 11.6 µHz.

Seasonal variations in demand are about 30 nHz though in mid-latitude locations. Seasonal electrical demand may fluctuate at 60 nHz, with higher demand in winter for heating and in summer for air conditioning, and with reduced demand spring and fall.

Radians and ω

The frequency is given both in Hz, cycles per second, and ω, radians per second. Radians have some interesting properties. The reciprocal, $1/\omega$, is the time constant, a useful number. For small angles θ in radians, $\sin(\theta) = \theta$, which greatly simplifies some approximations.

Frequency Response of Equipment

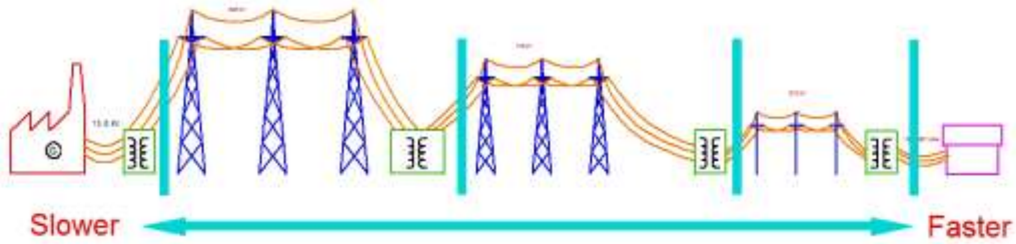

As a generalization, the equipment of the utilities is slower to respond to transients, while the consumer equipment is dominated by solid-state equipment (power converters), and can respond very quickly. It is logical, therefore, to solve higher-frequency compensation and regulation problems in or near the consumer equipment, a business opportunity for PSMA members.

Characterizing the Frequency Response of Equipment

Equipment connected to the grid can be characterized by frequency response, both physical limits and intended use. As an example, consider energy storage for power balance. The vertical axis is the log of power and the horizontal axis is the log of frequency.

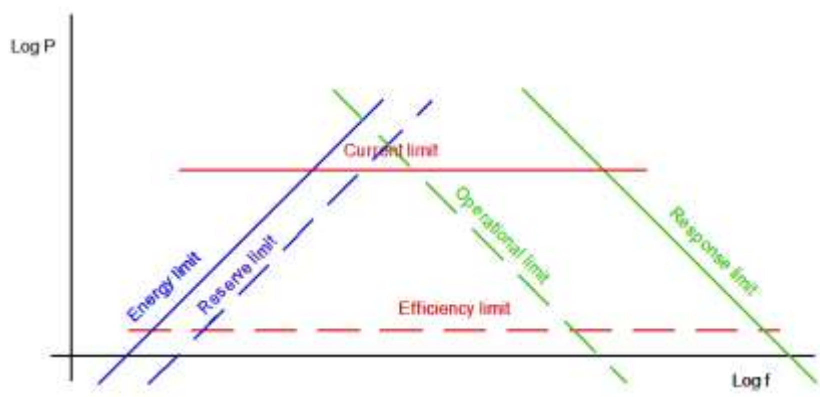

The physical limits of the equipment are the solid lines in the graph above. The energy limit is determined by the amount of energy the equipment can store, and that determines the lower limit of its frequency response. It has a plus-1 slope. There may be a higher reserve limit if it is undesirable to deplete the energy entirely.

The response limit is determined by the frequency response of the equipment. This may be very high for solid-state equipment, such as power converters or inverters, but there may be a much lower operational limit. As an example, battery storage may be capable of fast charge and discharge, in the order of minutes, but it may be preferred to limit its cycle rate to once a day to prolong the life of the battery. An emergency response to infrequent events may use the higher frequency response limit. As an example, if a transmission line is lost or a branch circuit disconnects, a very fast step response is appropriate.

The current limit caps the amount of power that can be handled. The efficiency limit is the lowest practical power to operate the equipment. This relates to the no-load power and it may be wasteful not to just shut it off.

For energy storage, the charge and discharge response is likely to be different. Energy storage to compliment solar panels may have to charge in a few hours around noon but may provide power from mid-afternoon to mid-morning.

It is best to sandbox, or restrict, the frequency response of equipment to the frequency range needed for its intended use. Equipment with frequency responses that do not overlap are unlikely to interfere with each other through unintended feedback loops.

We know, generally, that the large equipment for bulk power and bulk transmission tends to have a slow response. Going forward, we need to characterize the frequency response of all the equipment in the grid.

Regulation Services

An objective of control and compensation for the grid is to maintain a reasonably steady voltage and power variations (power noise) make this difficult.

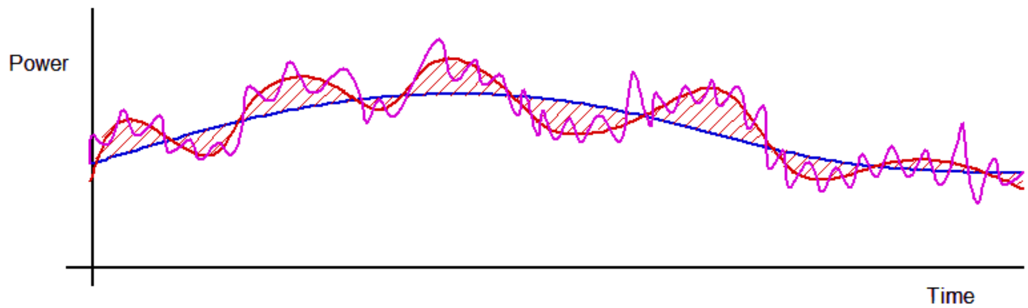

The graph above shows the power flow in a representative power regulator, as an illustration of frequency response. The regulator does not have the energy capacity to respond to the low-frequency variations, nor does it have sufficient frequency response to follow the highest frequency ripple. A band-pass filter sand-boxes its response to the mid-range shown by the crosshatched areas.

Frequency and Phase Angle

The grid is a synchronous system, so the frequency is the same throughout. Presently, the frequency varies slightly, but as more power sources are solid state, we expect that the frequency will be synchronized to a master clock.

Analyzing many of the relationships controlling power flow and compensation requires knowledge of the phase angle between various voltage and current phasors. Within the domain of the utilities (bulk generation, bulk transmission, sub-transmission), the Wide Area Measurement System (WAMS) provides this information. It is likely that the smart grid data communications can provide this information to distribution and end users when the computers and data-links are up and running. WAMS measures phasor data with reference to a precise clock provided by the global positioning system (GPS).

It would be very convenient if the phase angle of any voltage or current could be measured against a master clock. That may be the case eventually, but it is not presently possible and probably cannot ever be relied upon. Even with the smart grid, there needs to be a default mode to control the voltages and currents when the computers are down or the data-links are not working.

A phase angle reference useful for most equipment can be derived from the measured frequency, with suitable filters matched to the frequency domain of interest for the equipment.

The instantaneous frequency, ω, has two components: one due to underlying steady-state frequency, $\omega 0$, and the other due to fluctuations in the phase angle, α:

$$\omega = \omega 0 + \frac{d\alpha}{dt}$$

The frequency ω can be sampled twice each period of a sine wave by noting zero crossing. With three-phase power, ω can be sampled six times each period. Most power management functions operate at frequency domains that are well below line frequency, so this sampling method is more than adequate and is easy to implement.

If the frequency magnitude is then passed through a high-pass filter at the time constant of interest, the low-frequency components are eliminated and a reference frequency, ωr, can be synthesized. Its zero crossings can then be used as a reference for phase angle measurements and comparisons.

The Energy of Transients

Relationship of Frequency to Energy

Power management is largely about smoothing power fluctuations between power generation and power consumption. If there is no energy storage, the power must balance exactly. If power can be shunted to energy storage and returned, then the power generation and consumption can differ by the amount of power that can be diverted.

The product of power and time is energy. The power difference between supply and demand and its duration determine the amount of stored energy needed.

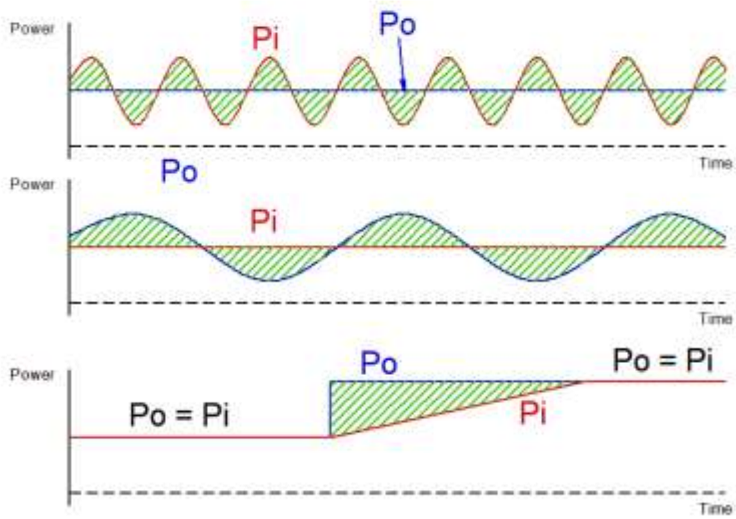

In the graphs above, power smoothing is implemented for three examples. In the first, the input power varies at a faster rate and the output power is smoothed, as in single-phase power factor correction. The energy that must be stored and returned each cycle is the cross-hatched area between the curves.

In the second, a load takes power at a varying rate, yet it is desired to keep the input power constant. Again, the energy that must be stored and returned each cycle is the area between the curves, but much more energy storage is required for lower-frequency power management.

In the third, there is a step change in the output power. Hypothetically, the power source is known to have a slow response and a step change in power results in a voltage droop. To avoid this, the change in the input power is limited to a rate that the power source can handle. The power difference must be supplied as energy from storage and the amount of energy needed is the cross-hatched area between the curves.

Fast-Recovery Energy Storage

Controlling energy storage is a powerful tool for correcting source-load imbalance and attenuating power noise. Fast-recovery energy storage is counter-productive and a serious problem to the recovery of the grid after power interruption.

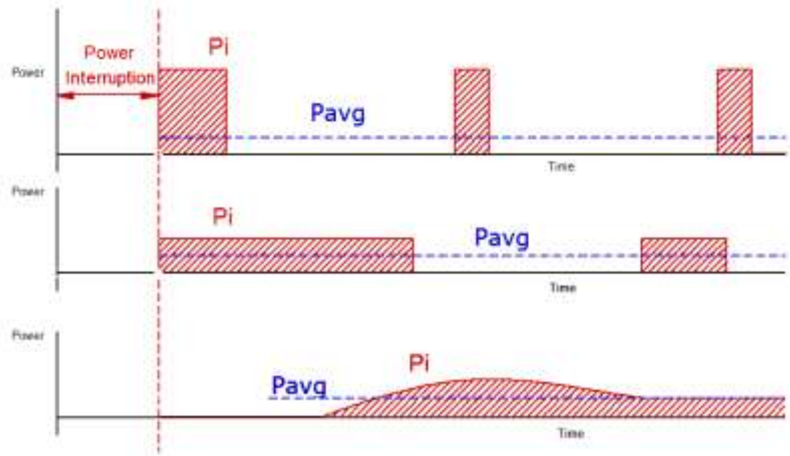

Consider three scenarios following a power interruption. As an example, a house has electric heat and it cooled down while the power was off. In this case, "energy storage" is virtual and heat is stored in the thermal capacity of the house. The first curve above is a fast-recovery space heater. It draws a large surge of power immediately following the restoration of power, then cycles normally. The on-going power noise is high.

The second curve is a space heater with a much smaller heating element. It also has a surge when the power is restored, but the level is much smaller, though it lasts longer. The subsequent power noise also is lower.

With controlled energy storage, there is a programmed delay until the heater starts, then it ramps up slowly, ramping down as the set-point temperature is approached, then settling down to a steady-state power draw. Some may complain about a delay in restoring heat, but a few-minute delay should not be objectionable and it would be very helpful to the grid when power is restored.

Voltage Measurement

Voltage measurement often ignores phase difference, a serious omission.

In an ac or dc circuit, a voltage difference causes current to flow through a resistance in accordance with Ohm's law. In a dc distribution system with any significant resistance, it is not possible to have equal voltages throughout or no current would flow.

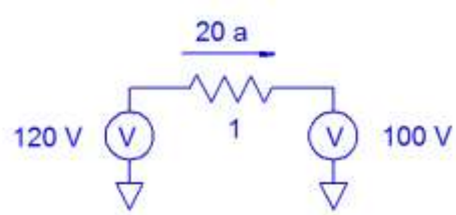

In an ac circuit, current flows between equal ac voltages if there is a phase difference. The ease and efficiency of changing voltage through transformers is one advantage of ac distribution, but being able to have nearly constant voltage throughout a distribution system is a significant advantage as well.

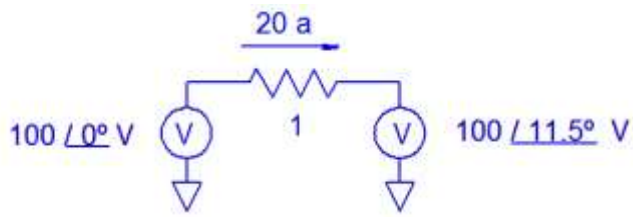

The phase difference actually causes an instantaneous voltage difference in accordance with the following expression:

$$Vr = 2 * V * Sin\frac{\theta}{2}$$

where Vr is the resultant voltage, V is the voltage magnitude, and θ is the phase angle difference.

For small angles θ, if θ is in radians, sin(θ) ≈ θ. The expression simplifies to:

$$Vr \approx V * \theta$$

If graphed as the percentage change in voltage versus the phase shift in degrees, the curve is very nearly a straight line for small phase angle differences. The graph below shows the percent difference voltage up to 30º phase angle difference.

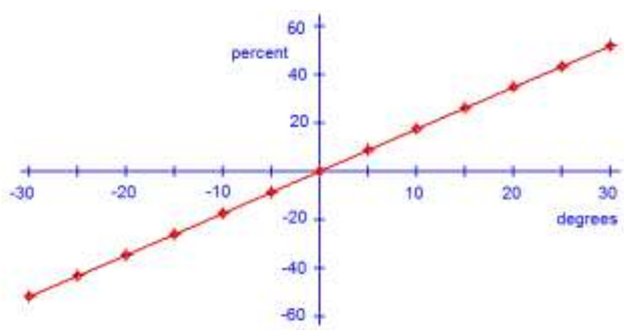

When the voltages are different in magnitude as well having different phase, the calculation is more difficult. A solution is to draw the phasors very accurately in CAD and measure the resultant.

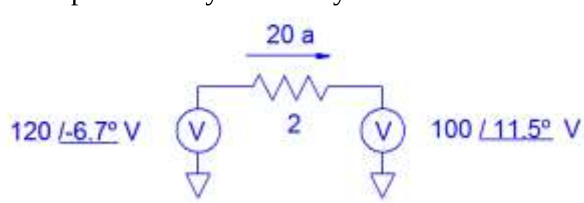

As an approximation:

$$Vr \approx \sqrt{\Delta V^2 + (V_a * \theta)^2}$$

In the example of the schematic above, the resultant Vr is 40, so 20 A flows in the 2 Ω resistor. The approximation gives Vr = 40.26, less than one percent error.

