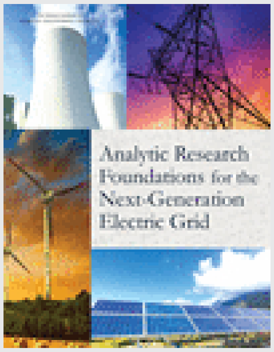

Analytic Research Foundations for the Next-Generation Electric Grid

DETAILS

160 pages | 8.5 x 11 | PAPERBACK
ISBN 978-0-309-39231-0 | DOI: 10.17226/21919

AUTHORS

Committee on Analytical Research Foundations for the Next-Generation Electric Grid; Board on Mathematical Sciences and Their Applications; Division on Engineering and Physical Sciences; National Academies of Sciences, Engineering, and Medicine

Analytic Research Foundations for the Next-Generation Electric Grid

Committee on Analytical Research Foundations for the Next-Generation Electric Grid

Board on Mathematical Sciences and Their Applications

Division on Engineering and Physical Sciences

The National Academies of
SCIENCES · ENGINEERING · MEDICINE

THE NATIONAL ACADEMIES PRESS
Washington, DC
www.nap.edu

THE NATIONAL ACADEMIES PRESS **500 Fifth Street, NW** **Washington, DC 20001**

This study was supported by Contract No. DE-PI0000010/DE-DT0005900 with the Department of Energy. Any opinions, findings, or conclusions expressed in this publication do not necessarily reflect the views of any organization or agency that provided support for the project.

International Standard Book Number-13: 978-0-309-39231-0
International Standard Book Number-10: 0-309-39231-4
Digital Object Identifier: 10.17226/21919

This report is available in limited quantities from the Board on Mathematical Sciences and Their Applications, 500 Fifth Street NW, Washington, DC 20001; bmsa@nas.edu, http://www.nas.edu/bmsa.

Additional copies of this report are available for sale from the National Academies Press, 500 Fifth Street NW, Keck 360, Washington, DC 20001; (800) 624-6242 or (202) 334-3313; http://www.nap.edu/.

Printed in the United States of America

Suggested citation: National Academies of Sciences, Engineering, and Medicine. 2016. *Analytic Research Foundations for the Next-Generation Electric Grid*. Washington, DC: The National Academies Press. doi:10.17226/21919.

The National Academies of
SCIENCES · ENGINEERING · MEDICINE

The **National Academy of Sciences** was established in 1863 by an Act of Congress, signed by President Lincoln, as a private, nongovernmental institution to advise the nation on issues related to science and technology. Members are elected by their peers for outstanding contributions to research. Dr. Ralph J. Cicerone is president.

The **National Academy of Engineering** was established in 1964 under the charter of the National Academy of Sciences to bring the practices of engineering to advising the nation. Members are elected by their peers for extraordinary contributions to engineering. Dr. C. D. Mote, Jr., is president.

The **National Academy of Medicine** (formerly the Institute of Medicine) was established in 1970 under the charter of the National Academy of Sciences to advise the nation on medical and health issues. Members are elected by their peers for distinguished contributions to medicine and health. Dr. Victor J. Dzau is president.

The three Academies work together as the **National Academies of Sciences, Engineering, and Medicine** to provide independent, objective analysis and advice to the nation and conduct other activities to solve complex problems and inform public policy decisions. The Academies also encourage education and research, recognize outstanding contributions to knowledge, and increase public understanding in matters of science, engineering, and medicine.

Learn more about the National Academies of Sciences, Engineering, and Medicine at **www.national-academies.org.**

COMMITTEE ON ANALYTICAL RESEARCH FOUNDATIONS FOR THE NEXT-GENERATION ELECTRIC GRID

JOHN GUCKENHEIMER, Cornell University, *Co-Chair*
THOMAS J. OVERBYE, University of Illinois, Urbana-Champaign, *Co-Chair*
DANIEL BIENSTOCK, Columbia University
ANJAN BOSE, Washington State University
TERRY BOSTON, PJM Interconnection, LLC
JEFFERY DAGLE, Pacific Northwest National Laboratory
MARIJA D. ILIC, Carnegie Mellon University
CHRISTOPHER K. JONES, University of North Carolina, Chapel Hill
FRANK P. KELLY, University of Cambridge
YANNIS G. KEVREKIDIS, Princeton University
RALPH D. MASIELLO, Quanta Technologies
JUAN C. MEZA, University of California, Merced
CYNTHIA RUDIN, Massachusetts Institute of Technology
ROBERT J. THOMAS, Cornell University
MARGARET H. WRIGHT, New York University

Staff

SCOTT T. WEIDMAN, Board Director
NEAL GLASSMAN, Study Director
ALAN CRANE, Senior Staff Scientist, Board on Energy and Environmental Systems
RODNEY N. HOWARD, Administrative Assistant

BOARD ON MATHEMATICAL SCIENCES AND THEIR APPLICATIONS

Acknowledgments

This report has been reviewed in draft form by individuals chosen for their diverse perspectives and technical expertise, in accordance with procedures approved by the Report Review Committee. The purpose of this independent review is to provide candid and critical comments that will assist the institution in making its published report as sound as possible and to ensure that the report meets institutional standards for objectivity, evidence, and responsiveness to the study charge. The review comments and draft manuscript remain confidential to protect the integrity of the deliberative process. We wish to thank the following individuals for their review of this report:

Mihai Anitescu, Argonne National Laboratory,
Seth Blumsack, Pennsylvania State University,
Karen Butler-Purry, Texas A&M University,
Michael Chertkov, Los Alamos National Laboratory,
Michael Ferris, University of Wisconsin,
Narain Hingorani, Independent Consultant,
Peter Sauer, University of Illinois,
Carson Taylor, Bonneville Power Administration (retired), and
Alan Washburn, Naval Postgraduate School.

Although the reviewers listed above have provided many constructive comments and suggestions, they were not asked to endorse the conclusions or recommendations of this report, nor did they see the final draft of this report before its release. The review of this report was overseen by M. Granger Morgan, Carnegie Mellon University, who was responsible for making certain that an independent examination of this report was carried out in accordance with institutional procedures and that all review comments were carefully considered. Responsibility for the final content of this report rests entirely with the authoring committee and the institution.

The committee also acknowledges the valuable contributions of the following individuals, who either provided input at its meetings, on which this report is based, or contributed in other ways:

Eyad Abed, National Science Foundation,
Mihai Anitescu, Argonne National Laboratory,
Emanuel Bernabeu, Dominion Technical Solutions, Inc.,

Gil Bindewald, Department of Energy (DOE),
Jan Caspary, Southwest Power Pool,
Hong Chen, Pennsylvania-New Jersey-Maryland Interconnection LLC (PJM),
Nicolai Dumitriu, PJM,
Timothy Heidel, Advanced Research Products Agency-Energy,
Adam Keech, PJM,
Sandy Landsberg, DOE,
Mark McGranaghan, Electric Power Research Institute,
David Meyer, DOE,
David Ortiz, DOE,
Russell Robertson, Grid Protection Alliance, and
Stan Williams, PJM.

The committee also organized the Workshop on Analytical Foundations for the Next-Generation Electric Grid, which took place on February 11-12, 2015, at the Arnold and Mabel Beckman Center of the National Academies of Sciences and Engineering in Irvine, California, and received valuable information from the speakers at that workshop. The agenda for that workshop is reprinted in Appendix A of this report, and a summary of the workshop, *Mathematical Sciences Research Challenges for the Next-Generation Electric Grid: Summary of a Workshop*, is available from http://www.nap.edu.

Contents

Summary

The electric grid is an indispensable critical infrastructure that people rely on every day. The Department of Energy (DOE) envisions that by 2030, the grid will have evolved into an intelligent energy system, a smart grid. By "smart," DOE anticipates that the grid will have the characteristics of (1) customer participation, (2) integration of all generation and storage options, (3) new markets and operations, (4) power quality for the 21st century, (5) asset optimization and operational efficiency, (6) self-healing from disturbances, and (7) resiliency against attacks and disasters.[1] The next-generation electric grid must be more flexible and resilient than today's. For example, the mix of generating sources will be more heterogeneous and will vary with time (e.g., contributions from solar and wind power will fluctuate), which in turn will require adjustments such as finer-scale scheduling and pricing. The availability of real-time data from automated distribution networks, smart metering systems, and phasor data hold out the promise of more precise tailoring of services and of control, but only to the extent that large-scale data can be analyzed nimbly.

Today, operating limits are set by off-line (i.e., non-real-time) analysis. Operators make control decisions, especially rapid ones after an untoward event, based on incomplete data. By contrast, the next-generation grid is envisioned to offer something closer to optimized utilization of assets, optimized pricing and scheduling (analogous to, say, time-varying pricing and decision making in Internet commerce), and improved reliability and product quality. In order to design, monitor, analyze, and control such a system, advanced mathematical capabilities must be developed to ensure optimal operation and robustness; the envisioned capabilities will not come about simply from advances in information technology. Within just one of the regional interconnects, a model may have to represent the behavior of hundreds of thousands of components and their complex interaction affecting the performance of the entire grid. While models of this size can be solved now, models where the number of components is many times larger cannot be solved with current technology. As the generating capacity becomes more heterogeneous due to the variety of renewable sources, the number of possible states of the overall system will increase. While the vision is to treat it as a single interdependent, integrated system, the complete system is multiscale (in both space and time) and multiphysics, is highly nonlinear, and has both discrete and continuous behaviors, putting an integrated view beyond current capabilities. In addition, the desire to better monitor and control the condition of the grid leads to large-scale flows of data that must in some cases be analyzed in real time. Creating decision-support

[1] U.S. Department of Energy, *Smart Grid Research & Development Multi-Year Program Plan (MYPP) 2010-2014—September 2012 Update*, September 2012, http://energy.gov/oe/downloads/smart-grid-rd-multi-year-program-plan-2010-2014-September-2012-update.

systems that can identify emerging problems and calculate corrective actions quickly is a nontrivial challenge. Decision-support tools for non-real-time tasks—such as pricing, load forecasting, design, and system optimization—also require new mathematical capabilities.

Mathematical modeling and control of the electric grid has been an active area of research for decades. However, in 1996 a major outage that affected 11 Western states and 2 Canadian provinces—coupled with emerging concerns that computers would malfunction after December 31, 1999—increased awareness of a lack of complete understanding of the overall system and its frailties. For several decades the Electric Power Research Institute funded a program of research to develop tools for recognizing early signs of instability and means to counter them. That research was largely of a mathematical nature.

More recently, DOE has been supporting research to develop the analytical and computational tools that will be necessary for the next-generation grid. Many frontier areas of the mathematical sciences are represented in that body of research. For example, the 2011 DOE conference Computational Needs for the Next Generation Electric Grid identified seven computational challenges associated with the operation and planning of the electric power system:

- Cloud computing,
- Hierarchical models,
- Analysis and planning for contingencies,
- Modeling of infrastructure interdependencies,
- Modeling and controlling multi-time-scale and multidimensional power systems,
- Optimization under uncertainty, and
- Unit commitment and economic dispatch.[2]

Other than the first of these, all require or could benefit from new tools from the mathematical sciences. In short, the future grid will rely on integrating advanced computation and massive data to create a better understanding that supports decision making. That future grid cannot be achieved simply by using the same mathematics on more powerful computers. Instead, the future will require new classes of models and algorithms, and those models must be amenable to coupling into an integrated system.

To complement this research specifically focused on tools for the next-generation grid, a range of potentially applicable research exists. Examples include research into the general topics of uncertainty quantification, simulation and analysis of complex adaptive systems, simulation and analysis of multi-time-scale systems, and methods for characterizing and controlling resilience and reliability. This research is taking place in a range of science and engineering disciplines. More generally, complex adaptive systems have been studied for several decades, and a good deal of "mathematical machinery" has been developed.

While many of the necessary tools are inherently mathematical, the best progress in these complex areas is achieved through multidisciplinary efforts, involving a community with diverse strengths and perspectives. In order to develop the next generation of tools required for the challenges of the smart grid, DOE commissioned the National Research Council (NRC)[3] to engage in a study with the following charge:

1. What are the critical areas of mathematical and computational research that must be addressed for the next-generation electric transmission and distribution (grid) system? Identify future needs. In what ways, if any, do current research efforts in these areas (including non-U.S. efforts) need to be adjusted or augmented?
2. Because this research frontier is best approached by a community that is truly multidisciplinary—including not only a cutting-edge knowledge of mathematics, statistics, and computation, but also a deep understanding of the emerging electric grid and of the questions that need answering to realize its potential—How

[2] U.S. Department of Energy, *Computational Needs for the Next Generation Electrical Grid. Proceedings April 19-20, 2011* (J.H. Eto and R.J. Thomas, eds.), LBNL-5105E, 2012, http://energy.gov/oe/downloads/proceedings-computational-needs-next-generation-electric-grid-workshop-april-19-20-2011.

[3] Effective July 1, 2015, the institution is called the National Academies of Sciences, Engineering, and Medicine. References in this report to the National Research Council are used in a historical context identifying programs prior to July 1.

can DOE help to effectively build this community? What mix of backgrounds is needed and how can the community be developed? How can DOE extend its reach beyond its existing ties?

To address this charge, the NRC assembled a committee of 15 members who collectively have academic, industrial, and national laboratory experience in both power systems and the relevant mathematical areas. In addition to meeting five times over the course of the study, a subset of the committee planned and ran a workshop on February 11-12, 2015, at the Arnold and Mabel Beckman Center of the National Academies in Irvine, California, to gain outside perspectives. The agenda of that workshop is appended to this report, and a published summary of that workshop is available at http:///www.nap.edu/21808.

The grid itself and the conditions under which it operates are changing, and the end state is uncertain. For example, new resources, especially intermittent renewable energy such as wind and solar, are likely to become more important, and these place new demands on controlling the grid to maintain reliability. At the same time, the increasing affordability of storage technology may ease controllability. Many technical improvements could be made to the grid, such as those noted below, but this report does not aim to cover them all nor does it presume one possible future grid scenario over another. The next-generation grid will require the efforts of many other scientific disciplines, including economics, social science, market planning, and risk analysis, to name a few, and some of these have significant mathematical content. After discussions with the study's sponsor, the committee interpreted its charge to focus on those mathematical research directions with broad impact, and which must be advanced in order to enable the next-generation grid, rather than to discuss the full range of possible improvements to the grid or mathematics that may play a secondary role in the next-generation grid's planning or management.

The committee also recognizes that acceptance of the conclusions and recommendations in this report by key industry segments—utilities, grid and market operators, market participants, software and system vendors, and the research community—is essential if productive research and development is to be conducted and successful results adopted. Some of the recommendations—for alternating current (ac) optimal power flow (ACOPF), for stochastic scheduling, and for integration of different time-scale models—need buy-in from key user segments to garner support for the research and development (R&D) efforts. Others, such as development and use of open data sets for testing, need buy-in to change in the way things are done. However, suggestions for obtaining this necessary buy-in were beyond the committee's charge.

This report contains the recommendations of the committee for new research and policies to improve the mathematical foundations for the next-generation grid.

In particular,

- New technologies for measurement and control of the grid are becoming available. Wide area measurement systems provide a much clearer picture of what is happening on the grid, which can be vital during disruptions, whether from equipment failure, weather conditions, or terrorist attack. Such systems send a huge amount of data to control centers, but the data are of limited use unless they can be analyzed and the results presented in a way suitable for timely decision making.
- Improved models of grid operation can also increase the efficiency of the grid, taking into account all the resources available and their characteristics; however, a systematic framework for modeling, defining performance objectives, ensuring control performance, and providing multidimensional optimization will be needed. If the grid is to operate in a stable way over many different kinds of disturbances or operating conditions, it will be necessary to introduce criteria for deploying more sensing and control in order to provide a more adaptive control strategy. These criteria include expense and extended time for replacement.
- Other mathematical and computational challenges arise from the integration of more alternative energy sources (e.g., wind and photovoltaics) into the system. Nonlinear alternating current ACOPF can be used to help reduce the risk of voltage collapse and enable lines to be used within the broader limits, and flexible ac transmission systems and storage technology can be used for eliminating stability-related line limits.
- Transmission and distribution are often planned and operated as separate systems, and there is little feedback between these separate systems beyond the transmission system operator's knowing the amount of power to be delivered and the distribution system operator's knowing what voltage to expect. As different

types of distributed energy resources, including generation, storage, and responsive demand, are embedded within the distribution network, different dynamic interactions between the transmission and distribution infrastructure may occur. One example is the synchronous and voltage stability issues of distributed generation that change the dynamic nature of the overall power system. It will be important in the future to establish more complete models that include the dynamic interactions between the transmission and distribution systems, including demand-responsive loads.

- In addition, there need to be better planning models for designing the sustainable deployment and utilization of distributed energy resources. Estimating future demand for grid electricity and the means to provide it entail uncertainty. New distributed-generation technologies move generation closer to where the electricity is consumed. Climate change will introduce several uncertainties affecting the grid. In addition to higher temperatures requiring increased air conditioning loads during peak hours, shifting rainfall patterns may affect the generation of hydroelectricity and the availability of cooling water for generating plants. The frequency of intense weather events may increase. Policies to reduce emissions of carbon dioxide, the main greenhouse gas, will affect generating sources. Better tools to provide more accurate forecasting are needed.
- Modeling and mitigation of high-impact, low-frequency events (including coordinated physical or cyberattack; pandemics; high-altitude electromagnetic pulses; and large-scale geomagnetic disturbances) is especially difficult because few very serious cases have been experienced. Outages from such events could affect tens of millions of people for months. Fundamental research in mathematics and computer science could yield dividends for predicting the consequences of such events and limiting their damage.

Ten years ago, few people could have predicted the current energy environment in the United States—from the concern for global warming, to the accelerated use of solar and wind power, to the country's near energy independence. Each of these developments, and others, will profoundly shape the future electric grid, and with it the analytic challenges and associated mathematical advances needed to cope with those developments. For that reason, the committee's recommendations do not focus on overcoming the inadequacies of specific algorithms or techniques. Rather, its recommendations are designed to help direct future research as the grid evolves and to give the nation's R&D infrastructure the tools it needs to effectively develop, test, and use this research. The committee's recommendations are in four areas: data availability, modeling capabilities, improved algorithms, and the organizational structure needed to integrate improvements in these areas and to make them accessible to a large community of researchers.