Extracting the Magnitude of Sine Waves

Ac voltage is a sine wave, ideally, having an instantaneous value of:

$v=V*\sin(\omega*t+\theta)$

where V is the voltage magnitude, ω is the frequency in radians per second, t is time, and θ is the phase angle in radians.

It is often useful to know the voltage magnitude, V. Because voltage is a sine wave and is continuously varying in instantaneous value, the voltage magnitude V cannot just be measured at any instant. V can be measured twice during the sine wave: at its peaks. However, there is complete information about the voltage magnitude V at any point in the cycle, including zero-crossing. At zero-crossing, the slope of a pure sine waveform fully defines its magnitude. Throughout the cycle, the dynamic voltage magnitude, V, can be derived using the following identity:

$$\sqrt{\sin^2(\alpha)+\cos^2(\alpha)}=1$$

If the voltage sine wave and the corresponding cosine wave are used:

$$V=\sqrt{(V*\sin(\theta))^2+(V*\cos(\theta))^2}$$

The cosine term can be found either by differentiating the voltage or by integrating it, and both are useful. The three expressions below can be solved in real time and are equivalent for a pure sine wave:

$$V=\sqrt{v^2+\left(\frac{dv}{dt}\right)^2}$$

$$V=\sqrt{v^2+\left(\int v\,dt\right)^2}$$

$$V=\sqrt{v^2+\left(\frac{dv}{dt}\right)*\left(\int v\,dt\right)}$$

Real ac voltages are not pure sine waves; the deviations from a pure sine wave result in a time-varying value for the voltage magnitude V includes all of the harmonics and sub-harmonics of the voltage waveform. The three expressions for the voltage magnitude, V, given above have different harmonic content due to the different ways that the cosine term is derived. The first expression exaggerates noise, and the second minimizes it. Once derived, the voltage magnitude may be filtered like any analog signal if a particular roll-off or bandwidth is needed. Rigorous analysis and optimization can provide term projects for many graduate students.

For some calculations, the rms voltage Vrms is preferred, and is derived as:

$$Vrms=\frac{V}{\sqrt{2}}$$

Step Change in Voltage

In a well-controlled system, voltage changes are not abrupt. A step-change in the voltage very likely indicates an equipment failure, such as a lost transmission line or disconnected load, requiring an emergency response.

Impedance Z and Admittance Y

Understanding impedance in its several forms is key to managing compensation and power. Percent impedance, %Z, is used by the utilities to characterize the impedance of transformers. It may be important for compensation. The conservation voltage reduction factor CVRf is actually a power admittance, pY. Because current cannot be used as a parameter for voltage control; its reciprocal, the power impedance pZ, may be important for power management. A modification of power impedance defines market impedance, possibly important for arbitrage and scheduling.

Ohm's Law

Ohm's law defines the relationship between impedance, voltage, and current:

$$V = I * Z$$

where V is the voltage in volts (V), I is the current in amperes (A), and Z is the impedance in ohms (Ω).

The reciprocal of impedance Z is admittance Y:

$$Y = \frac{1}{Z}$$

where Y is the admittance in siemans (S) (formerly mhos ℧).

Impedance in ac circuits is complex. The impedance may be resistance, R, or reactance, X. For a capacitor, C, or an inductor, L; the reactance is a function of frequency.

$$Z = R + jX$$

$$X = j * \omega * L = j * 2 * \pi * f * L$$

$$X = \frac{1}{j * \omega * C} = \frac{1}{j * 2 * \pi * f * C}$$

$$B = j * \omega * C = j * 2 * \pi * C$$

The reciprocal of resistance is conductance, G, and the reciprocal of reactance is susceptance, B.

Negative Impedance

There is frequent mention of "negative impedance;" particularly with reference to "constant power" loads. "Negative impedance" is not truly negative. If you measure the voltage and the current into a power converter and divide the voltage by the current, the result is a positive number. "Negative impedance" almost always is shorthand for negative dynamic impedance:

$$dZ = \frac{dv}{di}$$

To maintain constant power; if the voltage goes up, the current goes down and vice versa, so the dynamic impedance is negative.

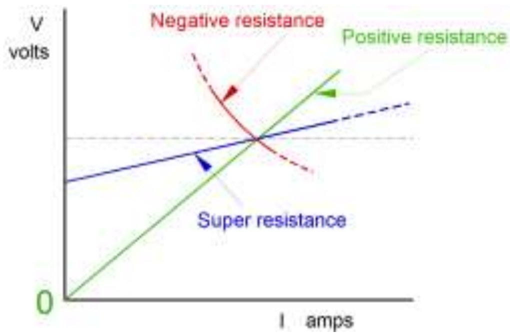

The reciprocal is the dynamic admittance, dY:

$$dY = \frac{di}{dv}$$

Percent Impedance, %Z, and Decimal Impedance, •Z

Percent impedance is the usual way utilities characterize the impedance of transformers. It is special case of “per unit calculation,” beyond the scope of this report. The readers are encouraged to familiarize themselves with per unit calculation, as it is a powerful tool.

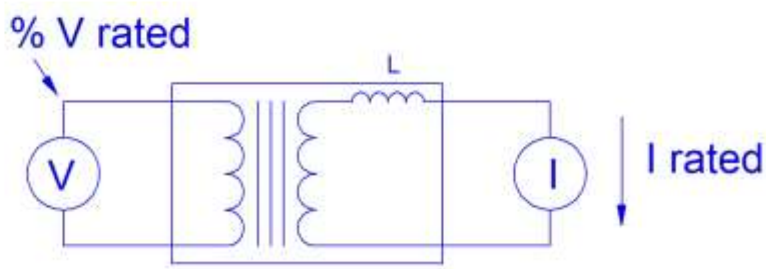

To determine the %Z of a transformer, a short circuit is applied to the secondary through an ammeter. A voltage is applied to the primary, and the primary voltage is increased until the rated current flows in the secondary. The voltage drop, V_{drop}, is measured on the primary and expressed as a percentage of the rated primary voltage, V_{rated}, or %Z. As an example, a transformer having an impedance of 5% would have a voltage drop of 5% of its rated primary voltage with 100% of its rated current flowing in the secondary.

$$\%Z = \frac{V_{drop}}{V_{rated}} * 100\,\% \text{ at } I_{rated}$$

Generally, useful for other operating points:

$$\%Z = \frac{I_{rated}}{V_{rated}} * \frac{V_{rated} - V}{I} * 100\,\%$$

The voltage drop, $V_{rated} - V$, is linearly related to the current, I.

Percent impedance, %Z, is a complex entity. Percent impedance, %Z, may be resolved into the percent resistance, %R, and the percent reactance, %X.

Percent impedance, %Z, is a normalized function. Unlike ohm impedance, Z; the percent impedance is the same into the primary winding and into the secondary winding. Therefore, in a schematic, the transformer can be represented as a series impedance.

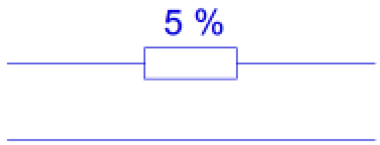

Decimal impedance, •Z, is the same except expressed as a decimal quantity, not percent, avoiding the factor of 100%. It is more useful for calculations:

$$\bullet Z = \frac{I_{rated}}{V_{rated}} * \frac{V_{rated} - V}{I}$$

The percent admittance, %Y, and decimal admittance, •Y, are the reciprocals. Be careful, %Y ≠1 / %Z. The factor of 100% causes confusion.

Given the decimal impedance, •Z, and the voltage magnitude, V; the current, I, can be found:

$$I = \frac{V_{rated} - V}{V_{rated} * \bullet Z} * I_{rated}$$

If the voltage magnitude, V, is 130 Vrms; the rated voltage V_{rated} is 120 Vrms; the rated current is 1,000 Arms; and the •Z is 0.05; then:

$$I = \frac{120 - 130}{120 * 0.05} * 1{,}000 = -1{,}666.7 \text{ A}$$

Percent Impedance Reveals Percent Loading

The percentage impedance, %Z, indicates the voltage drop as a percentage of the rated voltage when the rated current is flowing. When a lower current is flowing, the voltage drop is proportionately lower. Accordingly, in a transformer having a %Z of 5%, if the voltage drop is 2.5% of the rated voltage, the current is 2.5/5 or 50% of the rated current:

$$\%I_{rated} = \frac{\%V}{\%Z} * 100\,\%$$

where %V is the measured voltage drop as a percentage of the rated voltage.

If the voltage magnitude, V, is 130 Vrms; the rated voltage V_{rated} is 120 Vrms; and the %Z is 5%; then:

$$\%I_{rated} = \frac{\left(\frac{120 - 130}{120}\right) * 100}{5} * 100 = -166.7\,\%$$

We prefer the load ratio, using •Z, to avoid the factors of 100%.

$$\frac{I}{I_{rated}} = \frac{120 - 130}{120 * 0.05} = -1.667$$

If applied to the example of a transformer with a %Z of 5% with a secondary voltage of 130 V and the rated voltage of 120 V on the primary; the current is flowing back into the secondary and there is a large overload, 1.667 times the rated current(166.7%).

Percent Impedance %Z is a Normalizing Function

Because the base of %Z is the rated current, it is a normalizing function. With equal voltage difference, parallel currents divide in proportion to the respective ratings.

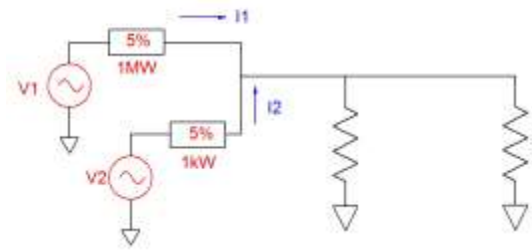

The schematic above shows two voltage sources, V1 and V2, with two currents, I1 and I2, into a load. V1 and V2 have source impedances, both 5%Z, but one voltage source is rated for 1 MW and the other is rated 1 kW, three orders of magnitude apart. If both voltage sources are incremented by the same percentage of the rated voltage, the currents divide so that each has the same proportion of its rated current. As an example, if both voltages are raised 2.5%, then each voltage source supplies 50% of its rated current. The power divides proportionately to the respective ratings, 500 kW for V1 and 500 W for V2.

Percent impedance, %Z, does not vary as more sources with the same %Z are added in parallel. Adding more sources increases the rated current and the same percent voltage drop occurs at the same percentage of the increased rated current.

Reactive Load Line

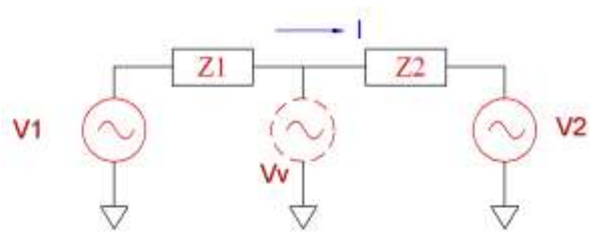

In the previous example, we assumed that the impedances Z1 and Z2 were resistive. It is interesting to compare the results if the impedance is 5% reactive (inductance). If a volt meter is used to measure the voltages across the impedances, the voltage drop is the same as before. However, if the voltage at V1, Vv, and V2 are measured, all are approximately the same because the voltage drop in inductors leads the current by 90º (π / 2 radians). The voltage is a quadrature voltage and the voltage difference is almost entirely due to phase shift. (As the phase angle increases, this approximation becomes less accurate.)

The phase shift at full load through a 5% reactance is:

$$\theta = \tan^{-1}(0.05) = 0.04996\,\mathrm{rad} = 2.86^{\circ}$$

As with the sine, if the phase angle θ is in radians, tan(θ) ≈ θ for small angles. At lower loads, the phase angle is proportionately lower. If the voltage source V1 and the impedance Z1 are loaded to 30% of its rating, the phase angle across Z1 is 0.3 * 0.05 rad = 0.015 rad, or 0.86º.

We want to maintain the voltage within a narrow limit at the output, but phase variation is relatively unimportant. Therefore, an impedance that is mostly reactive seems indicated.

An impedance with %R = 2% and %X = 10% may be a good compromise. The resultant is 10.8%Z. From full rated current out of the voltage source to full rated current into it, the voltage is + 1% to –1% and the phase angle varies from – 0.05 rad to + 0.05 rad. This is a reasonable phase shift for compensation as well.

Conservation Voltage Reduction Factor

There is a concept used by the utilities called “conservation voltage reduction.” The conservation voltage reduction factor is defined as:

$$CVRf = \frac{\text{Percent demand reduction}}{\text{Percent voltage reduction}} = \frac{\Delta P}{\Delta V} = pdY$$

The conservation voltage reduction has a strong similarity to dynamic admittance, dY, with the substitution of ΔP for ΔI, so we can rename it dynamic power admittance, pdY This not as farfetched as it might seem. Nominally, voltage does not change much, especially at the distribution end of the grid,” so changes in power are mostly due to changes in current. It is also a useful figure of merit for voltage regulation. It is also normalizing, as it does not change going through transformers or modulators, and it is not negative for “constant power” loads and sources.

Its reciprocal, power impedance, pdZ, may useful as well:

$$pdZ = \frac{\Delta V}{\Delta P}$$

Percent Power Impedance and Per-Unit Power Impedance

For controlling power, a key parameter is the rated power. Knowing the percent of rated power being supplied or consumed is key to knowing the reserve capacity to supply or absorb more power. The same transformation of power for current is used for the conservation voltage reduction factor, so we can derive a percent power impedance:

$$\%pZ = \frac{P_{rated}}{V_{rated}} * \frac{V_{rated} - V}{P} * 100\,\%$$

The per-unit power impedance is:

$$\bullet\ pZ = \frac{P_{rated}}{V_{rated}} * \frac{V_{rated} - V}{P}$$

The percent power admittance and the decimal power admittance are:

$$\%pY = \frac{V_{rated}}{P_{rated}} * \frac{P}{V_{rated} - V} * 100\,\%$$

$$\bullet\ pY = \frac{V_{rated}}{P_{rated}} * \frac{P}{V_{rated} - V}$$

As applied to equipment, knowing the percent power impedance, %pZ; the voltage drop, $\Delta V = V_{rated} - V$; and the power, P, being consumed or provided: the rated power, P_{rated}, can be determined for either a power source or a load. It does not seem strange to specify the percentage of power that a power source is supplying, but it is not customarily applied to loads, as they tend to be on or off. However, for a controlled load, the ratio of the power actually being consumed to the total power that the load could consume if turned on 100% is a useful concept.

Market Percent Impedance

In percent power impedance, as defined above, the basis is the rated load in load, P_{rated}. If the amount of powe. P_{sale}. that is for sale is substituted for the rated power, P_{rated}; some of the control equations can be solved to determine a market response.

Transformers and Impedance

With a transformer, impedance on one winding reflects to the other as the square of the turns-ratio. On the other hand, percent impedance or decimal impedance is unaffected.

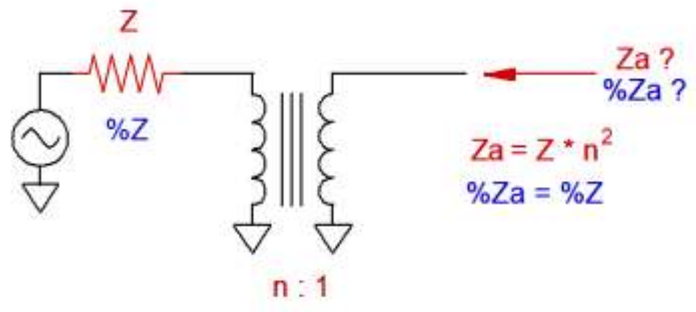

However, be careful of the respective ratings. For example, if a 100 A rated distribution line connects through a sub-station 10:1 step-down transformer to a 100 A rated sub-transmission line, the answer is very different. If the impedance of the transformer itself is neglected, the 100 A rating of the primary sub-transmission line looks like a 1,000 A rating from the secondary side.

For a given load, the ohm/per km of the primary looks like ohm/n2 per km, vanishingly small if n is large. For the same %Z to make sense, the assumption is that n2 similar loads are paralleled into the primary circuit. The primary and secondary are both loaded at the same percent of rating if all n2 loads are at the same percent of rating.

Modulators

Modulators look like transformers with a variable turns-ratio:

$$Vo = \frac{1}{n} * Vi$$

$$Vo = \frac{1}{D} * Vi$$

where:

Vo is the output voltage;

Vi is the input voltage;

n is the transformer turns ratio, and:

D is the duty-ratio of a boost converter.

The ohm impedance reflected from the output to the input changes radically as n or D are varied; %Z and pY do not.

Changing Impedance to Control Power

In many circuits, power control is through changing the impedance of the circuit. It may be as simple as an on-off switch with a resistance heater. Power-factor-corrected power converters simulate a resistor at line frequency and the power is controlled by varying the value of the simulated resistor.

Be careful… the impedance may be very different at other frequencies. For example, the input impedance of a power factor controlled power converter is “negative” (dynamic impedance) at lower frequencies. The frequency response is important in characterizing power converters.

Synthetic Impedance

Synthetic resistors are widely used in power converters, though not by that name. A common example is power factor correction (PFC), in which the input is made to look resistive. The equivalent circuits are shown below; the synthetic resistor on the left and the way it is implemented on the right.

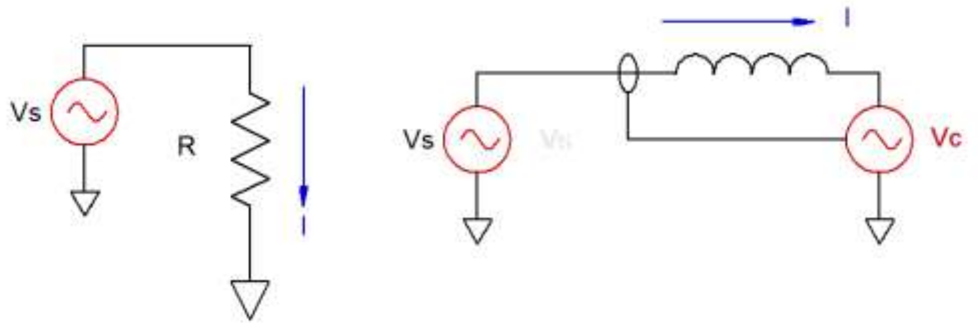

In a PFC power converter, the input current is made to "look resistive," that is, I = V/R. The value of the resistor, R, changes as the load on the power converter changes. In steady-state condition, the input current varies only as the input voltage Vs in accordance with Ohm's law. In the actual circuit, there is a very, very small inductor on the input with a pulse-width modulated boost converter to a storage capacitor. The pulse-width modulator is controlled by a current sense so that the voltage Vc seen on the load side of the inductor is almost identical to the supply voltage, Vs; so nearly so that the inductor current barely changes, but does so enough to replicate the current that would be seen in the resistor.

Current and Power

Using Ohm's law, given the voltage, V, and the impedance, Z, or admittance, Y; solve for the current, I:

$$I = \frac{V}{Z} = V * Y$$

The current between equal voltages with a phase difference θ is given by:

$$I = \frac{2 * V * Sin\frac{\theta}{2}}{Z} \approx \frac{V * \theta}{Z}$$

where θ is the phase angle in radians.

Where there is both a voltage magnitude difference and a phase difference, it is more complicated. Plotting the phasors to find the solution graphically is suggested. As an approximation for small-phase angles is given as:

$$I \approx \frac{\sqrt{\Delta V^2 + (V_a * \theta)^2}}{Z}$$

With two variables, volts and phase angle, we can determine the current. In many cases, for example in power converters, Z can be varied as well.

The same techniques described for voltage may be used to extract a value for the dynamic current magnitude, I, so the details are not repeated here.

Conservation of Charge

Conservation of charge requires that the sum of the currents into a system is zero, often called Kirchoff's law. Current passes through a system without any delay or phase change and, if there is a quadrature component, it passes through the system unchanged as well. The phase of the current relative to a fixed reference angle does not change. (If multiple currents are injected at different points, each is conserved and the vector sum of all currents into the system is zero.)

A current source drives a current through any impedance without change.

Power and VARs

Instantaneous power, p, is the product of the instantaneous voltage, v, and instantaneous current, i:

$$p = v * i$$

If the current and the voltage are in phase:

$$p = \frac{V * I}{2} * (1 - \cos(2 * (\omega * t + \theta))$$

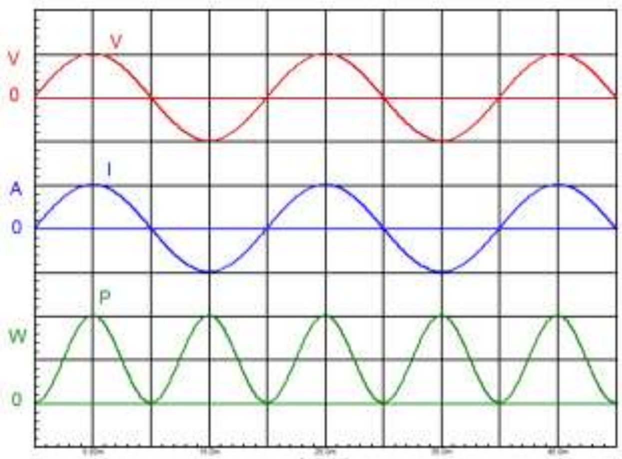

If the voltage magnitude, V, and the current magnitude, I, are multiplied, the product is volt-amps, VA:

$$VA = V * I$$

If the voltage and current are in phase, the power magnitude, P, equals the volt-amps, VA. If there is a phase angle, α, between the voltage and the current; then the volt-amps, VA, has "real" and "imaginary" components corresponding to the power and the VARs:

$$P = V * I * \cos(\alpha)$$

$$VARs = V * I * \sin(\alpha)$$

The voltage magnitude, V; the current magnitude, I; and the phase angle, α, all may vary with time and have harmonic and sub-harmonic content. A real-time analysis of the respective waveforms is probably the only way to extract a dynamic power magnitude, P, for single-phase ac.

Conservation of Energy

The conservation of energy requires that the power into a system, P_i, be equal to the power out, P_o, plus losses, P_L:

$$P_i = P_o + P_L$$

The same is true of reactive power, VARs, except that there are no losses in the reactance:

$$VAR_i = VAR_o$$

VARs are sometimes considered to be generated in capacitors and dissipated in inductors. This is a very narrow interpretation and is easily misunderstood. If a capacitor is used in parallel with a load resistor to generate a leading current, the phase angle of the current leads the phase angle of the voltage by an amount, α. If the voltage source and the voltage at the load are comparable, there is a phase angle gradient through the impedance that separates the load from the source and, at some point, the voltage and current have the same phase angle. It is a mistake to assume that the VARs are used up at that point and more VARs should be injected. The point at which the phase angles are equal could be called a virtual voltage and that phase angle could be called zero as a reference. It is a summing junction, not a load point. With no power or VARs added or subtracted, it makes no contribution to power and VARs balance.

Continuing through the impedance, the phase-angle gradient continues and the phase angle becomes increasingly negative until at the source, the current leads the voltage by the same amount that it lagged at the load. The VARs are equal, but have a different sign because the current is flowing into the system on one end and out of the system at the other. The law of the conservation of energy requires that the sum of the energy flowing into the system is zero, so is the sum of the power and VARs.

It may be useful to add VARs at an intermediate point in a distribution system, but not because the VARs are "used up." It is because more uncompensated loads have been added along the way. However, current injected at an intermediate point in a system can flow both ways, so can power and VARs. If VARs flow toward the load, they are counter-productive.

Power in Series Circuits

If power flows through a series circuit, the loss in each section is the square of the current times the resistance of that section:

$$P_L = I^2 * R_i + I^2 * R_1 + I^2 * R_2 + I^2 * R_n + I^2 * R_o +$$

We know that the losses in the grid are approximately seven percent, but we do not know the proportion that is series loss nor under what conditions. That needs to be defined going forward. We know that some of the loss is core loss in the transformers and that loss is relatively fixed. If the resistances do not change with load, then the percent loss in any section is probably fairly constant, and we could postulate 3.5% total series losses. That suggests that the %R of the distribution network is about 3.5%R on average.

Step Change in Current or Power

In a well-controlled system, current and power changes are not abrupt. A step-change in either very likely indicates an equipment failure, such as a lost transmission line or disconnected load. The power must be adjusted to compensate.

Load-Lines

A good understanding of load-lines is fundamental to voltage and power control. A textbook example of a load line is a voltage source with series resistance and a resistive load, shown below.

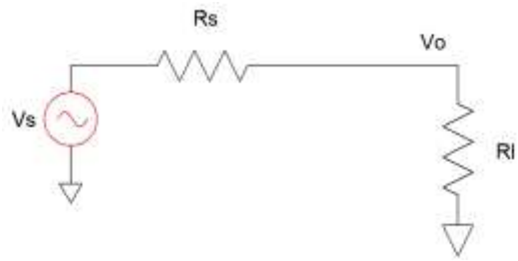

If the voltage increases, so does the current and, therefore, the power. A voltage increase inherently increases the loading on the power source, a stabilizing effect. Analytically, resistive loads are well known to damp oscillations.

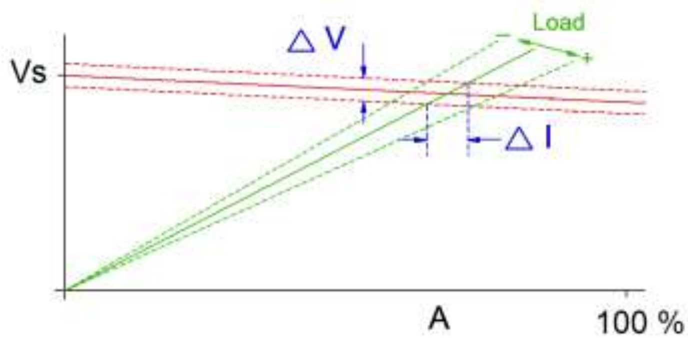

All controlled loads have "negative resistance," jargon for "negative dynamic resistance." The power is regulated to be "constant." This does not mean that the power does not change – it does change in response to load changes, as shown below – but it does not change in response to voltage variation. To keep the power, p, constant; if the voltage, V, increases; the current, I, decreases so that their product, $P = V * I$, remains constant.

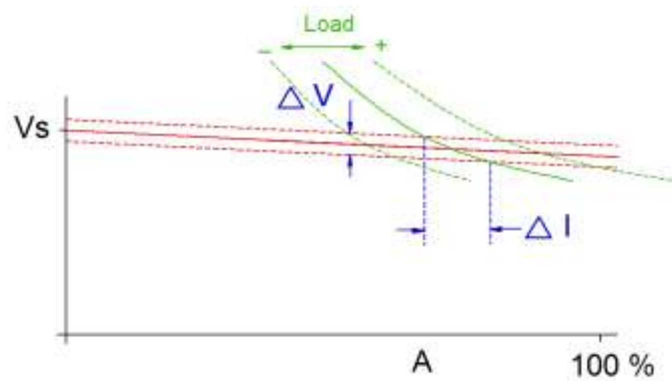

One consequence of controlled loads is that current cannot be used as a parameter for power management, as a change in current is impossible to interpret, and another is that the damping effect on voltage variation is lost.

Another consequence is that conservation voltage reduction is ineffective for power management with controlled loads. In fact, with constant power loads, secondary effects actually make attempts at saving power through conservation voltage reduction have the opposite effect. With ideal constant power control, reducing the voltage has no effect. With real power converters and a lossy distribution system, it is counter-productive. As the voltage goes down, the current goes up, and losses increase as the square of the increased current, both in the power converters and the power delivery system.

Most "solid state" loads "look resistive" at line frequency due to power factor correction (PFC), so the negative resistance is not a serious problem for compensation and power quality correction. Every controlled load has negative resistance at low frequency, most loads are controlled, and a trend of modernization is to control more of the loads. Variable speed drives are replacing induction motors. LEDs are replacing incandescent bulbs.

Controlled loads also make momentary interruption of power an ineffective strategy for power management in most cases. Once the power is restored, many of the controlled loads go into overdrive to make up for the lost power. Since efficiency is probably less, it may actually increase demand through increased losses.

Induction Motors

Probably the most effective equipment for maintaining regulation is induction motors, as they tend to regulate both the voltage and the frequency. Synchronous generators and motors are easier to understand, so we will discuss them first, because they are important as well, but also as a pedagogical first step.

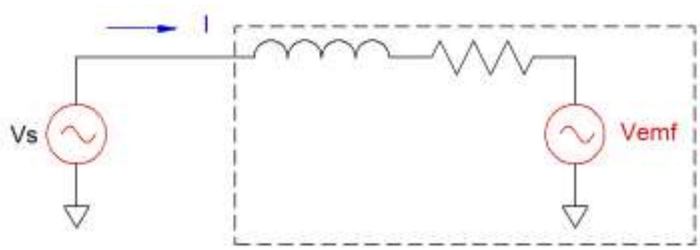

In synchronous machines (motors and generators) and induction machines (mostly motors), an important part of the load line is the opposing electromotive force (counter-emf), modeled as a voltage source. It is very closely linked to the rotor speed and excitation.