RECOMMENDATIONS

The recommendations and their discussion that follow are grouped so that those concerning the same subject are discussed together. For that reason, some are listed out of sequence.

Current algorithms do not scale well to the anticipated growth in the number of nodes in a large marketing area. One algorithm that does so is of particular importance: the mathematical programming formulation of the ACOPF problem. The problem is discussed in Chapter 2 and formulated mathematically in Chapter 7.

Recommendation 1: The Department of Energy should develop and test a full ac optimal power flow (ACOPF) model with an optimization algorithm using all nodes in the market area, taking advantage of supercomputers and parallel processing and respecting all thermal and voltage constraints. (Chapter 2)

The committee believes that available data are not sufficiently used by either the power industry or other potential researchers. It found that data used by the community of power engineers to develop and test algorithms are not available to the larger community because specialized software is needed to access them. To make the data available to a larger research community, the committee makes the following two recommendations:

Recommendation 2: The Federal Energy Regulatory Commission (FERC) should require that all text file formats used for the exchange of FERC715 power flow cases be fully publicly available. (Chapter 3)

Recommendation 3: The Federal Energy Regulatory Commission should require that descriptions of all models used in system-wide transient stability studies be fully public, including descriptions of any associated text file formats. (Chapter 3)

Most of the data being generated by the electric power industry are viewed as proprietary, both because they would reveal information about company operations and because they might reveal information useful to terrorists. For this reason, synthetic data that are sufficient to mirror real operations are required for future research.

Recommendation 4: Given the critical infrastructure nature of the electric grid and the critical need for developing advanced mathematical and computational tools and techniques that rely on realistic data for testing and validating those tools and techniques, the power research community, with government and industry support, should vigorously address ways to create, validate, and adopt synthetic data and make them freely available to the broader research community. (Chapter 6)

Furthermore,

Recommendation 9: The Department of Energy should sponsor additional efforts to create synthetic data libraries to facilitate studies of, and tool building for, the reliability and control of the future electric grid. (Chapter 8)

The committee believes that, for reasons that are not completely clear, the power industry is not making sufficient use of the data available to it—perhaps because it does not fully recognize the value of such data for both prediction and control. The power industry hires very few data scientists.

Recommendation 7: The Department of Energy should support research on data-driven approaches applied to the operations, planning, and maintenance of power systems. This would include better machine-learning models for reliability, comprehensible classification and regression, low-dimensional projections, visualization tools, clustering, and data assimilation. A partial goal of this research would be to quantify the value of the associated data. (Chapter 6)

The two mathematical areas that the committee believes will yield greatly improved capabilities are dynamical systems and mathematical programming, particularly nonlinear and nonconvex programming.

Recommendation 6: The Department of Energy should support research to extend dynamical systems theory and associated numerical methods to encompass classes of systems that include electric grids. In addition to simulation of realistic grid models, one goal of this research should be to identify problems that exemplify in their simplest forms the mathematical issues encountered in simulating nonlinear, discontinuous, and stochastic time-dependent dynamics of the power system. The problems should be implemented in computer models and archived in a freely available database, accompanied by thorough documentation written for both mathematicians and engineers. Large grid-sized problems that exemplify the difficulty in scaling the methods should be presented as well. (Chapter 4)

Recommendation 8: Order-of-magnitude improvements in nonlinear, nonconvex optimization algorithms are needed to enable their use in wholesale electricity market analysis and design for solving the ac optimal power flow problem. Such algorithms are essential to determine voltage magnitudes. Therefore the Department of Energy should provide enhanced support for fundamental research on nonlinear, nonconvex optimization algorithms. (Chapter 6)

There is a similarity between the electric grid and the climate system—both are sufficiently complex as to defy precise analysis. For that reason, the use of various kinds of machine learning, along with improved control and optimization algorithms, is important.

Recommendation 5: Integration of theory and computational methods from machine learning, dynamical systems, and control theory should be a high-priority research area. The Department of Energy should support such research, encouraging the use of real and synthetic phasor measurement unit data to facilitate applications to the power grid. Establishment of test-beds for physical experiments would provide valuable additional sources of data. (Chapter 6)

The committee believes that the electric generation research community would benefit from the availability of new open-source software.

Recommendation 10: The Department of Energy and National Science Foundation should sponsor the development of new open-source software for the next-generation electric grid research community. (Chapter 8)

Finally, the committee has found a need for coordination among a community broader than the national research laboratories.

Recommendation 11: In view of the importance of its efforts to coordinate power grid research at the national laboratories, the Department of Energy should broaden this coordination to include academic and industry researchers. (Chapter 8)

Recommendation 12: The Department of Energy should establish a National Electric Power Systems Research Center to address fundamental research challenges associated with analysis for the future electric system. The center would act as an interface between the power industry, government, and universities in developing new computational and mathematical solutions for data and modeling issues and in sharing valuable data. (Chapter 8)

1

Physical Structure of the Existing Grid and Current Trends

INTRODUCTION

Electricity is the lifeblood of modern society, and for the vast majority of people that electricity is obtained from large, interconnected power grids. Engineered to offer the ultimate in plug-and-play convenience, the wall outlet is actually the gateway to one of the world's largest and most complex machines. Starting in the early 1880s with Thomas Edison's Holborn Viaduct system in London and the Pearl Street Station in New York, serving a total of just 59 customers in lower Manhattan,[1] central station power rapidly developed so that within a decade electricity was ubiquitous in many cities around the world. In the decades that followed, high-voltage, interconnected power grids developed and many rural areas were electrified as well.

While the grid was initially fueled to a large extent by hydro (and still is in some countries such as Canada), in the United States coal was king, and the 20th century was powered by fossil fuels with up to 20 percent nuclear. Economies of scale resulted in most electric energy being supplied by large power plants. Control of the electric grid was centralized through exclusive franchises given to utilities, which in turn had an obligation to serve all existing and future customers. This relatively stable arrangement allowed numerous technical challenges to be overcome, resulting in the creation of the modern electric grid. Named by the National Academy of Engineering as the greatest achievement of the 20th century,[2] electrification has truly changed the world.

However, the grid that was developed in the 20th century, and the incremental improvements made since then, including its underlying analytic foundations, is no longer adequate to completely meet the needs of the 21st century. The next-generation electric grid must be more flexible and resilient. While fossil fuels will have their place for decades to come, the grid of the future will need to accommodate a wider mix of more intermittent generating sources such as wind and distributed solar photovoltaics. Some customers want more flexibility to choose their electricity supplier or even generate some of their own electricity, in addition to which a digital society requires much higher reliability. The availability of real-time data from automated distribution networks, smart metering systems, and phasor measurement units (PMUs) holds out the promise of more precise tailoring of the performance of the grid, but only to the extent that such large-scale data can be effectively utilized. Also, the electric grid is

[1] Con Edison, "A Brief History of Con Edison," http://www.coned.com/history/electricity.asp, accessed February 19, 2016.
[2] National Academy of Engineering, "Greatest Achievements of the 20th Century," http://www.greatachievements.org/, accessed February 19, 2016.

increasingly coupled to other infrastructures, including natural gas, water, transportation, and communication. In short, the greatest achievement of the 20th century needs to be reengineered to meet the needs of the 21st century.

Achieving this grid of the future will require effort on several fronts. Certainly there is a need for continued shorter-term engineering research and development, building on the existing analytic foundations for the grid. But there is also a need for more fundamental research to expand these analytic foundations. The purpose of this report is to provide guidance on the longer-term critical areas for research in mathematical and computational sciences that is needed for the next-generation grid. This chapter and Chapters 2 and 3 set the stage by providing a brief overview for the physical structure of the existing grid and some of the analyses that are essential for planning, evaluating, and operating the grid. Given the complexity of the existing grid, this introduction can only touch on some of the major topics and is certainly not comprehensive. More information is available in power systems textbooks such as Glover et al. (2012), Wood et al. (2013), Kundur (1994), or Van Cutsem and Vournas (2007).

BASIC GRID CONCEPTS

Interconnected Alternating Current Power Grids

While Edison's original power grid was direct current (dc), it soon became apparent that power having low dc voltages (about 100 V) could be distributed over only a few blocks. This is because power is equal to the product of the voltage and current, and with 19th century technology there was no practical way of changing dc voltages. Hence higher currents were required, and since the transmission losses are proportional to the square of the current times the line resistance, power could not be efficiently transmitted over long distances. The alternative, alternating current (ac) power, championed by Nikola Tesla and George Westinghouse, soon won out because the voltage could be easily changed by devices known as transformers (though lower Manhattan remained fully dc until 1928, with the last dc service remaining until 2007[3]). So by the 1890s, ac transmission lines operating at tens of kilovolts (kV) were transmitting electricity dozens of miles, such as from a hydro power plant at Niagara Falls to Buffalo. Twenty years later electricity was transmitted hundreds of miles at voltages of about 100 kV, reaching 735 kV in the 1960s. Today the highest ac voltage used in North America is 765 kV, while a 1,000-kV grid is being developed in China, and countries in the former USSR have operated lines up to 1,150 kV (Huang et al., 2009).