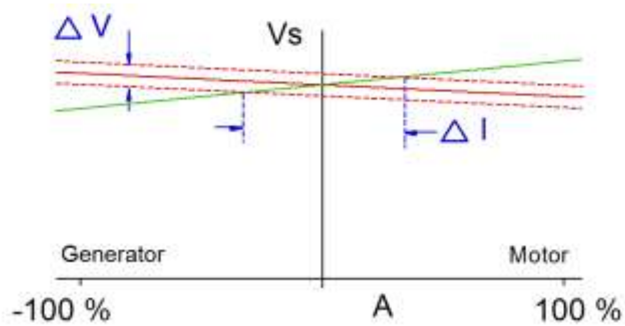

At no load (ignoring losses), the counter-emf equals the input voltage and there is no current. The current increases as the load (if the voltage and frequency are steady), so that mechanical power out equals electrical power in. The details of electrical machines are beyond the scope of this report. The important part is that the counter-emf is linked to the shaft speed and excitation, neither of which can change quickly. At equilibrium, the counter-emf equals the line voltage, less the voltage drop due to losses.

If the line voltage, Vs, changes; there is significant lag until the counter-emf catches up and a new equilibrium is established. During that time, the current is limited only by the series winding impedance, which is quite low. A step increase in voltage results in a disproportionately large step increases in the current, initially. At steady-state, the increase in voltage results in a decrease in current if the load has not changed.

The large initial increase in current means that an increase in voltage causes an even larger increase in power. The effect is to resist the change in voltage, a damping effect much like an R-C snubber.

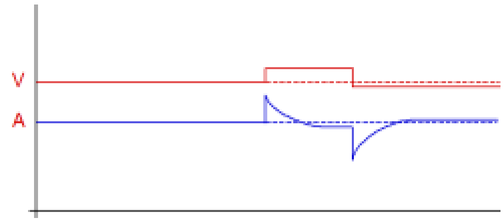

Notice that when the voltage decreases, the steady-state current is higher, so the motor too has "constant power" characteristics. Conservation voltage reduction is ineffective.

Three-Phase Synchronous Machines

The equivalent circuit model above is for a single-phase synchronous motor or generator. The three-phase model is shown below. There are three equivalent voltage sources corresponding to the emf of each of the three phase windings; but, because they are from a common rotor and share the excitation, the voltage magnitude of each is exactly the same and the phase difference of the emf voltage is exactly 120º. That is great for a generator, as the three-phase output voltage is balanced and precise. However, both generators and motors are connected to the three-phase power lines and the line voltages may not be balanced.

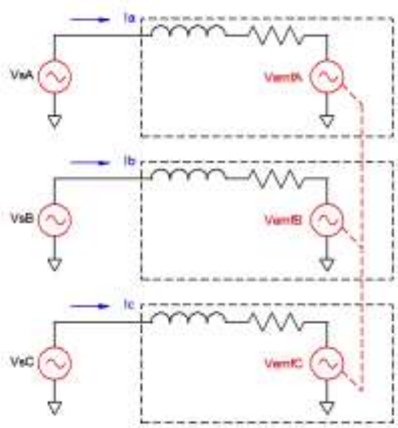

As can be seen by looking carefully at the model above, the counter-emf voltage cannot simultaneously be optimum for unbalanced voltages. If one phase is high and another is low, there is excess current into one and out of the other. In one sense, this is good, because it tends to load the phase with excess voltage and absorb the excess energy. However, it also causes excessive currents in the stator winding, with accompanying excessive I^2R losses. Even though the currents are large, the net change in torque is nil, so the large currents are not balanced by a change in rotor speed as in the single-phase motor. The excess current and excess losses persist as long as the unbalanced state continues.

This emphasizes the importance of good power quality at the distribution level so that it does not reflect unbalanced voltages and currents to the generators or motor loads. Unbalanced currents are known to cause circulating currents in the distribution system, with increased losses. The increased losses in the generators and motors is less recognized, but perhaps more wasteful – a loss that may damage the motors and increase the power bill.

For the purposes of this discussion, all of the above applies to induction motors as well. However, the shaft speed of an induction motor (or generator) is not locked to line frequency, except at no load, and then only theoretically if losses are ignored. The excitation of an induction motor is self-induced through a transformer action between the stator and rotor and the currents can vary in the squirrel-cage winding of the rotor. The "slip-speed" varies with torque.

Both an induction machine and a synchronous machine tend to stabilize the frequency due to their high rotating inertia. A change in frequency causes a momentary increased phase displacement between the emf voltage and the line voltage, with resulting changes in the current just as a comparable voltage change. This tends to resist a change in frequency, and it takes energy to accelerate the rotor to a new speed. Unfortunately, once the rotation has adjusted to a new speed, it tends to stabilize the new erroneous frequency, and it takes the same energy in reverse to correct it.

A little-recognized phenomenon is that induction motors can become generators if a branch circuit is disconnected. Without the stabilizing effect of the grid on frequency, the frequency drops rapidly as the smaller induction motors decelerate. The decelerating inertia of the larger motors then provides significant energy and can cause a very high voltage surge for a short while. A portion of the grid that could become an island and that has a large number of induction machines must have a way of absorbing the energy to limit the voltage.

It is strongly suggested that the control algorithms of "solid-state" inverters and power converters emulate synchronous machines. However, unlike real synchronous machines, the frequency change can be resisted if all of the loads are synchronized to a common clock. This cannot be achieved with the present grid, but if the algorithms of "solid-state" inverters and power converters are biased to return to a clocked frequency if perturbed, once a critical mass of new equipment is installed, the grid will synchronize.

Power Load-Lines

Just as the familiar load-lines are based upon impedance in ohms, they can be drawn using other impedances as well. If power, P, in W is substituted for current, I, in A; the load line of a resistor becomes as shown below.

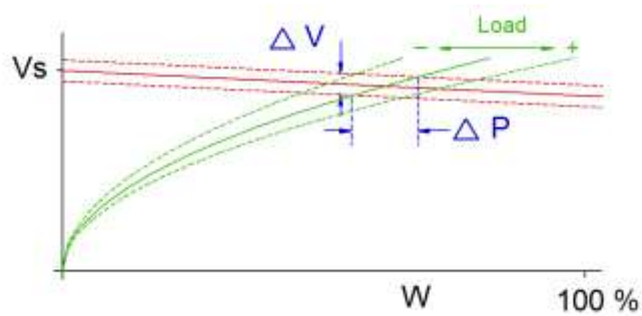

The usefulness of this transformation is apparent when the power load-line for a "constant power" load is plotted, as shown below.

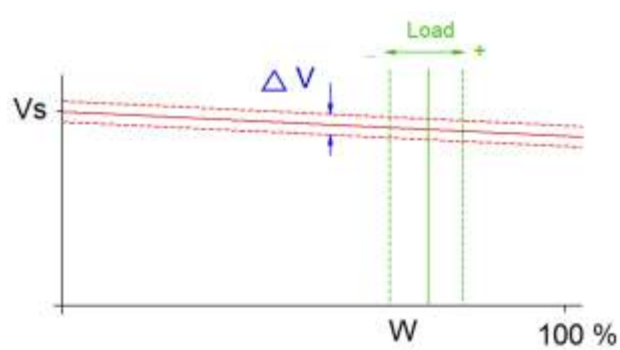

The negative slope is avoided. Assuming some mix of resistive and constant power loads, the net load is always positive resistive when plotted against power.

Three-Phase AC Voltage, Current, and Power

For balanced three-phase ac voltage, current, and power; there are some elegantly simple relationships because the sine and cosine components cancel when they are added:

$$\sum v = 0$$

$$\sum i = 0$$

$$V = \frac{3}{2} * \sqrt{v_a^2 + v_b^2 + v_c^2}$$

$$I = \frac{3}{2} * \sqrt{i_a^2 + i_b^2 + i_c^2}$$

$$P = v_a * i_a + v_b * i_b + v_c * i_c$$

$$P = \frac{1}{2} * V * I * \cos(\theta)$$

$$VARs = \frac{1}{2} * V * I * \sin(\theta)$$

where:

V is the voltage magnitude;

v is the ac voltage;

V*sin (ω*t+θ+α);

I is the current magnitude;

i is the ac current; and

P is the power magnitude.

With three-phase power, the cosine terms cancel. θ is the phase angle between the voltage and the current, α is the phase displacement between the phases and usually is 0º, 120º, and 240º, respectively, for the three phases.

In the above relationships, usually only the fundamental sine and cosine terms that cancel, so solving the equations in real time result in dynamic quantities with harmonics and sub-harmonics that are useful for power quality correction and power management.

For power quality, feedback can be used to make $\sum v = 0$ and $\sum i = 0$.

Power Quality and Phase Balance

Considering the extra losses that unbalanced voltages and current cause in three-phase motors and generators, power quality is important. Control of the third harmonic is important as well, as it results in high neutral currents.

For rotating machines, only the power in the fundamental harmonic of the sine wave contributes power to the shaft. All energy in the higher harmonics is dissipated as heat. This is not true of a resistor, nor is it true of most power converters. Therefore, it is beneficial to keep higher harmonics out of motors, but they can be absorbed in power converters.

Both phase imbalance and higher harmonics cause increased losses in motors and generators, and are best avoided. The best way to avoid higher harmonics is to require power factor correction in all loads and power sources. The best way to avoid imbalance is to use three-phase power exclusively. Three-phase rectification is simpler and more efficient as well.

We recommend that all larger power converters and inverters operate on three -phase power. If it is not available, the utilities should consider providing a rotary converter on loan.

Some equipment used by utilities to improve balance is probably not familiar to power converter engineers, and beyond the scope of this report. Those who want further knowledge should familiarize themselves with buried tertiary windings and transformers with zig-zag windings. Though not primarily intended for that purpose, synchronous converters absorb voltage imbalance through increased losses just as in a motor.

Power Noise

For power management, particularly regulation services, it is useful to compare the dynamic power magnitude to a reference power. Unintended and unwanted differences are power noise. Mitigating power noise is a business opportunity for power converters.

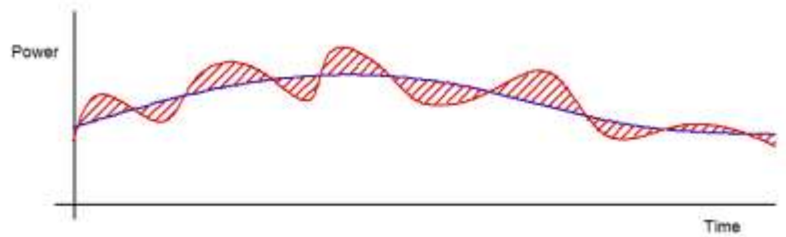

In the graph above, power noise is the higher frequency power fluctuation above and below the baseline power curve.

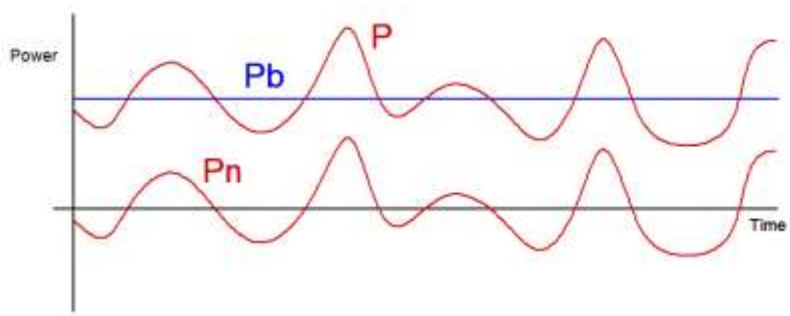

Power noise, Pn, can be quantified by subtracting the baseline power, Pb, from the actual power, P:

$$Pn = P - Pb$$

where Pn is the power noise, P is the actual power, and Pb is the baseline power.

Power noise can be cancelled by using energy storage, with complementary energy being stored and returned as Ps.

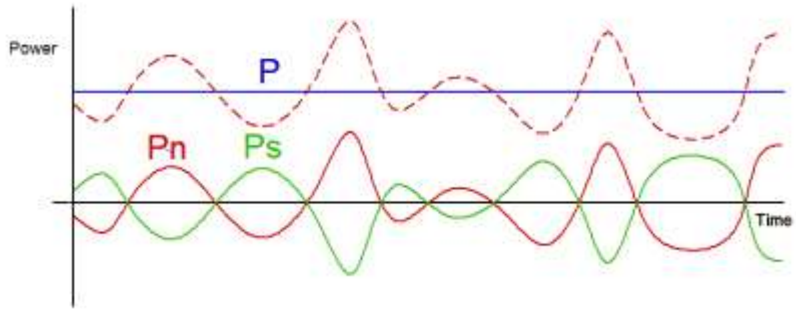

Virtual Energy Storage

Electrical storage; such as batteries, flywheels, capacitors, and so forth; is expensive. An alternative is to modulate a load or a power source to have intentional complementary power ripple that cancels the original power noise. In the graph below, a noisy power source, P1, is matched with a second power source, P2, that is controlled so that its output power has intentional complementary power ripple that cancels the power noise in P1.

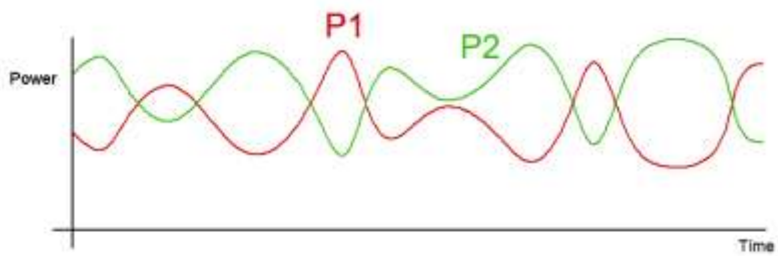

With a smart grid, it would be possible to measure the power at one site and transmit the quantified power noise measurement to one or more other sites so that they could modulate their power to cancel noise. However, that has a number of disadvantages, the most obvious being data-link failure or failure to validate the data. If there were a delay in the data, the cancellation would be compromised and a 180º lag could make the noise twice as bad. In addition, the noise may be from multiple sources, and any given complementary site may be able to handle only a portion of the power noise. A preferred implementation may be to have loads be responsive to power variations with strong load-lines.

Power noise can occur at any frequency. Close to line frequency, it is a power quality issue and should be resolved at the source. At a somewhat lower frequency, it is regulation services. At much lower frequency, it may be a daily power variation or even a seasonal one.

Sources of Power Noise

In the workshop, the power variation of “alternative” power sources was illustrated and discussed in several of the presentations. A lot of power noise originates in the loads and it can be synchronized there as well, to compliment and cancel other sources of noise.

Anytime that the power into a load varies from its baseline or scheduled power, that is power noise. It is analogous to emi, and we are used to identifying sources of emi in any part of a circuit. Conducted emi often originates in a load (a power supply), and many techniques for controlling emi are well known and most are adaptable to power noise control.

A thermostat that switches heaters off and on creates power noise, at the fundamental frequency of the thermostat (in the order of 0.5 mHz to 20 mHz), but also there is very rich higher frequency power noise content in the sharp rise and fall times of the power. The trend to variable-speed compressors and fans tend to quiet this source of power noise.