Excepting islands and some isolated systems, North America is powered by the four interconnections shown in Figure 1.1. Each operates at close to 60 Hz but runs asynchronously with the others. This means that electric energy cannot be directly transmitted between them. It can be transferred between the interconnects by using ac-dc-ac conversion, in which the ac power is first rectified to dc and then inverted back to 60 Hz.

Any electric power system has three major components: the generator that creates the electricity, the load that consumes it, and the wires that move the electricity from the generation to the load. The wires are usually subdivided into two parts: the high-voltage transmission system and the lower-voltage distribution system. A ballpark dividing line between the two is 100 kV. In North America just a handful of voltages are used for transmission (765, 500, 345, 230, 161, 138, and 115 kV). Figure 1.2 shows the U.S. transmission grid. Other countries often use different transmission voltages, such as 400 kV, with the highest commercial voltage transmitted over a 1,000-kV grid in China.

The transmission system is usually networked, so that any particular node in this system (known as a "bus") will have at least two incident lines. The advantage of a networked system is that loss of any single line would not result in a power outage. In some regions, a 69- or 46-kV subtransmission system, which may be networked, is also used.

The lower-voltage distribution system is usually radial, meaning there is just a single supply; networked distribution is sometimes used in urban areas. Typical distribution system voltage levels include 34.5, 13.8, 12.4, 4.16, and 2.4 kV. Distribution lines are often called feeders. Additional transformers step the voltage down to the load supply voltages of usually less than 1 kV (commonly 480 V for commercial and 120/240 V for residential customers).

[3] J. Lee, "Off Goes the Power Current Started by Thomas Edison," New York Times blog, March 4, 2011, http://cityroom.blogs.nytimes.com/2007/11/14/off-goes-the-power-current-started-by-thomas-edison/?_r=0.

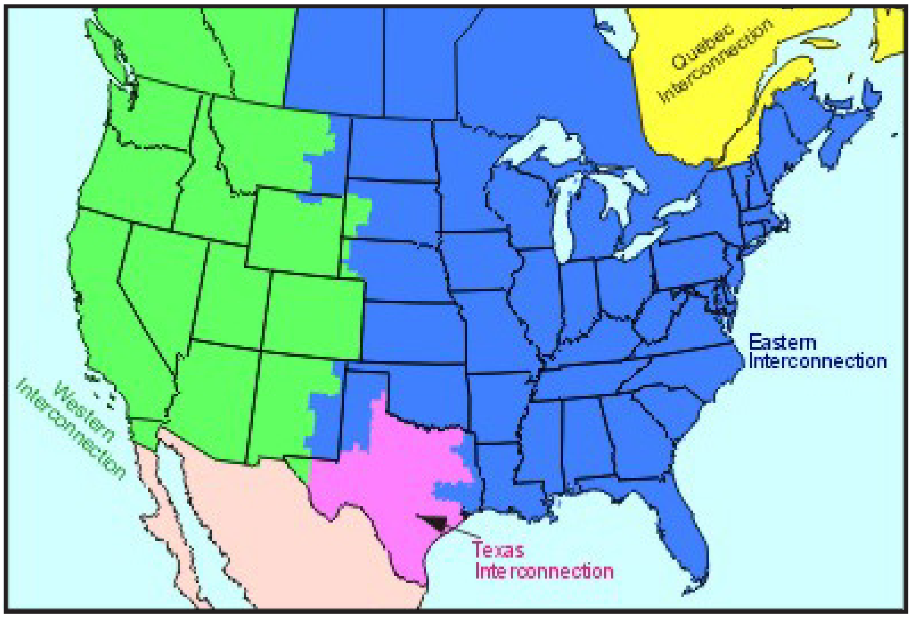

FIGURE 1.1 North America's electricity interconnections. SOURCE: NERC. This information from the North American Electric Reliability Corporation's website is the property of the North American Electric Reliability Corporation and is available at http://www.nerc.com/docs/oc/rs/NERC%20Balancing%20and%20Frequency%20Control%20040520111.pdf. This content may not be reproduced in whole or any part without the prior express written permission of the North American Electric Reliability Corporation.

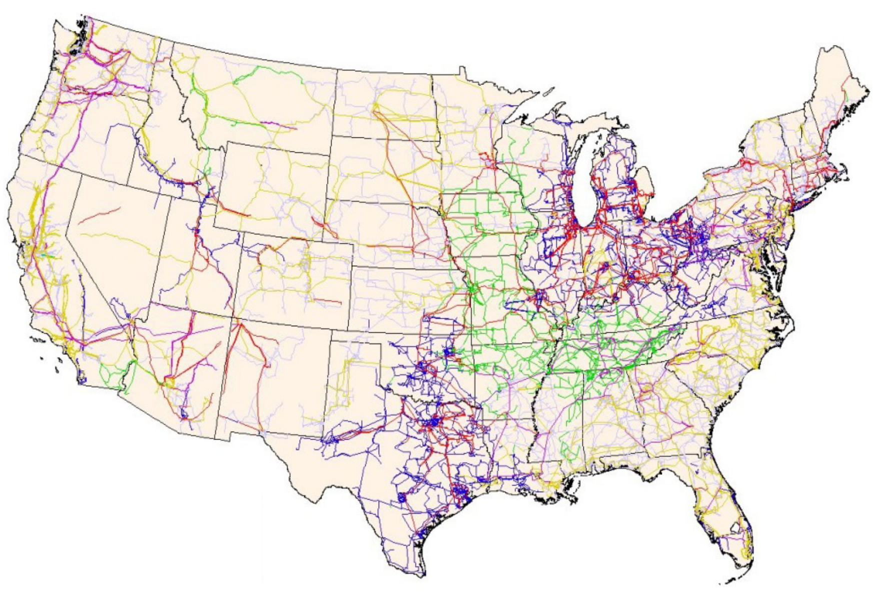

FIGURE 1.2 U.S. transmission grid. SOURCE: Created by Rolypolyman, available at https://commons.wikimedia.org/wiki/File:UnitedStatesPowerGrid.jpg.

While ac transmission is widely used, the reactance[4] and susceptance[5] of the 50- or 60-Hz lines without compensation or other remediation limit their ability to transfer power long distances overhead (e.g., no farther than 400 miles) and even shorter distances in underground/undersea cables (no farther than 15 miles). The alternative is to use high-voltage dc (HVDC), which eliminates the reactance and susceptance. Operating at up to several hundred kilovolts in cables and up to 800 kV overhead, HVDC can transmit power more than 1,000 miles. One disadvantage of HVDC is the cost associated with the converters to rectify the ac to dc and then invert the dc back to ac. Also, there are challenges in integrating HVDC into the existing ac grid.

Commercial generator voltages are usually relatively low, ranging from perhaps 600 V for a wind turbine to 25 kV for a thermal power plant. Most of these generators are then connected to the high-voltage transmission system through step-up transformers. The high transmission voltages allow power to be transmitted hundreds of miles with low losses—total transmission system losses are perhaps 3 percent in the Eastern Interconnection and 5 percent in the Western Interconnection. With the advent of distributed photovoltaics, more generation is being directly connected to the distribution system, sometimes with supply voltages as low as 120/240 V for residential connections. Figure 1.3 shows the general distribution of load (white) and generation (magenta) in North America. Notice that in the East the load is more evenly distributed, with the generation closer to the load (except in Northeast Canada); in the West, much of the load is on the coast, with the generation spread throughout the interconnect.

Large-scale interconnects have two significant advantages. The first is reliability. By interconnecting hundreds or thousands of large generators in a network of high-voltage transmission lines, the failure of a single generator or transmission line is usually inconsequential. The second is economic. By being part of an interconnected grid, electric utilities can take advantage of variations in the electric load levels and differing generation costs to buy and sell electricity across the interconnect (a topic that is more fully discussed in Chapter 2). This provides incentive to operate the transmission grid so as to maximize the amount of electric power that can be transmitted. However, large interconnects also have the undesirable side effect that problems in one part of the grid can rapidly propagate across a wide region, resulting in the potential for large-scale blackouts such as occurred in the Eastern Interconnection on August 14, 2003. Hence there is a need to optimally plan and operate what amounts to a giant electric circuit so as to maximize the benefits while minimizing the risks.

Power Grid Time Scales

Anyone considering the study of electric power systems needs to be aware of the wide range in time scales associated with grid modeling and the ramification of this range on the associated techniques for models and analyses. Figure 1.4 presents some of these time scales, with longer term planning extending the figure to the right, out to many years. To quote University of Wisconsin statistician George Box, "Essentially, all models are wrong, but some are useful. However, the approximate nature of the model must always be borne in mind" (Box and Draper, 1987, p. 424). Using a model that is useful for one time scale for another time scale might be either needless overkill or downright erroneous.