A PID temperature controller with a solid-state relay is a popular arrangement. With very good PID controllers available on eBay for $15, their use is becoming widespread. A PID controller may switch quite a large load several times a minute, generating noise in the 2 mHz range, again with sharp rise and fall times.

Power Noise Filters and Attenuators

It would be very difficult to completely cancel power noise from on-off switches, but a system with filters and attenuators can be helpful. In-line compensators can be effective filters and controlled loads can be effective power noise attenuators.

Emi is often filtered through a combination of series filters to block noise and shunt filters to attenuate emi or snubbers to absorb it. That is a good guiding principle for power noise filtering as well.

As with emi, the best defense against power noise is not to generate it in the first place. If it cannot be avoided, try to limit it to as narrow a bandwidth as possible. By contrast, the bandwidth of the shunt paths should be as great as possible.

Loads and power sources that generate power noise preferably have their response sandboxed over a narrow frequency range that is as low as practical. Virtual energy storage in the form of loads and power sources that generate sympathetic power variations that cancel power noise should have as wide a bandwidth as practical, one that extends to higher frequencies.

Compensation

As the load current flows through the distribution system, it causes voltage drops in the impedance elements of the network. Without compensation, the voltage at the load would droop badly, making it very difficult to maintain voltage regulation.

Compensation overcomes the voltage drops so that power can flow, while maintaining nearly constant voltage at the load. There is a direct correlation between the load current and the amount of compensation needed. The compensation must change as the current changes, so the dynamic characteristics of the current are important. Preferably, the current is not allowed to change more rapidly than the compensation can be adjusted. Adjustment can be made quite quickly in the power converters and inverters, as they have a fast response. High di/dt should be kept away from the slower responding equipment in the bulk power and transmission end of the grid.

There are two ways of overcoming the voltage drops of the load current in the distribution system, current compensation and voltage compensation. Usually, both are used.

Current compensation injects current into the system at one or more nodes. The compensation current flows through the impedances of the distribution system and causes compensating voltage drops that cancel the voltage drop of the load current. Current compensation is ineffective if the ratio of reactance to resistance %X/%R is low. Unfortunately, injected current sometimes flows the wrong direction and can be counter-productive.

Voltage compensation adds voltage to the system that compensates for the voltage drop of the load current as well as any misapplied current compensation. Because the voltage is nominally fixed at the load end, voltage compensation is applied at the source end or in series within the distribution system. Quadrature voltage compensation can simulate a series capacitor and can cancel the effect of series inductance.

Source voltage variation can also be cancelled by the varying the compensation, within limits, but it is preferred that the source voltage is controlled so that this is not necessary.

VARs are Tricky

It is very likely that VARs are second nature for the utilities and ISOs. VARs are unfamiliar to power converter designers, and they definitely are not intuitively obvious, so we will explore them in some depth.

Current Compensation

Current compensation, or shunt compensation, injects current into the power transmission system at one or more nodes. Capacitors are widely used to inject quadrature currents. Often special equipment is used to inject currents deep inside the grid, far from the load. The compensation may be fixed capacitors that are switched in and out, but synchronous condensers and static compensators are being used more and more often. While much more expensive, they have the advantage of being continuously variable as the load current changes.

SPICE Model

We are giving emphasis to compensation because it is so important to the operation of the grid. A SPICE model helps illustrate what is happening.

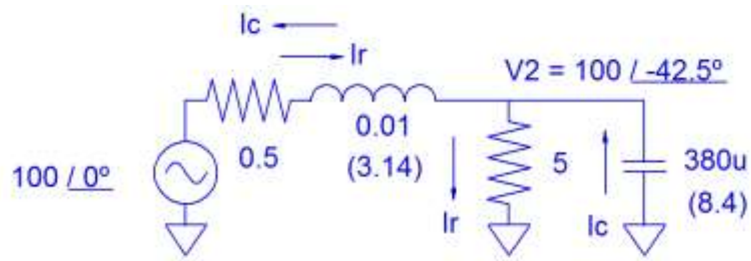

The circuit shown above was modeled in SPICE with a series resistance and reactance (inductance) to simulate a very simple power distribution system. The capacitor in parallel with the load was varied by trial and error until the output voltage V2 equaled the input voltage. 380 µF provided the correct quadrature current. The phase difference was noted and the vector diagram of the voltages was drawn. The leg y1 is the voltage drop in the series impedance, X = 3.14 Ω, and R = 0.5 Ω. Taking the vector sum, the resultant is 3.18 Ω.

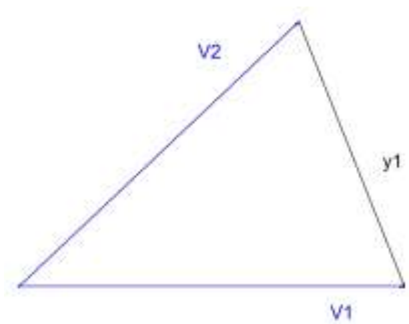

By measuring the voltage drop, y3, across the resistance and drop, y2, across the capacitance; the vector y1 can be resolved into quadrature voltage components as shown below.

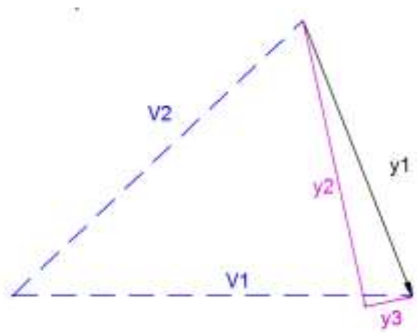

The SPICE model was manipulated to show how the load current and compensation current resolve in the series impedance. The method is beyond the scope of this report, but y2 and y3 can be resolved into their respective quadrature voltages, as shown on the left below. The vector y2 is the voltage across the inductance, and it resolves into the voltage y2r due to the load resister current and the voltage y2c due to the compensation capacitor current. The vector y3 is the voltage across the resistance (so it is lossy) and it resolves into the voltage y3r due to the load resister current and the voltage y3c due to the compensation capacitor current.

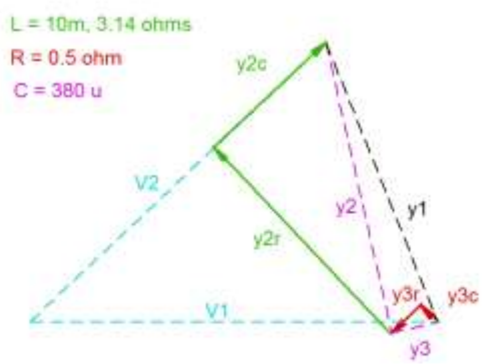

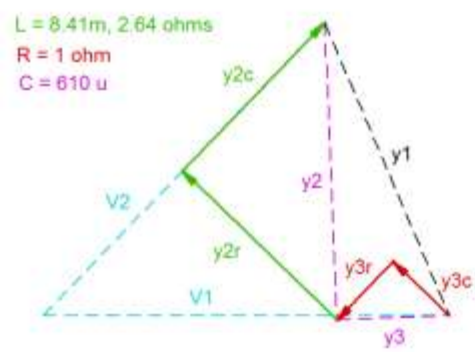

The values of the series impedances were changed in the SPICE model to X = 2.64 Ω and R = 1 Ω. The resultant impedance dropped from 3.18 Ω to 2.82 Ω. The compensation capacitor was then adjusted by trial and error until V2 once again equaled 100 V. It took 610 µF, compared to 380 µF with the original series impedances. The new vector diagram is on the right above.

The ratio of inductance, L, to resistance, R, makes a big difference. In the vector diagram on the left, the L/R ratio is 6.28 and the voltage across the series resistance is quite small, and so are the losses, noting that the loss varies as V2. In the vector diagram on the right, the L/R ratio is 2.64. The capacitor has to be much larger even though the total impedance remained almost the same, and the losses are 4.4 times greater. Unless the series reactance (inductance) is much larger than the series resistance, current compensation is ineffective. This is because the effect of the voltage drop, y3r, of the compensation current in the series resistance bucks the voltage drop, y2c, of the compensation current in the series inductance in the vector of the output voltage, V2.

Compensation at the Load

We explored applying the compensation current at the load. This is seldom done in practice, but if it is successfully adopted, it has market potential for PSMA members.

There are several reasons to consider this. The most important, in our view, is that the compensation current must vary as the load current changes, and what better place to do that than at the load itself? Another consideration is that separate hardware usually would not be needed to generate the quadrature currents – adjusting the phase angle does that nicely.

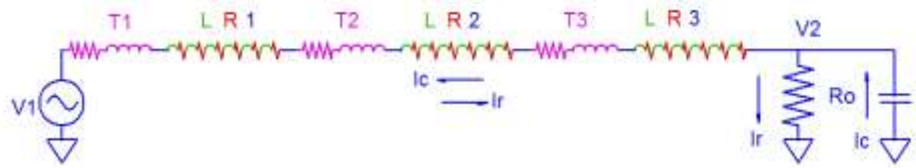

A problem is that the current flows the entire length of the circuit, with widely varying impedance and impedance ratios, L/R. The very high-voltage transmission transformers typically are 25 %X and 1 %R. Both the transmission lines and the distribution lines tend to have a much lower L/R ratio, typically in the order 2:1 or 1:1. Distribution transformers also have a relatively low L/R ratio. It is impossible to generate a compensation current that is right for the entire circuit. However, if the compensation current is correct for the higher L/R, it is moderately helpful for the rest of the circuit.

Tentatively, we chose a compensation current that is correct for an impedance of %X = 10% and %R = 2%. Theoretically, this is too much compensation for the high-voltage transformers, but there very likely will be uncompensated parallel loads that will dilute the quadrature compensation currents.

Another SPICE Model

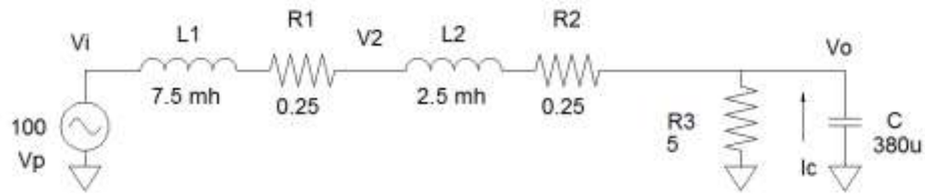

We set up a model circuit with series reactance and resistance, but this time used two inductors and two resistors, with a much higher L/R ratio at the source end. We then used trial and error to find the optimum compensation capacitor so that the output voltage, Vo, equaled the input voltage, Vi. Not surprisingly, it was once again 380 µF, as the total inductance and resistance matched the original SPICE model above.

When the simulation was run, we plotted the input voltage, Vi; the output voltage, Vo; and the intermediary node voltage, V2; noting that they are all the same.

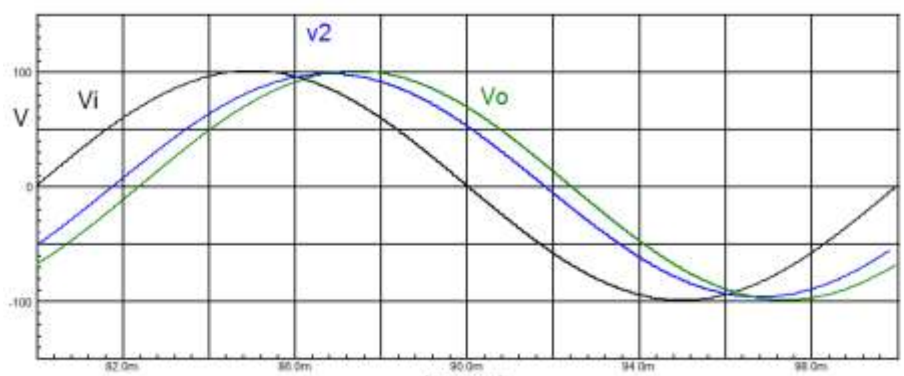

We then moved the capacitor to the node in the middle of the series impedances and again optimized the capacitor by trial and error until the output voltage, Vo, equaled the input voltage, Vi. It required a slightly larger capacitor, 395µF vs. 380µF, not a significant difference. However, it did not work as well.

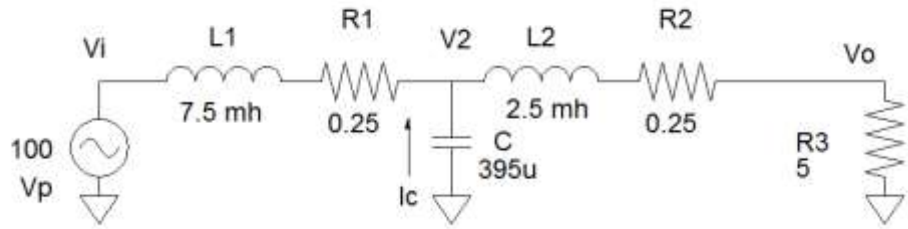

By putting the capacitor at the intermediate node, V2 at that node was boosted significantly. With the capacitor at the load, there was no voltage drop through the second pair of impedances, from V2 to Vo, but there is significant voltage drop in the revised circuit, on the order of 10%, 10 V in a 100 V circuit.

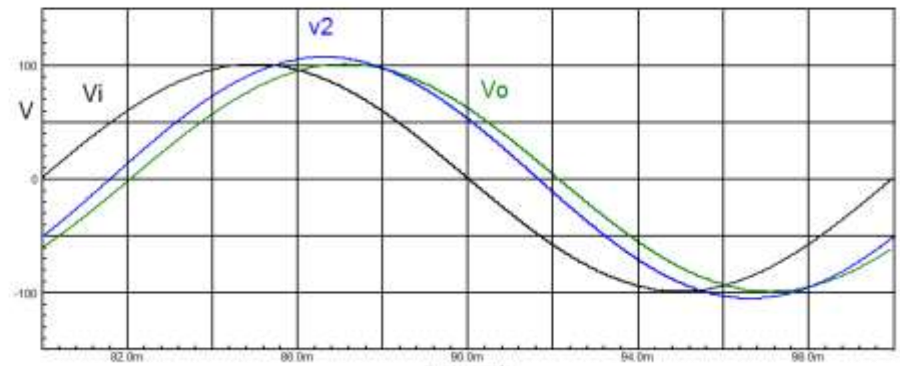

A graph of the currents shows that the capacitor current is significantly higher as well. Ica is the capacitor current in the first SPICE model with the capacitor at the load, and Icb is the capacitor current with the capacitor at the intermediary node. The explanation for the difference is that the current injected into the intermediary node flows both ways and it is counter-productive in the output circuits.