As an example, a key aspect of power system design is what is known as insulation coordination—designing the grid to adequately protect power system equipment from the transient overvoltages caused by lightning strikes and switching surges. This requires dynamic models of the system response using time steps on the order of microseconds. Since voltages and currents propagate down the transmission lines at velocities near the speed of light (186,000 miles per second or 300 m/μsec), it is important to model the delays that occur as these waves propagate down the lines. Thus on a microsecond time scale the response of the grid becomes decoupled since what occurs at one location on the grid does not instantaneously affect more distant locations, allowing for the use of distributed simulation, a technique that makes simultaneous use of multiple arithmetic processors to reduce the time required to complete the simulation. This is quite different from the coupled algebraic equations that will be introduced in the next sections to model the transmission system in the slower time frames.

[4] Reactance is the opposition of a circuit element to a change in current or voltage.
[5] Susceptance is the reciprocal of reactance.

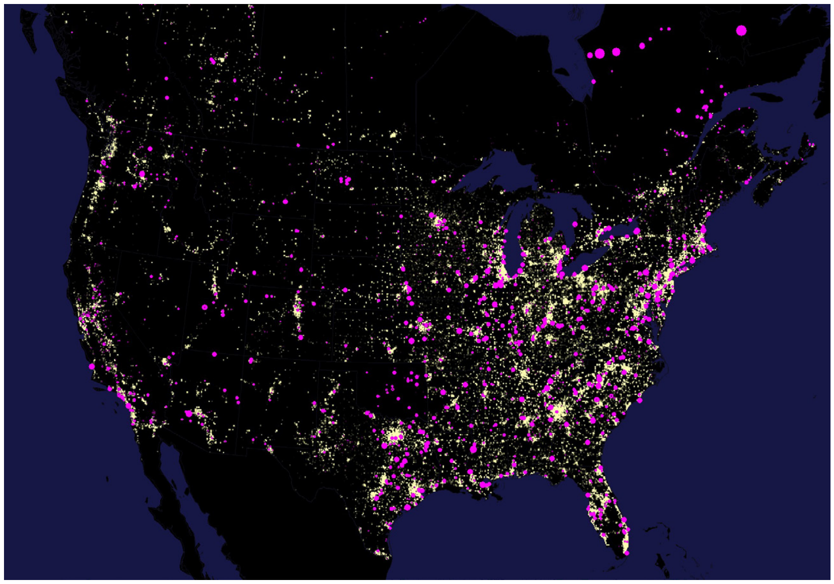

FIGURE 1.3 North American electric load (yellow) and generation (magenta). SOURCE: Courtesy of Thomas J. Overbye, University of Illinois, Urbana-Champaign.

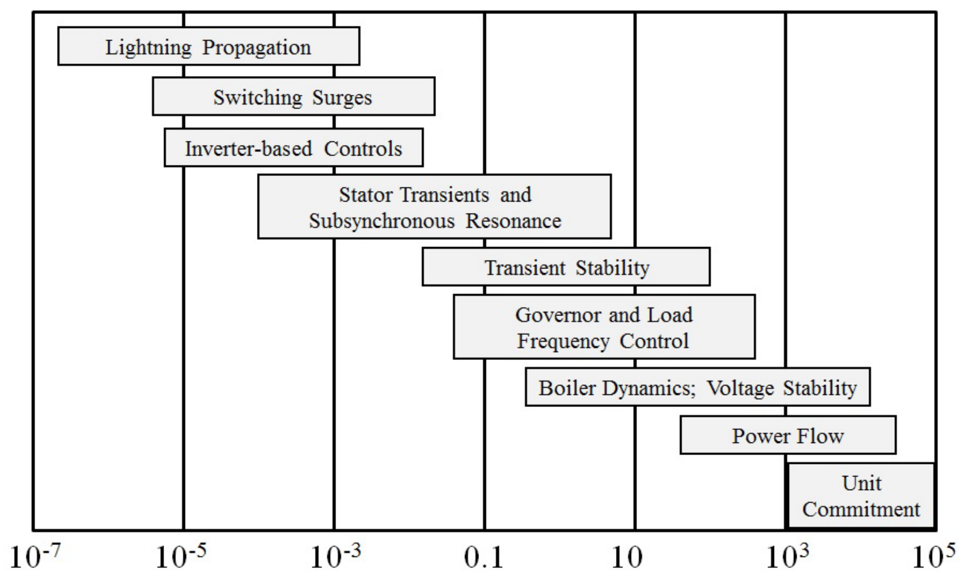

FIGURE 1.4 Power system time scales. SOURCE: Modified from Pai et al. (2006), copyright 2006 Springer-Verlag, Berlin, Heidelberg. Courtesy of Thomas J. Overbye, University of Illinois, Urbana-Champaign.

Basic Circuits—Quasi-Steady-State Time Frame

A good place to start the development of power system models is in what is known as the power flow time frame, or quasi-steady state. This is the time frame that would be perceived if one were to go into a utility control center in the common situation when there are no disturbances on the grid. Being an ac system, the voltages and currents would actually be varying at close to 60 Hz. But the displayed average power consumed by the load would not show this variation. Rather, it would be slowly changing as it goes through its broader daily, weekly, and seasonal variation. Figure 1.5 shows an example of the weekly variation in the total aggregate electric load for PJM (a regional transmission organization in the Eastern Interconnection) during the summer, whereas Figure 1.6 shows an example of the same variation in winter. Likewise, the average generation dispatch would be slowly changing to match the variation in the electric load. So, even though the load might change by close to 100 percent in a single day, the change is slow—to the casual observer the grid would appear to be in near steady state.

To develop the models consider a sinusoidal voltage $v(t)$ or current at a constant frequency, f (say, 60 Hz), so that

$$v(t) = \mathrm{V}_{\max} \cos(\omega t + \theta_V) \tag{1}$$

where $\mathrm{V}_{\max}$ is the maximum voltage level, $\omega = 2\pi \mathrm{f}$, and θ_V is a phase angle offset. The root mean square (rms) for this constant frequency sinusoidal is

$$\mathrm{V} = \mathrm{V}_{\max} / \sqrt{2} \tag{2}$$

If one were to model a network of voltage sources, current sources, resistors, inductors, and capacitors in which all the voltage and current sources were sinusoidal with the same frequency, then all the voltages and currents in the system would be sinusoidal at this frequency. The steady-state response of this uniform frequency network could then be modeled using phasor analysis in which

(1) Each voltage and current is represented by a complex phasor value with the magnitude equal to its rms value and the angle equal to its phase angle.
(2) Each resistance R is represented by an impedance $Z_{\mathrm{R}} = \mathrm{R}$.
(3) Each inductance L is represented by an impedance $Z_{\mathrm{L}} = j\omega\mathrm{L}$.[6]
(4) Each capacitance C is represented by an impedance $Z_{\mathrm{C}} = 1/(j\omega\mathrm{C})$.
(5) The relationship between the phasor voltage and current in a device with impedance Z is given by Ohm's law, $V = ZI$; admittance is defined as the inverse of impedance, so $Y = 1/Z$.

The instantaneous power consumed in a device with a sinusoidal voltage $v(t)$ across the device and sinusoidal current $i(t)$ into the device is

$$p(t) = v(t)i(t) = \mathrm{V}_{\max} \cos(\omega t + \theta_V)\mathrm{I}_{\max} \cos(\omega t + \theta_I) = 2\mathrm{VI} \cos(\omega t + \theta_V) \cos(\omega t + \theta_I) \tag{3}$$

which can be rewritten by applying trigonometric identities as a nonzero average power and a component with double the original frequency:

$$p(t) = \mathrm{VI}[\cos(\theta_V - \theta_I) + \cos(2\omega t + \theta_V + \theta_I)] = \mathrm{P} + \mathrm{VI} \cos(2\omega t + \theta_V + \theta_I) \tag{4}$$

in which V is the rms voltage, I the rms current, and P the average power over a period. Since the average value of the second (sinusoidal) component over a period is zero, for the quasi-steady-state time frame only the average power is of interest. The complex power can be defined as

$$S = VI^{*} = \mathrm{P} + j\mathrm{Q} \tag{5}$$

where the magnitude of S is known as the apparent power, P as the real power, and Q as the reactive power. Real power is usually expressed in megawatts (MW), reactive power in megavars (Mvar), and apparent power in megavoltamperes (MVA). The physical significance of the reactive power is difficult to describe. Reactive power is defined mathematically in (5). Its physical significance is difficult to describe, but roughly speaking, reactive

[6] Where j is defined as the imaginary unit using electrical engineering notation to avoid confusion with the symbol i used for current.