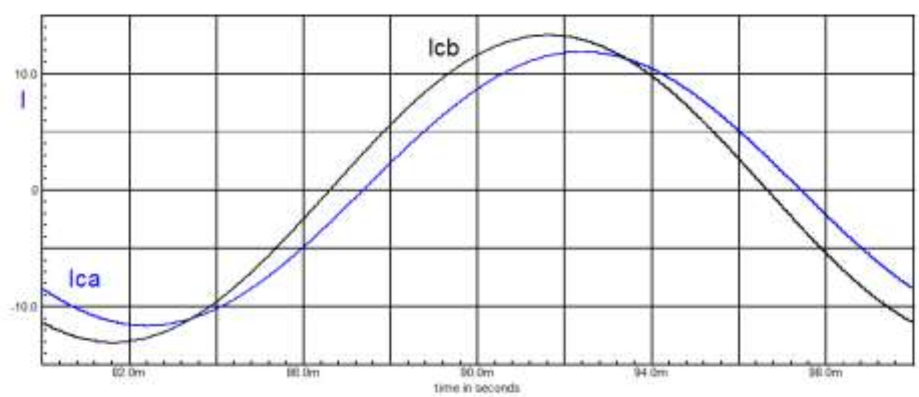

This validates that current compensation injection at the load is preferred.

Compensation for a Distributed Power Source

A power source injects current into the transmission network and it needs to be compensated as well, but with opposite polarity.

More important, if distributed power sources have compensation that is complementary to the load compensation, when they are both on the distribution network, the compensation currents cancel to the extent that the current is equal, leaving only the net compensation current corresponding to the net current from the utility.

Voltage Compensation

Voltage compensation, or series compensation, adds voltage to the system to compensate for voltage drops in the impedances due to the load current. Often, voltage opposing the voltage drop in the inductance is generated using series capacitors, and it is possible to tune a particular transmission to have no net reactance.

The voltage compensation can be applied all at once at the bulk generation end of the line, but since the power transmission system is a network, distributed voltage compensation is more effective. It can be used effectively to steer power from a particular source or through a particular path.

In the distribution network, the ratio of the reactance to resistance is too low for current compensation to be effective. However, voltage compensation can be used to cancel both the resistance and the inductance as a function of current in a particular path, ensuring minimum impedance if the distribution system is a virtual voltage source.

Series Compensator

A series compensator that is inductively coupled into the transmission line allows four quadrant compensation and is an effective filter. It is highly desirable to keep higher frequency components off the transmission lines.

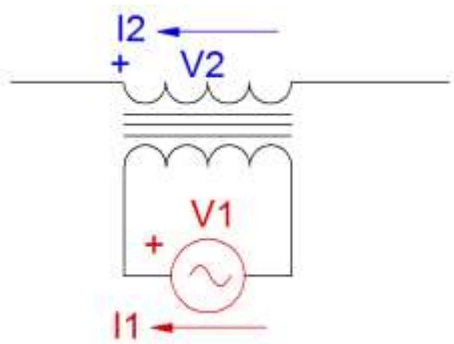

The voltage of the compensator V1 reflects as a series voltage in the transmission line as the turns-ratio; but, more important, is that in a transformer, the sum of the currents is zero. I2 is a faithful replication of I1, including all of its frequency components or lack thereof. If I1 has no higher frequency components, neither does I2. Moreover, the filtering can be selective, blocking higher frequency components in one direction, but passing them in the other direction. It is all in the control of the compensation voltage V1 and current I1.

One Percent 1% Compensator

The voltage that can be coupled is limited by the flux capacity of the transformer. Conversely, the size of the transformer and the voltage of the compensator can be limited if the compensation is limited. This has the advantage of making the compensator economical, but it also limits the effect on the circuit if it is hacked or otherwise malfunctions. Smaller units are easier to place strategically and several can be used along a power line for optimization.

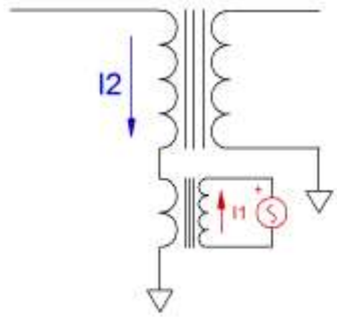

The series compensator has the disadvantage that the secondary has a high voltage to ground, so the insulation has to be rated accordingly. However, there may be two places where the series circuit is at ground potential, and that is at the neutral of Y connected transformers. A 1% compensator sees no more than 1% of the line voltage, so its insulation can be rated at that voltage. Also, its flux capacity is limited, so it is small. It does, however, have to conduct the full primary current.

The series voltage, however, must never be allowed to saturate the transformer core, as it would badly distort the waveform and introduce unacceptable noise. Rather, the voltage must be controlled so that the maximum volt-second product is limited, while maintaining the correct waveform for compensation. If the limited waveform cannot cancel all of a particular harmonic, it at least attenuates it.

Also, higher-frequency noise requires less flux and tends to balance much of its flux requirement (its volt-seconds) in the fundamental line frequency. Even a one-percent compensator should be effective at attenuating significant higher-frequency line noise. For example, a one-percent compensator at line frequency should be able to compensate third-harmonic noise up to 3%.

Virtual Compensator

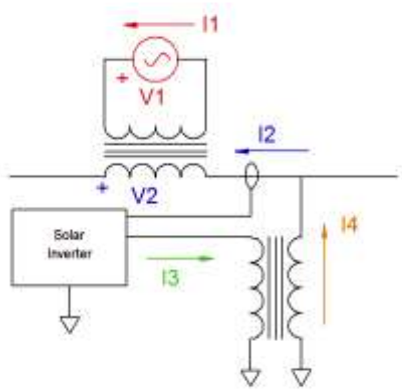

It is fairly obvious to generate compensation current at a power source or load that is at the end of the power line. It is less obvious to co-locate series compensation and corrective shunt compensation where a power source or load taps into a distribution system. However, if there is an inverter or power converter already in operation, it makes sense to do so.

As an example, the solar inverter in the figure above generates power represented by current I3. The step-up transformer conducts the power into the distribution line at high voltage. There is no reason that a current sense cannot measure the currents in the line, assess whether the compensation is correct for the measured load current, and generate supplemental current in the inverter that augments current I3 so the net current I2 is correctly compensated.

Although it is an independent function, a series compensator can be at the same location, powered by the same control circuits. Accurate compensation and filtering of the line current can be accomplished without requiring separate equipment by modifying the control algorithm of equipment already operating.

At the house level, a battery charger could compensate itself. With a current sense on the power line, it can compensate downstream loads as well, with its current capacity. The same is true of a solar inverter. They may need higher rating, but the compensation currents are relatively small, so it is minimal.

Power Management

Much of the motivation for the smart grid is to manage power more effectively and to maintain a regulated voltage. Our concern is that we need to have a robust default mode for when the smart grid central computer is offline or the data cannot be validated.

Unfortunately, much of the equipment on the distribution end of the grid that is inherently most effective at voltage regulation, resistors and induction motors, is being replaced with “solid-state” loads that are largely ineffective for regulation, or worse, may cause instability. “Alternative energy” is the bogeyman, but it is only one example of a much larger problem.

Load-Line

Control by load-line is preferred, as it is spontaneous, stable, and absorbs changes in power without the delay or the stability issues of adjusting power sources through feedback. Some caution is needed, however, as we design the load-line with synthetic impedances. The stability of their controls needs to be reviewed carefully. First, we derive a load-line for an ideal grid.

We start very simply, with apologies to the expert, because it is such an important topic. We develop our idealized grid load-line through a series of baby-steps, starting with a voltage source, Vs, and a source resistance, Rs. Assuming that voltage regulation should be controlled to ±5%, for discussion purposes, we propose a voltage load line of ±1%, or 2% total, no load to full load. The source resistor, Rs, is chosen so that the voltage drop is 2% at full-load current. That is, %R = 2%.

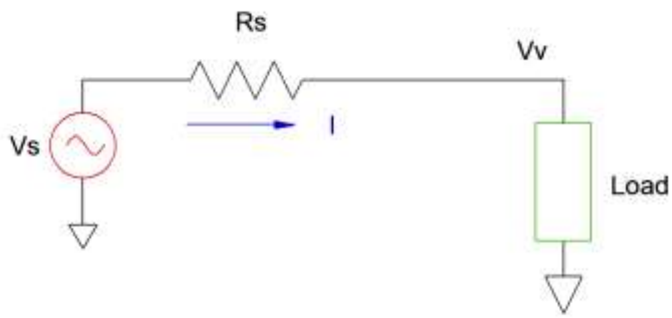

We want the load-lines of sources and loads to be complementary, so we define the load-line for reverse current flow as well as 2%R.

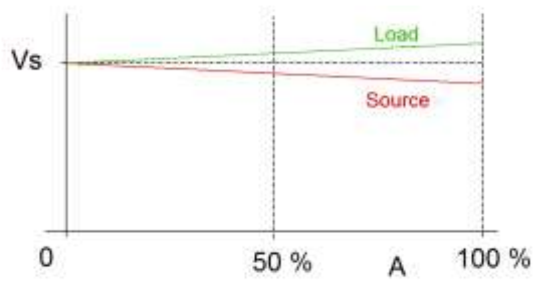

We need to maintain nominal voltage on the grid with balanced source and load power, so in our control algorithm, we adjust the power sources to account for the load current. Each power source, each load, and each energy storage unit adjusts its voltage source to equal the nominal voltage, plus the voltage drop in its series impedance. A power source adjusts its voltage upward to drive a current toward the virtual voltage, Vv, and a load adjusts its voltage downward to drive a current out of the virtual voltage, Vv.

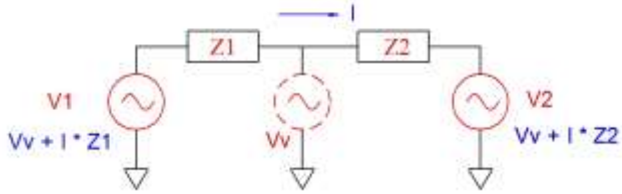

With one source and one load, this scheme provides a stable virtual voltage, Vv; but we need a load line for the sum of all power sources, power loads, and energy storage units. We use percent impedance, %Z, for the source and load impedances. In this way, regardless of the size of the source or load, a certain percent voltage drop relates to a specific percent of its rated load. If the percent impedance, %Z, is 2%R; the voltage drop is 1% at 50% load regardless of size. We assume that all controlled voltage sources know the percent load that they should provide and that they adjust their source voltage accordingly.

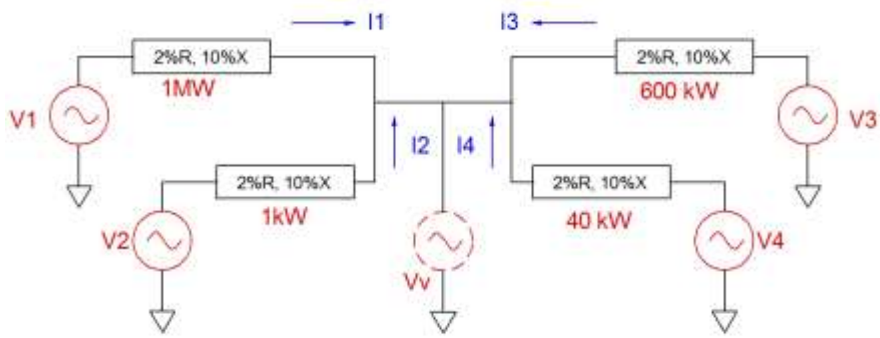

Many controlled sources, loads, and energy storage units can be parallel. As long as each has a designated percent loading and follows the algorithm Vn = Vv + In*Zn, the virtual voltage is maintained. However, we need to add a method for each source to determine what its percent loading should be.

Information about the global percentage loading is accomplished by biasing the load-line voltage by the percent loading. A representative value of 2%R + 10%X is used here to develop the example. A future study can optimize them, but once a suitable algorithm is found, it should be the model for all controlled sources, loads and energy storage units. We select a 2% voltage bias 1% above the nominal virtual voltage, Vv, for no load and 1% below Vv for 100% load.

We also bias the phase shift and we can be much more aggressive with the slope of the load-line for phase shift because phase shift is largely transparent to legacy equipment and does not affect the voltage magnitude. For discussion, we select a 10% bias, 5% leading for no load, and 5% lagging for 100% load. A 10% phase angle is sin-1(0.1) = 0.10 rad = 4.5 º.

Remembering that $Vr \approx \sqrt{\Delta V^2 + (V_a * \theta)^2}$, a load-line change from 0 to 100% is equivalent to approximately 10.8% change in voltage across the series impedance, which is synthesized to be 2%R + 10%X.

Load-Line Extremes

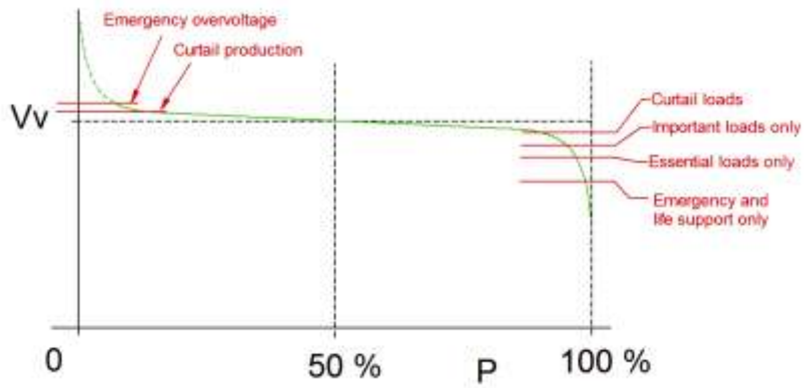

The voltage versus power load-line should droop significantly as overload is approached, though the graph above is exaggerated for illustration. If the algorithm is symmetrical about the 50% point, the voltage will rise if there is excess power production, but, at some point, the voltage must be clamped to prevent damage to equipment.

As overload is approached; production, of course, should be maximized and the load should be curtailed. As the voltage droops more, all but essential loads, then even emergency loads, should be disconnected. A smart meter responsive to the load-line could enforce that. The most likely scenario for extreme overload is a circuit that has been "islanded" – disconnected from the grid with limited local resources. Even a small local generator could provide enough power for emergency services such as life support, emergency communication, exit lighting, and cell-phone charging; though the magnetizing currents of all of the connected transformers may be a significant load.

The voltage droops if the simulated source voltage is reduced, but another option is to change the impedance as full load is approached. If the 2%R + 10%X transitions to 10%R, 0%X, the voltage drop across the synthetic impedance is real (in phase) with the same result. That also reduces the compensation current, so the line drop is greater.

Percent Power Impedance

The percent impedance, %Z, is the right choice for determining the compensation current, as explained earlier, because compensation is based on the load current. However, we cannot use current to control power because of the impedance inversion of controlled sources and loads, so the algorithm for power control is based upon the percent power impedance, %pZ, as defined earlier. Fortunately, controlled loads tend to be power factor corrected, so they "look resistive" at line frequency, where current compensation is a factor. Power control tends to be at a lower frequency, so the frequency response of the algorithm can be modified to make the transition. A transition frequency of about 5 Hz may be about right.