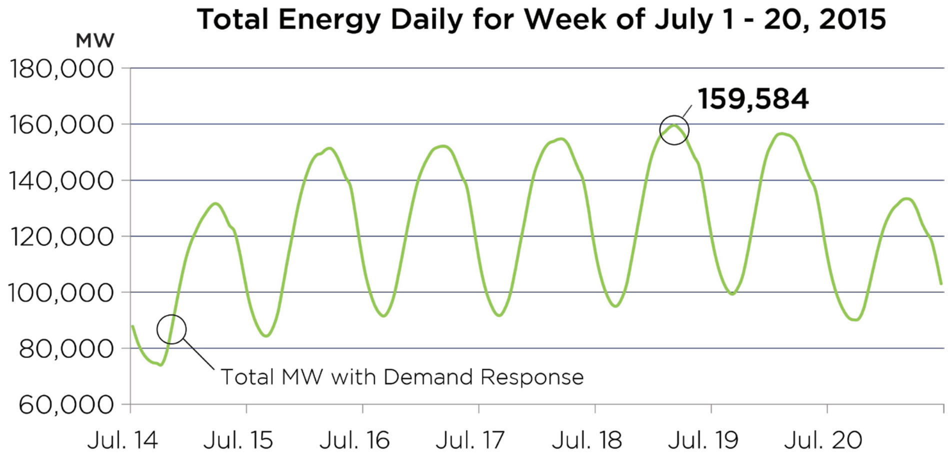

FIGURE 1.5 Weekly aggregate electric load variation for PJM, July 2015. SOURCE: Courtesy of PJM Interconnection.

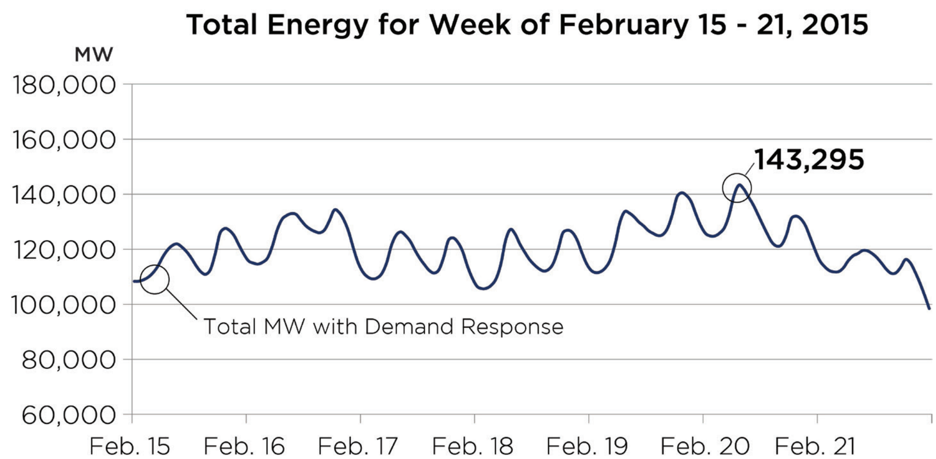

FIGURE 1.6 Weekly aggregate electric load variation for PJM, February 2015. SOURCE: Courtesy of PJM Interconnection.

power represents energy stored for part of an electrical cycle in the magnetic field and released later in the cycle. It is required in order to make many electrical devices, such as the induction motors used in air conditioners and refrigerators, function correctly. The concept of reactive power is quite useful for power system analysis and it is treated in a manner analogous to the real power. It is easy to show that resistors always consume real power, inductors always consume reactive power, and capacitors always generate reactive power.

Three-Phase Power Systems and Per-Phase Analysis

High-voltage power systems are almost always three phase. In a three-phase system there are three conductors instead of the two conductors found in dc circuits or single-phase circuits. A three-phase system is considered balanced if the voltages and currents (respectively) have equal magnitude but are shifted in phase from each other by 120°. The phases are usually labeled A, B, and C. Two key advantages of balanced three-phase systems (compared to single-phase) are (1) for the same amount of wire twice the power can be transferred and (2) three-phase electric devices such as generators and motors are more efficient and hence more economical than single-phase devices with the same power rating.

Three-phase systems can either be Y-connected (wye-connected) or Δ-connected (delta-connected). Figure 1.7 shows an example of a Y-connected voltage source on the left supplying a Y-connected load on the right; in a balanced three-phase system, the neutral current, I_n, would be zero, so this conductor could be omitted. In a balanced three-phase system the line-to-line voltages (e.g., E_a-E_b in the figure) are square root of 3 greater than the line-to-neutral voltages (e.g., E_{an}). Since nominal transmission line voltages are expressed in line-to-line values, a 345 kV transmission line would have line-to-neutral values of 200 kV. Figure 1.8 shows a Y-connected voltage source and a Δ-connected load. Both wye and delta connections are commonly used in the power grid.

The actual power grid is never perfectly balanced. Most generators and some of the load are three-phase systems and can be fairly well represented using a balanced three-phase model. While most of the distribution system is three-phase, some of it is single phase, including essentially all of the residential load. While distribution system designers try to balance the number of houses on each phase, the results are never perfect since individual household electricity consumption varies. In addition, while essentially all transmission lines are three phase, there is often some phase imbalance since the inductance and capacitance between the phases are not identical. Still, the amount of phase imbalance in the high-voltage grid is usually less than 5 percent, so a balanced three-phase model is a commonly used approximation.

In order to model the interconnected power network, appropriate models need to be developed for the transmission lines, transformers, generators, and loads. The analysis of a balanced three-phase system can be greatly simplified by using a technique known as per-phase analysis, in which the system is treated as an equivalent single-phase system.

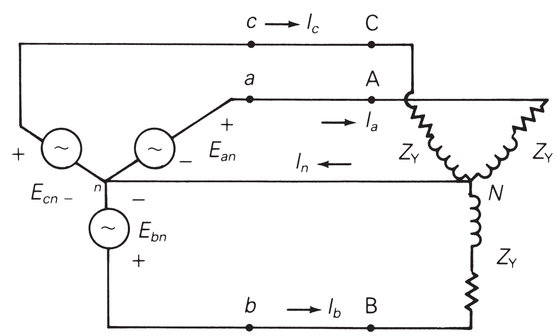

FIGURE 1.7 Wye-connected voltage source and load. SOURCE: Glover et al. (2012), *Power System Analysis and Design*, 5E, © 2012 Cengage Learning, a part of Cengage Learning, Inc., reproduced with permission, www.cengage.com/permissions.

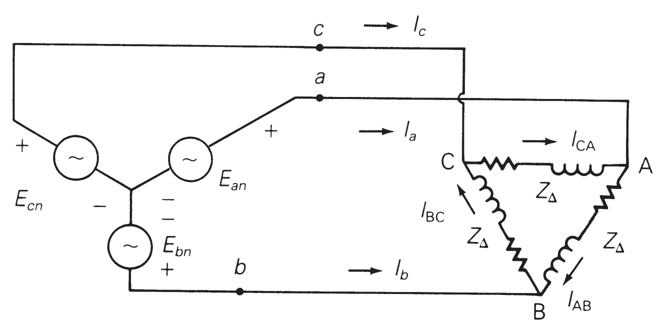

FIGURE 1.8 Wye-connected voltage source and delta-connected load. SOURCE: Glover et al. (2012), *Power System Analysis and Design*, 5E, © 2012 Cengage Learning, a part of Cengage Learning, Inc., reproduced with permission, www.cengage.com/permissions.

In the steady-state time frame a reasonable per-phase model for a transmission line is what is known as the π-equivalent circuit, consisting of a series impedance Z' between two shunt admittances $Y'/2$ (Figure 1.9). The maximum amount of power that can be transferred through a transmission line, sometimes due to thermal constraints, is often represented as the maximum power in megavolt amperes, or MVA limit.

Likewise, a reasonable steady-state transformers model consists of a series impedance and shunt admittance, except now in series with an ideal transformer model (shown in Figure 1.10). In an ideal transformer model, the ratio of the voltages, E_1/E_2, is identical to the turns ratio of the windings, $a_t = N_1/N_2$, and the ratio of the current into the E_1 side versus the current out of the E_2 side is the inverse of the turns ratio. The maximum amount of power that can be transferred through a transformer is also often represented as an MVA limit.

In order to easily analyze networks with transformers it is helpful to introduce what is known as per-unit (PU) analysis, in which the system values are normalized using base values that depend on a systemwide power base and voltage bases that differ by the turns ratios of the ideal transformers. PU analysis can be used with either single-phase systems or, as presented here, three-phase systems. For three-phase PU, first select a single three-phase base power for the entire system, $S_{b,3\varphi}$; 100 MVA is typical. Then select line-to-line voltage bases that differ by the ideal transformer turns ratios, $V_{b,LL}$; these values are typically the nominal transmission voltages (e.g., 500, 345, 138 kV). Current, impedance, and admittance bases can then be defined as

$$I_b = \frac{S_{b,3f}}{\sqrt{3}V_{b,LL}}, \ Z_b = \frac{V_{b,LL}^2}{S_{b,3f}}, \ Y_b = \frac{1}{Z_b} \tag{6}$$

All network complex powers, voltages, currents, and impedances are converted to PU by normalizing by their corresponding base values, which are quantities in the nominal steady-state operating conditions. PU values are still complex numbers but are dimensionless. When using PU, the ideal transformers are eliminated from the

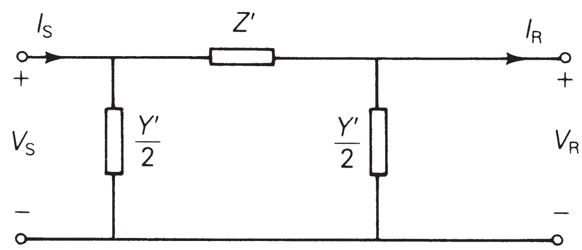

FIGURE 1.9 Transmission line π-equivalent circuit. SOURCE: Glover et al. (2012), *Power System Analysis and Design*, 5E, © 2012 Cengage Learning, a part of Cengage Learning, Inc., reproduced with permission, www.cengage.com/permissions.

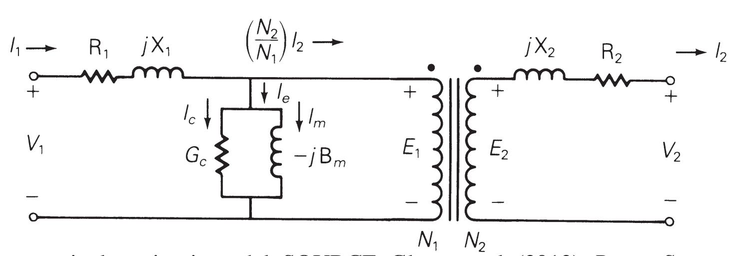

FIGURE 1.10 Transformer equivalent circuit model. SOURCE: Glover et al. (2012), *Power System Analysis and Design*, 5E, © 2012 Cengage Learning, a part of Cengage Learning, Inc., reproduced with permission, www.cengage.com/permissions.

transformer models. This results in a model of the network consisting of just PU impedances and admittances, greatly simplifying the network analysis. By using a three-phase PU base, a balanced three-phase system can be solved as though it were a single-phase system. With proper accounting of the 30° phase shift in transformer voltages for the wye-delta connection, the analysis is the same whether a device is connected as wye or as delta.

To study an interconnected system the relationship between the PU phasor voltages and bus (node), current injections can be obtained by applying Kirchhoff's current law (KCL) at each bus in the system. That is, the equality constraints are obtained by recognizing that the net current being injected into each bus must be equal to the current flowing out of the bus into the rest of the network. Using matrix notation for a network with N buses gives

$$\boldsymbol{I} = \boldsymbol{Y}\,\boldsymbol{V} \tag{7}$$

where $\boldsymbol{Y}$ is the $N \times N$ bus admittance matrix, $\boldsymbol{I}$ is an N-dimensional column vector of the net phasor current injections at each bus, and $\boldsymbol{V}$ is the N-dimensional column vector of the bus voltages. In a large network, $\boldsymbol{Y}$ will be quite sparse since there is only a nonzero off-diagonal entry $\boldsymbol{Y}_{kn}$ if there is a direct connection between buses k and n.

If the generator outputs could be represented as complex current injections and the loads as shunt admittances, then this equation could be used to determine $\boldsymbol{V}$, and by using $\boldsymbol{V}$ with the branch models (with "branch" used generically to refer to the transmission lines and transformers), all the system complex powers could be determined. Unfortunately, the generator outputs cannot be represented as current injections, and the loads are not well modeled as shunt admittances. Rather, to determine $\boldsymbol{V}$ the nonlinear power flow equations need to be formulated and solved.

ILLUSTRATIVE TYPES OF ANALYSIS NEEDED FOR THE GRID

Power Flow—Steady-State Analysis

The power flow or load flow (the two terms have been used interchangeably since at least the 1960s) is the most widely used power system analysis technique either as a stand-alone application or embedded in other applications. The goal of the power flow is to determine the quasi-steady-state $\boldsymbol{V}$ vector, given a specified set of generation and load values. To develop the power flow equations, it is necessary to first present time-scale-appropriate generator and load models.

On the power flow time scale, generators are usually most appropriately modeled as a constant real power injection (P) into the system at a specified per unit (PU) bus voltage magnitude (V). Hence the generator is assumed to be modifying its reactive power output (Q) to keep its terminal (bus) voltage magnitude constant. This is known as a PV bus. Loads are often represented as constant negative real power (P) and reactive power (Q) injections into the system and are known as PQ buses. However, because in the quasi-steady-state time frame the total real power generation must exactly match the total real-power-load bus losses, the outputs of all the generators cannot be independently specified. Rather, at least one generator is designated as the slack (or swing) bus, in which the voltage magnitude and angle at the generator's bus is specified, and the power flow algorithm determines the generator's real and reactive power output.

Power flow equation derivation starts with applying KCL at each bus so that the net current injection into the bus must equal the current going into the network. And since the complex power is the voltage times the conjugate of the current, the net complex power injection into the bus must equal the complex current into the network. For bus k, the relationships are

$$\begin{aligned} I_k &= \sum_{j=1}^{n} Y_{kj} V_j \\ \mathrm{P}_k + j\mathrm{Q}_k &= V_k \left[\sum_{j=1}^{n} Y_{kj} V_j \right]^* \end{aligned} \tag{8}$$

where P_k and Q_k are the specified real and reactive power injections at bus k. Expressing the complex numbers with the following notation,

$$
\begin{aligned}
V_k &= \mathrm{V}_k \angle \theta_k \\
Y_{kj} &= \mathrm{G}_{kj} + j\mathrm{B}_{kj}
\end{aligned}
\tag{9}
$$

these complex equations are typically written as the real-valued power balance equations

$$
\begin{aligned}
\mathrm{P}_k &= \mathrm{V}_k \sum_{j=1}^{n} \mathrm{V}_n \left[\mathrm{G}_{kj} \cos\left(\theta_k - \theta_j\right) + \mathrm{B}_{kj} \sin\left(\theta_k - \theta_j\right) \right] \\
\mathrm{Q}_k &= \mathrm{V}_k \sum_{j=1}^{n} \mathrm{V}_n \left[\mathrm{G}_{kj} \cos\left(\theta_k - \theta_j\right) - \mathrm{B}_{kj} \sin\left(\theta_k - \theta_j\right) \right]
\end{aligned}
\tag{10}
$$

Hence the power flow problem is the solution of a set of $2n$ nonlinear algebraic equations given by (10). For PV buses the reactive power balance equations are not included since the voltage magnitude at these buses is specified; the reactive power outputs of the PV generators are dependent variables. For the slack bus, neither equation is included since both the real and reactive power output of the generators are dependent variables.

Power flow models can come in all different sizes, from just a few buses for academic systems, to a representation of an entire interconnect. When modeling a large system it is certainly not possible nor is it needed to represent each individual load. Rather, since the distribution system is usually radial it is often sufficient in power flow studies to represent all the devices on a distribution feeder as a single, aggregate load. In addition, equivalent models can be developed that further reduce the number of buses that need to be represented. Currently the Eastern Interconnection, with a total load of about 650 GW, is modeled with about 65,000 buses, while 20,000 buses are currently used for the Western Interconnection, with a total load of about 170 GW.

Figure 1.11 shows an example power flow solution for a small seven-bus system. The results are shown on a "oneline diagram" (oneline) in which the actual three-phase transmission lines are represented by a single line. As is common for engineering studies, voltages are reported in PU, so assuming a 138-kV voltage base, the actual line-to-line voltage magnitude shown in the figure in yellow at bus 1 as 1.05 PU would be 1.05 × 138 kV = 144.9 kV. The oneline shows the real and reactive power outputs for all the generators (shown as black circles), the aggregate

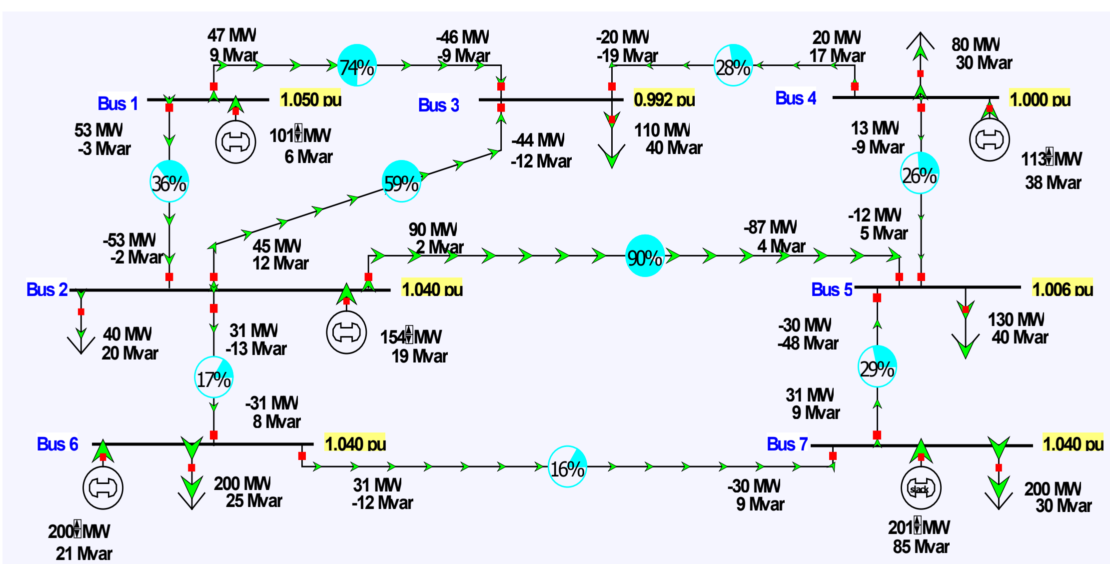

FIGURE 1.11 Example of a seven-bus power flow solution. SOURCE: Courtesy of Thomas J. Overbye, University of Illinois, Urbana-Champaign.

bus loads (shown as black arrows), and the transmission lines. Note that both the real and reactive powers at each bus sum to zero, with a sign convention that power into each transmission line at the bus is assumed positive. The green arrows show the flow direction of the real power. Because of line resistance the amount of real power out of a line is always less than the power into it. With reactive power this is not always the case since the transmission line model includes capacitance terms (which create reactive power). The PU bus voltage magnitudes are shown with the yellow fields. Buses 1, 2, 4, and 6 are modeled as PVs with a fixed voltage magnitude, and bus 7 is the system slack. The pie charts show the percentage loading for each of the transmission lines, with the limit for each line specified in terms of either maximum MVA or maximum amps. Oftentimes these limits are due to thermal considerations, recognizing that as the lines' conductors heat up they expand, resulting in increased sag for overhead conductors.