Voltage Regulation vs. Voltage Compensation

It is not unusual to see voltage regulators proposed for a branch circuit. If the voltage of a branch circuit is controlled to a precise value, regardless of the input voltage, the information of the voltage load-line is lost and equipment that cannot sense phase angle has no parameter that shows incipient overloading. The correct voltage compensation for a series regulator is to compensate only for voltage drop due to the current, I * Z. This requires knowledge of the impedance and Dr. Sun showed us the importance of quantifying it. As a control parameter, likely it can be optimized in the smart grid using voltage and current sensing in the power distribution system. The information will serve the system as a default setting when the computer is down.

Synthesizing the Algorithm and Control

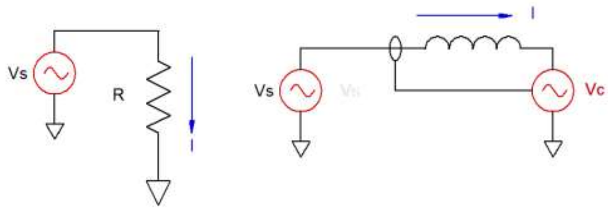

A generic "solid-state" power converter block diagram is shown above, the familiar power factor controlled (PFC) power converter. PFC converters control the input current to make the input "look resistive," that is, the current varies proportionately to the voltage. Any other algorithm for the current can be substituted. To generate quadrature currents, four-quadrant operation is necessary, or a small parallel current source can be used to supply the reverse current when the resultant is negative.

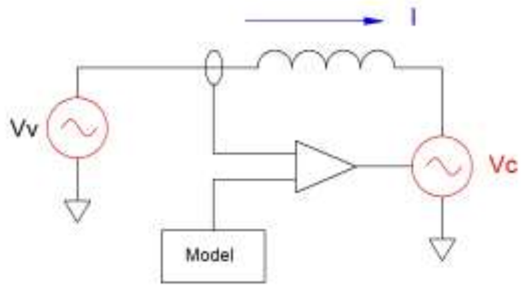

The model can be applied to inverters as well. The direction of the current is immaterial to the control – the voltage can decrease the control voltage, Vc, relative to the virtual voltage, Vv, to cause a controlled current to flow into the device as a power supply or it can increase the control voltage, Vc, so that it flows out. In a PFC power converter, the inductor is a tiny impedance compared to the resistance being modeled, so tiny variations of Vc with respect to Vv are sufficient to generate any current within the limitation of the power electronics.

Ideally, the control of the current loop can overcome any impedance variation, so the resulting circuit is a current source.

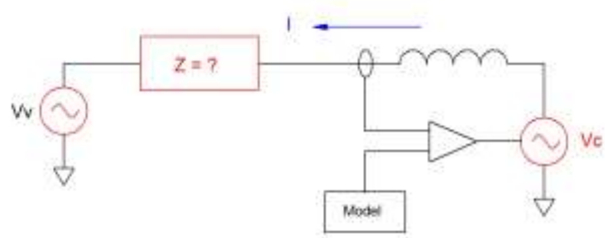

If a current is modeled to be the correct contribution to the virtual voltage, Vv; the unknown intervening impedance theoretically is irrelevant, within the limitation of the control voltage, Vc, to drive the current, I. Assuming that the unknown impedance may have significant inductance, however, rolling off the current may be worthwhile.

Most of our attention is directed toward local power sources, loads, and energy storage in the distribution grid, where the impedance is quite low and can be ignored or compensated. More of an issue will be the connections to the grid. If the local distribution is well controlled and accommodates most of the power noise that is locally generated, the load from the grid should be much smoother and fairly well compensated. The utilities use the wide area measurement system (WAMS) and sophisticated controls backed up by ISO, so a compatible interface algorithm should not be a great problem at the sub-station.

Managing Power Noise

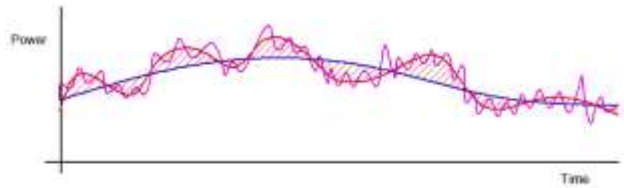

To review, power noise is unwanted variations in the power, either from a power source, such as a solar collector, or into a load. Higher-frequency variations are more the domain of power noise management, as the lower-frequency variations are really resource management and more likely to be determined by non-engineering considerations, like marketing. At line frequency, compensation and power quality are issues, and these problems should be corrected at the source. So, power noise management might cover a range from 5 mHz to 5 Hz, more or less.

As with emi, the best strategy for managing power noise is not to generate it or to incentivize the source of power noise, be it a power source or a load, to control it on site. This might be substituting proportional control for on-off control, or it could be using integrated energy storage to cancel the power noise. To the extent possible, the frequency response of power noise generators should be sandboxed to the smallest frequency range practical, preferably at the lowest frequency.

Solar Energy

At the workshop, the variation in power production from solar panels was featured prominently as a source of power noise, so we will use it as an example for discussions.

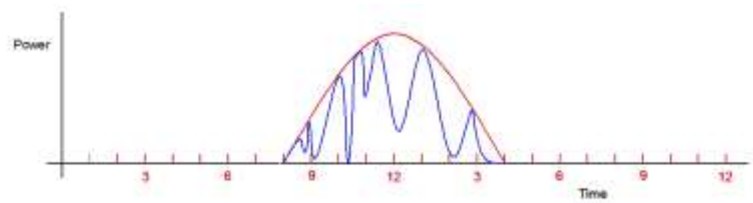

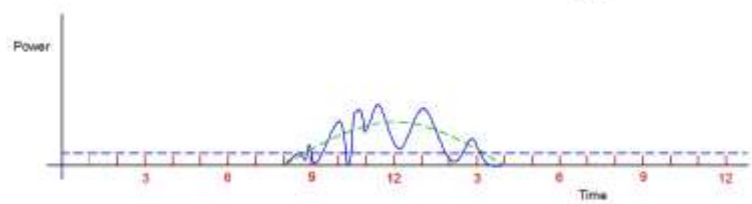

The top graph shows the envelope of maximum power production on a bright cloudless day. Within the envelope, there is a representative curve of power production on a stormy day. The second graph shows a load configured for virtual energy storage to cancel a part of the power noise of the solar panel. Within the distribution grid; other power sources, loads, and energy storage units should absorb the rest.

There are some significant issues illustrated here.

First, the solar panel does not generate power noise when it is dark. When it is operating, the power noise is limited by its maximum power envelope. So, if a load is to be configured to absorb part or all of the power noise, scheduling is important.

Second, a load that absorbs power noise must have a quiescent input power as an operating point to accept variations above or below the average input. If power noise regulation is to be ensured for the entire day, the quiescent input power must be scheduled for the same timeframe. In the second graph, the dashed line from 8:00 am to 4:00 pm might represent a scheduled power input to the load.

A load may absorb power noise in an opportunistic way, perhaps just shaving the peaks. However, if the rate is not consistent with the need for regulation, the regulation service may end. As an example, a homeowner may heat a tank of water to provide domestic heat and hot water through the night. If it reaches its maximum temperature at noon, he can absorb no more. So, to ensure regulation services as needed, the power input must be scheduled for the duration needed and in the amount needed.

Finally, to absorb maximum power noise, the peak power input should be two times the scheduled power, more or less. With regulated voltage, that means the peak current will be two times, more or less. (The numbers look even worse with an option to have an even load over the entire day.) The equipment must have at least twice the current capacity and twice the thermal dissipation. It will cost a lot more, which is good for the PSMA member who sells the equipment, but it is something that needs to be rewarded.

Conversely, the user who generates the power noise should pay for the cost of equipment that absorbs it. Power noise should be metered.

Synchronizing Power Noise Absorption

For power noise cancellation, the timing must be right as there is a potential to make it worse if it gets out of phase. Obviously, power should be used sympathetically when there is a surge in power production, and use should be curtailed when production sags. This can happen at quite a high frequency, perhaps as high as 1 Hz with a fast PID controller. In the workshop, power noise from solar panels was more of a concern, and it looked as if the dominant frequency was in the range of 500 µHz to 5 mHz.
A change in percent loading of “The Grid” shows as a voltage modulation of the load-line, and that modulation can be detected. If the voltage is averaged over a period roughly corresponding to the lower frequency of interest, and that is subtracted from the instantaneous voltage, the difference is reflective of the power variation, and the resulting signal can be used to synchronize the virtual energy source.

However, the phase angle modulation signal is much stronger. Absolute phase angle is hard to measure, but again if the phase angle magnitude is averaged over a period roughly corresponding to the lower frequency of interest, and that is subtracted from the instantaneous phase angle, the difference reflects the power variation and can be used for control.

Power Marketing

Most of the relationships discussed above are based upon using the percent power impedance, %pZ. The basis is the full-load rated power, and the resulting voltage and phase angle load lines reflect the percentage power as a percentage of available power.

If the amount of power for sale is substituted for the maximum power rating to determine a percent sale impedance, %$Z; all of the equations can be used. If at some point, a supplier decides to make more power available, the power basis is revised. If many providers decide to sell more power, the entire load line can shift, reflecting the more open market.

This raises the possibility of using the percent sale impedance as a buy-sell tool for establishing a spot market for power. It also raises the possibility that a sudden shift in the spot price of power could cause sudden shifts in the grid, just what power management is trying to avoid. It also raises the specter of price manipulation or hacking, with all that entails. As engineers, we would do well to avoid all of that. However, having recognized the possibility of using a percent market impedance %$Z, it is disclosed.

The Economies of "Free"

If the California initiative for zero net energy houses is implemented on a wide scale, there will be times that excess energy is produced, at least regionally.

Most of the zero net energy houses will use solar panels, as that is about the only energy source available for a suburban environment. Solar panels have peak production at noon, solar time, on a sunny day, falling off significantly by mid-afternoon and not resuming significant power production again until mid-morning the next day. If most of the energy production is during those six hours, more or less, the peak energy production will be about four times the average energy consumption. On a bright, sunny day, there may be a significant excess of "free" energy.

The economies of "free" are interesting to contemplate. Normally, we would be very concerned about conversion losses if we were buying energy to store and recover later. When it is "free," the conversion losses can be much greater and still have benefit. After all, any recovered energy is a bonus.

Whether excess energy from solar panels is "free" may be a matter of debate, but the operating cost is very low – most of the expense is capital expense that is amortized at a steady rate. Unlike a coal- or gas-fired plant, no fuel is used. Solar power that remains uncollected is lost forever.

It might not be economical to produce hydrogen from electricity from conventional power plants… but if the electricity is "free," any hydrogen production is a bonus.

If nothing else, use the excess electricity to heat or cool something. Melt some salt. Heat some water. Make some ice. Take a long shower. If it is not used or stored, the energy is lost forever. So, even a marginal use is worthwhile.

Leaky Bucket

Using much the same logic, it is apparent that storing "free" energy in a "leaky bucket" is worthwhile as well. Some of it may be recovered before it leaks away and any energy recovered is a bonus.

Cleanness and Greenness

In 1903, New York City passed laws, effective in 1908, banning coal-fired trains within the city. This was an early example of smog prevention. California has been increasingly aggressive in limiting pollution. It is likely that the cities of the future will limit energy consumption to "clean" fuels and the cleanest (at the point of consumption) are electricity and hydrogen. (Hydrogen can be made from water and electricity, electricity can be made from hydrogen in a fuel cell, so they are somewhat fungible. Is it possible that electric power transmission of the future will be partly through gas pipelines?)

Concerns about global warming will lead to increased pressure to use energy from "green" sources.

Cleanness and greenness were not discussed much at the workshop, but several questions address it. It is foreseeable that cleanness and greenness will be increasingly important in the future. We need to know how to quantify them.

There are two processes of interest, we think. One is the amount of carbon extracted from the atmosphere in the production of a unit of energy with its antithesis being the amount of carbon extracted from the ground. The second is the amount of carbon sequestered when fuel is burned with its antithesis being the amount that is released into the atmosphere.

The present dominant cycle is the least green and clean, extracting carbon from the earth and releasing it into the atmosphere. It should get a strong negative score. Two popular alternatives being discussed are burning biofuels. Theoretically, that is null, but in practice it may be somewhat ungreen due to the dirty energy being used in its production.

Extracting carbon from the ground and sequestering the resulting carbon dioxide is also, theoretically, a null.

The cycle that has real benefit is to burn biofuels and sequester the carbon.

The metric could be straight engineering quantities, kg per kwh, for example. We like to put scores on measurements, like degrees Fahrenheit or the octane rating for gasoline. So, we propose letting burning hydrogen score a zero, and burning methane with release to the atmosphere be a negative 100. Sequestering the same amount of carbon would be a positive 100. When methane is extracted from the ground, it too is a negative 100, but if it is derived from plant based material, it is a positive 100.

Extracting methane from biomass and sequestering the carbon dioxide would score 200.

Islands and Emergency Response

Some of the discussion at the workshop was about sustaining power during a storm and restoring it quickly afterwards. We have given it some thought.

A likely scenario during and after a major storm has many downed power lines and large parts of the distribution system separated from the bulk power sources. Yet it is likely that many parts of this system have some capacity to generate electricity, and that should be exploited. Therefore, a first effort should be to disconnect areas with downed power lines from those where they are intact, a smart grid might automate much of that.

One important island may be subdivisions with buried power. Regardless, once separated, those lines can be energized, even with meager power. At a minimum, everyone should be able to charge their phones and have some light. If there is solar power, they may be able to run important loads at noon, such as a water pump or a freezer. Having the intelligence to manage power under these circumstances may be an excellent use of smart meters.

Under this scenario, the first priority would be to remove downed power lines. If that was all that they did during the first sweep, it would go quite quickly and, once cleared, many islands could be reenergized at some level. With planning and investment in some infrastructure, it should be possible to provide sufficient power to large neighborhoods sufficient for life-support equipment and cell phone charging.

A major problem in restoring power is that a lot of the attached equipment turns on immediately. Every refrigerator, every water pump, every hot water heater… all come on at once. The resulting power surge very likely exceeds normal usage. Smart meters could wait after power restoration before reconnecting and then do so with staggered delay. The smart meter could also sense if power were limited, as the voltage would be low, and connect only essential services.

One consideration in restoring power to isolated islands already energized is synchronization. This will be greatly facilitated once the frequency is locked to one clock. In the meantime, the smart grid will have extensive data-link communication. It is quite possible that the data-link will be down during the storm, and the islands should be able to function in a strong default mode, but it is not unreasonable to require that communications be restored before reconnecting islands. That could ensure synchronization.

Appendix A — Questions and Answers

Appendix A — Questions and Answers

Introduction

During the workshop, the panel discussions were short due to extra time given to the presentations. This was good in that the presentations were valuable, but it reduced the opportunities for questions and discussion. In addition, we did not schedule time for "break-out" sessions, as many workshops do.

We asked all participants to provide questions for the presenters. There are a variety of the questions, and many seem to be off-target or naive. Some questions clearly have an "agenda" and are inappropriate. Questions broadly fit into a few categories:

1. **Dumb Questions** – ones that are trivial, off target, or where the answers are found in the workshop presentations. However, Google provides some insight: "*The only stupid question is the one left unasked.*" Carl Sagan said it well: "*There are naive questions, tedious questions, ill-phrased questions, questions put after inadequate self-criticism. But every question is a cry to understand the world. There is no such thing as a dumb question.*"

   So, we take every question seriously and to try to find a context and re-phrasing that makes it worthwhile. An elementary question sounds dumb to the expert, but don't belittle the novice trying to learn.

2. **Repeated questions**. Many questions are on cyber-security. With so many, questions are consolidated to offer fewer, more targeted questions.

3. **Overly technical questions**. Complete answers are not expected for the report, but the interest suggests that some further discussion is worthwhile. Mostly, they identify a need for a good reference. Some will become topics of future workshops.

4. **Missed questions**. Some topics are missed entirely, so we added some questions.

Questions that had answers in the presentations are answered by reviewing the transcripts, slides, and the video recording of the workshop. These answers are incorporated into the appropriate appendix.

Unanswered questions may be the most important. An objective of the workshop is to bring out topics that are not being addressed in other forums. The unanswered questions show where work is needed. They may become topics of a future workshop.

The Emperor's clothes. Nobody asked why the Emperor had no clothes. Some questions come close.

Questions

These are questions asked of the workshop presenters. They are filtered and consolidated. Questions that had answers in the presentations are not included here.

Names in [brackets] are references to speakers or their slides. They may or may not be the one to whom the questions were addressed.

- Why have a grid at all, let alone a "Smart Grid"? Wouldn't we be better off investing in local power production with independent islands?
  - *We need a smart grid that includes centralized and distributed generation for the most efficient and reliable power delivery*.
- Why have an ac grid? Isn't dc more efficient for transmission?
- Is there an accurate predictive model available for the grid? The effect of independent energy sources does not seem to be well understood.
- [Dr. Abdul-Rahman, slide 16] "Single Contingency Coverage –The interconnected power system must be operated at all times so that instability, uncontrolled separation, cascading outages, and uncontrolled voltage deviation will not occur as a result of any *single contingency*."

  Is computer failure a "single contingency?"

  Is loss of the data-link a "single contingency?"

  Is hacking or sabotage a "single contingency?"
- [Dr. Gelling, talk, no slide] "One of the constraints —- in a power system is — [moving] power for economic reasons. —Unfortunately, power moves because of physical laws.

  Dr. Gelling was probably talking about choosing among suppliers or transmission lines, but of more concern to us, the consumers, is our economies. The idea of shutting off our appliances so that the utilities can sell the power for a higher price to those who "opt out" (likely to be big banks, corporate headquarters and millionaires) does not seem fair. Is this a concern in planning the smart grid?

  It seems likely there could be large-scale disconnects if power prices hit a very high price tier. With demand reduced, the price would go lower, and there may reconnections in large numbers, with the cycle repeating. Is this a concern in planning the smart grid? Has this been observed?
- [Dr. Abdul-Rahman] "The grid operator is responsible for maintaining reliability and continuity of service on the electric grid at all times." [Dr. Gellings] "One of the constraints in a power system is people like to move power for economic reasons."

  Who does an ISO answer to?

  Who owns the ISO; who / what finances it?

  How does an ISO mediate economic interests if they may compromise reliability and continuity of service?

- [J. Pollet, presentation] "Saudi Aramco had 30,000 computers wiped out in one day. It was an insider that did it. They set it all on a certain time so that it all happened on the same day."

  A large percentage of down-time in data centers is attributable to rouge employees.

  Will the smart grid ensure that no person or group in a utility can set "logic bombs" or rogue code to cause a contingency?

- [J. Pollet, presentation] "Making a dinner reservation has more security than to send a command to an RTU."

  Would it be feasible to route the more critical commands through a bank or other institution that has robust security?

- [D. Symanski, slide 10] A bullet under "Secondary Topics" is "–Cyber Security & Privacy."

  In view of J. Pollet's scary talk, is cyber security still considered to be a "secondary topic"?

  - All electric power utilities are concerned about and are working on cyber security. Slide 10 shows the order of topics being addressed during the smart grid demonstrations being conducted at this time.

- If the smart grid turns off my refrigerator "for 15 minutes," but the data connection is lost before it is turned back on, will it stay off? Can I get it going again, before everything spoils?

  - *The "default" operation when communication is lost is to be "on."*

- What are the liability implications for manufacturers whose equipment is hacked and causes problems? Could equipment that is too easy to hacked be subject to recall?

- [Dennis Symanski (paraphrased)] "We need to find out what works well and what doesn't; what works well but is too expensive to implement; what technologies actually cancel the benefits of another technology."

  We could add "what works now, but won't work in the future" and "what doesn't work now, but may work in the future."

  - *Ok to add these situations*

  Is there a tabulation or score-card? (If not, we should start one.)

- [Dr. K. Smedley] "So by DOE, what is smart grid? The technology by two-way communication technology and computer processing to make this grid smarter. …One of the most daunting challenges facing the utility is the surge of data that will result as we modernize, managing thousands of times more data than we do today."

  How many data centers will it take for Southern California Edison to keep track of all of the programmable thermostats and "plug-loads"?

  How much power will they consume?

- In the future smart grid, what performance is expected as a default when the computer and/or the data-link is down or data cannot be validated?

  What discrimination is expected of the "Smart Meter?" For example, if it receives a command from the central computer to shut off the refrigerator, but the voltage is on the high side of normal, should the command be carried out anyway?

  - *These smart grid demonstrations will help to get to these answers*

- [Dr. Rajagopalan, slides 26 and 27] In monetizing "regulation services," please define "regulation services?" How are regulation services measured and paid for?

- [Dr. Sun] "We cannot control distributed power sources in the same way we do today because of the large number and the speed of response. …A large number of fast-response units necessitates autonomous control. We can't do it with central management anymore.

  What is the loop delay, from sensing a remote sensor to controlling a remote load? Is it fast enough to react to emergencies, such as the sudden loss of a transmission line?

  Has this been analyzed for phase margin and stability, using tools such as Nyquist plots and Bode diagrams?

- [Dr. Gelling, slide 34] The slide shows a frequency shift due to energy imbalance with a range of 58 to 62 Hz.

  Is it really that bad?

  Will there still be a frequency shift as more "solid-state" inverters are providing power? [Dr Gelling, slide 57 shows 50 % "digital power" by 2020].

- [Dr. Rajagopalan] "We have never had a single energy storage system where you plug it in and it works. …Power electronics, surprisingly, is an area that needs more research. …50% of the energy storage system cost is the power converters. …System integrators, sometimes, have no idea of either power electronics or battery management systems. …The electric utilities have no experience putting these systems on their feet. …We have had some high-priority fires in the last year or so and it's the power converter that appears to be the source of many of these fires."

  [Dr. Sun] We can characterize the grid by impedance. …By connecting the converters with the grid, we effectively create a local minor feedback loop.

  Energy storage inverters are often located near alternative energy inverters, essentially putting two or more voltage sources in parallel and expecting them to share the load in a stable manner. Power converter designers know that paralleling voltage sources is very problematic.

  Has this been analyzed for phase margin and stability, using tools such as Nyquist plots and Bode diagrams?

- [Dr. Von Dollen, slides 3 and 4] The slides clearly show more weather-related events.

  How much of this is attributable to more severe weather itself? How much is that more equipment is exposed?

- Putting power lines underground certainly is more expensive in many places. However, where soil conditions are right, distribution cable laid with a ditch-witch might be less expensive than placing poles and stinging wires overhead, and the reactance per mile may be much lower. Considering hurricane resistance might tip the balance in places like the Mississippi Delta. Is this being done or is it even being considered?
  - *Utilities are continuously analyzing the cost-benefit of putting distribution circuits underground. Some will be put underground and some will remain overhead.*

- [Dr Gelling, slide 18] The slide shows a 3-volt drop in the primary feeder; 2 volts in the transformer and 2 volts more getting to the house.

  Is this mostly reactance? It suggests that distribution transformers have about 2 % impedance.

- [Dr. Gelling talk, no slide] "85% of all outages are from the distribution system. By the way, 85% of those are due to squirrels and other animals. The squirrel is public enemy number one."
- 85% of 85% is 72% of outages are due to squirrels. So, does 72 percent of your "outage research" go to squirrel damage prevention? It seems like a solvable problem. Maybe if this were crowd-sourced, some good, low-tech solutions would emerge.

- [Dr. Gellings wrap-up, slide 11]. "Establish a Joint PSMA-EPRI Task Force on Smart Products for the Smart Grid?"

  Please give us further information on this idea. It sounds like a very good suggestion.

Why Have "The Grid"?

Ed Herbert

In an ideal world, each house or business would have its own power source, sufficient for its needs. Ideally, it would be free, "clean," and "green." If not each house, at least each community or region would be a self-sufficient island.

So, why have "The Grid," let alone a "Smart Grid"? Wouldn't we be better served building up local resources?

"The Grid" exists and will remain because the best sources of power are not located where power is consumed. Value is a consideration and "cleanness" and "greenness" will become increasingly important. If power is less expensive, cleaner, and/or greener at a distant location, we need "The Grid" to move it to the population centers.

It does not seem likely that cities will ever be able to generate enough power locally for the needs of dense concentrations of people and businesses, particularly if it must be clean and green.

The grid also is a backup power source. We are accustomed to thinking of "backup" as being a generator set at the facility, but where do we get power if the local generator fails or a storm destroys most of the local solar panels and wind turbines? It must be imported, by "The Grid."

Why Have an AC Grid?

The "war of the currents" rages still, with many proponents of a dc grid. However, the ease of shifting voltage in transformers remains a compelling reason for ac. Another factor, less recognized, is that the ac grid allows power flow between nodes of nearly equal voltage by controlling the phase difference. This is not possible with dc.

15 % Renewables? — 33 %? —— It will be 400 %!

California leads the way and often the rest of the world follows. A requirement that new homes and businesses have "zero net energy" suggests, as its logical conclusion, that peak renewable energy production will be at least 400% of average energy consumption.

Consider a new bedroom community with many large new housing tracts. Each of the houses must be zero net energy. The only widely available power source for individual houses in a suburban setting is solar panels, and they only make power when the sun shines. Most of their power production is between 9:00 a.m. and 3:00 p.m. (sundial time), with a peak power production at 12:00 noon.

If a whole day's energy production is concentrated in 6 hours, then, on a per hour basis, the power production must be four times the average per-hour power consumption, or 400%. Subtracting the power use during that time leaves an excess of 300%. Many suburban dwellers will have excess power and they will be eager to sell it to the population centers, so 400% probably is a low estimate.

There is an accelerated adoption of "alternative energy" now. Once the original goals are met, be they 35% by 2020 or some modification thereof, the solar panel producers and installers are not going to shut down and go out of business. They will keep producing and installing them as new houses are built and there will be increased incentives to retrofit older buildings and provide off-site alternative energy production on those sites where solar panels cannot be installed.

What Works?

This is the most important question!

Dennis Symanski (paraphrased): "*We need to find out what works well and what doesn't; what works well but is too expensive to implement; what technologies actually cancel the benefits of another technology*." We could add, "*what works now, but won't work in the future*" and "*what doesn't work now, but may work in the future*."

Picking and choosing, or getting on board?

This workshop focuses on the smart grid and, in particular, the interface between the grid and power converters. We aspire to give guidance to the PSMA members and we may be able to give advice to the smart grid, whatever it is, and whoever controls it.

As engineers, we all want to modify it, to make it better. Can we? Or has the train left the station and we better just jump on board? How do we know? As Reinhold Niebuhr (1892-1971) said:

"God, give me grace to accept with serenity the things that cannot be changed, courage to change the things which should be changed, and the wisdom to distinguish the one from the other." — Reinhold Niebuhr (1892-1971)

Our goal is to provide some courage and wisdom and to challenge the boundaries of what must be accepted with serenity if it is bad engineering. It is a daunting task.

Looking Back

It is interesting to look back at technologies that are commonplace now, but that once were dismissed as *"too expensive"* or not useable for some reason. A partial list:

- Synchronous rectifiers – too expensive
- Power factor correction – too expensive
- Si power MOSFETs – at $65 each, way too expensive
- PWM power supplies – too noisy ever to be used with processors
- Power supplies with power density greater than 3 W/cu in – too dense to be reliable

It is risky business predicting the future, but it is worthwhile to try. If some technology is doomed by foreseeable evolution, then it is best not to spend time and resources on it. Trends in power semiconductors, particularly GaN and SiC devices, suggest that power electronics can play an increasingly important role in the future grid, so solutions that are "too expensive" may be viable soon.

The most challenging is to identify problems for which there are no known solutions, or problems that have known solutions that are seriously flawed. There is a substantial risk that this latter category will be implemented on a wide scale because there are no obvious alternatives. A flawed solution can become "cast in concrete."

We need a table of what works and what doesn't. We start with one example. Others may be found by digging in the presentations and this report.

Example: Frequency Regulation

Frequency regulation works reasonably well now. There are problems with it that we live with, but there are clear signs that that the present way of controlling frequency will not work for the future as more sources become "solid state" and more loads become controlled ("constant power").

For reasons explained in more detail in the section on frequency regulation, it is likely that frequency regulation will evolve to the point that frequency does not vary at all. The grid will become synchronized to a master clock.

Some equipment exploits frequency variations as an indication of grid loading or overloading. It is very likely that this will not work in the future. It would be very unfortunate if this equipment and its algorithms became so well established that frequency variation was preserved just to keep it working.

Synchronization to a master clock may be a long way off, but during the evolution, the variations in frequency will be significantly reduced. Equipment will have to be able to handle frequency variation for the foreseeable future and may have to handle it as a "contingency" forever.

Appendix B — Smart Grid Links

Appendix B — Smart Grid Links

To assist readers who want to learn more about the Smart Grid and current activities by government and industry organizations, the following links lead a plethora of information, limited only by the inquisitiveness of the reader.

EPRI	http://www.epri.com/Pages/Five-Years-Later---EPRI-Smart-Grid-Demos-Continue-to-Inform.aspx
NIST	http://www.nist.gov/el/smartgrid/sgprogram.cfm
UCI	http://www.apep.uci.edu/3/research/partnership_ISGD.aspx
Caiso	http://www.caiso.com/Pages/default.aspx
Emerson	http://www.emersonnetworkpower.com/en-US/Pages/EmersonSmartGrid.aspx
Smart Grid News	http://www.smartgridnews.com/
US Government	http://www.smartgrid.gov/the_smart_grid
Department of Energy	http://energy.gov/oe/technology-development/smart-grid
PG&E Smart Grid Deployment Plan	http://www.pge.com/includes/docs/pdfs/shared/edusafety/electric/SmartGridDeploymentPlan2011_06-30-11.pdf