The power flow is commonly used to determine how modifications to the generation, load, or system topology would affect the flows throughout the system. Figure 1.12 shows the previous example, except now with the transmission line between bus 2 and bus 3 out. If this outage had occurred on an actual system there surely would have been transient changes to the system (as per Figure 1.4), including switching surges, and transient stability oscillations. But assuming the system remained stable, within several seconds it would have settled back to the quasi-steady-state power flow solution shown in Figure 1.12. Except for the topology change, all the power flow inputs (i.e., the load real and reactive values, and the generator real power and voltage setpoint values) remained constant. The only changes were to the power-flow-dependent variables, including the PQ bus voltage magnitudes, the PV generator reactive power outputs, and the slack bus real and reactive outputs. A single contingency, such as opening the line between buses 2 and 3, also changes the flows throughout the system, albeit with the largest changes usually closest to the contingency.

This example also illustrates that the transmission line flows are dependent variables—they cannot be directly controlled. In general, they can only be indirectly controlled, such as by changing the generator real power outputs (exceptions are phase-shifting transformers and HVDC transmission lines that do allow direct flow control). This is illustrated in Figure 1.13, where the transmission line overloads from the previous contingency are removed by reducing the real power output of generator 6 from 200 to 101 MW. Note a corresponding increase in the bus 7 (slack) generation, from 203 to 300 MW; the net change in the two generators does not sum exactly to zero

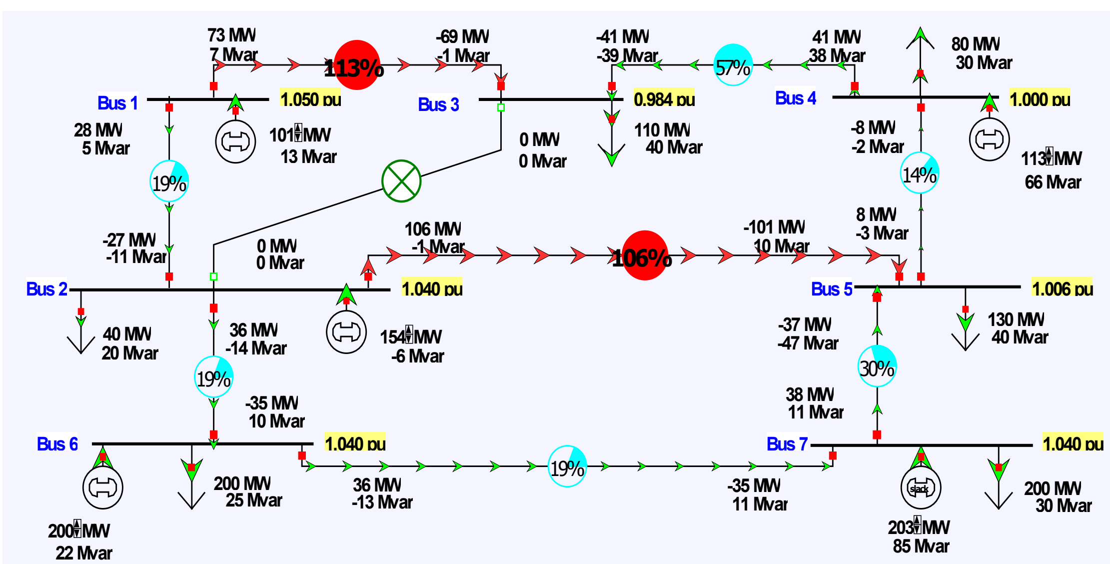

FIGURE 1.12 Seven-bus power flow solution with the transmission line between buses 2 and 3 out. SOURCE: Courtesy of Thomas J. Overbye, University of Illinois, Urbana-Champaign.

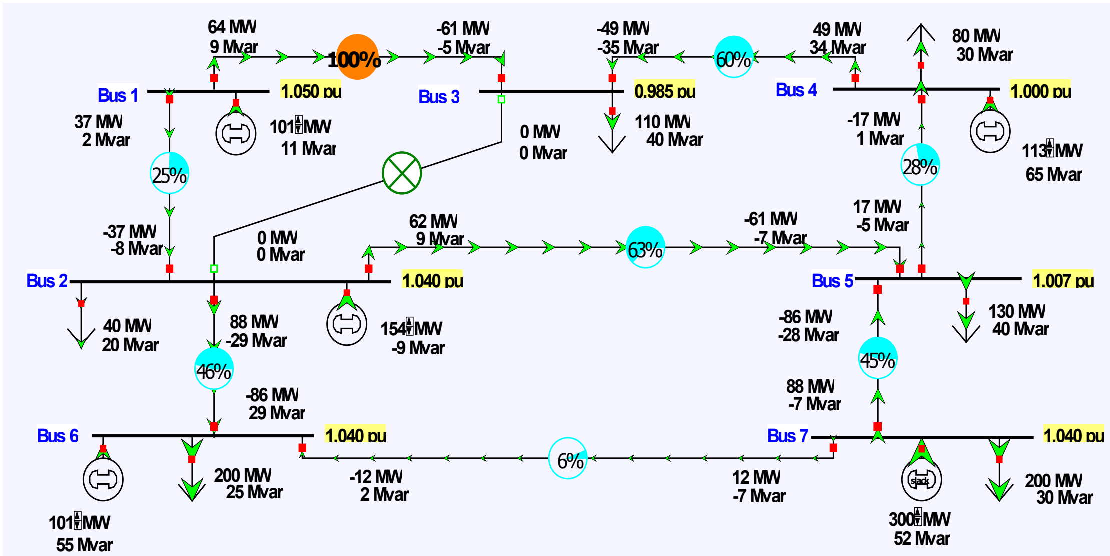

FIGURE 1.13 Generation re-dispatch to remove the transmission line overloads. SOURCE: Courtesy of Thomas J. Overbye, University of Illinois, Urbana-Champaign.

because of a slight change in the real power losses. In addition to branch limits, reliable power system operation also requires that the bus voltage magnitudes be within a reasonable range, usually between about 0.95 and 1.05 PU.

When modeling large-scale power systems, the basic power flow algorithm presented here is augmented to model the response of various continuous and discrete power system controllers. While the details are beyond the scope of this brief introduction, examples include load-tap-changing (LTC) transformers, phase-shifting transformers, switched capacitor banks, automatic generation control, HVDC transmission lines, and more advanced generator voltage control. Hence the power flow is solving a set of nonlinear algebraic equations,

$$\mathbf{0} = \mathbf{g}(\mathbf{y}, \mathbf{u}) \tag{11}$$

where $\mathbf{g}$ is a vector of algebraic constraints including the real and reactive power balance equations, $\mathbf{y}$ is the solution variable vector such as the PQ bus voltage magnitudes and angles, and $\mathbf{u}$ is the input parameter vector such as the load real and reactive power values. Both $\mathbf{y}$ and $\mathbf{u}$ might contain a mixture of continuous and discrete values.

One common approach to avoid solving the nonlinear equations of (11), used particularly with the market analysis discussed in Chapter 2, is to assume the approximations shown in (12). First, since the resistances of the transmission lines are often much less than their reactances, the conductance terms are assumed to be zero. Second, since the voltage magnitudes are usually close to 1.0 PU, they are assumed to be just that. Third, given that the angle differences across the lines are small, the cosine terms are assumed to be unity and the sine terms are approximated as the angle differences. Last, the reactive power constraints are ignored. This reduces the power flow to a set of linear equations,

$$\mathbf{P} = \mathbf{B}\,\boldsymbol{\theta} \tag{12}$$

with the inputs $\mathbf{P}$ and $\mathbf{B}$ used to solve for $\boldsymbol{\theta}$. This approximation is known as the dc power flow, with the nonlinear power flow often referred to as the ac power flow. Note that both are still ultimately providing solutions to an ac circuit, with the dc power flow just a linear approximation. The validity of the approximations is quite system specific. To illustrate, Figure 1.14 contains the dc power flow solution for the seven-bus system whose ac power flow solution is given in Figure 1.12.

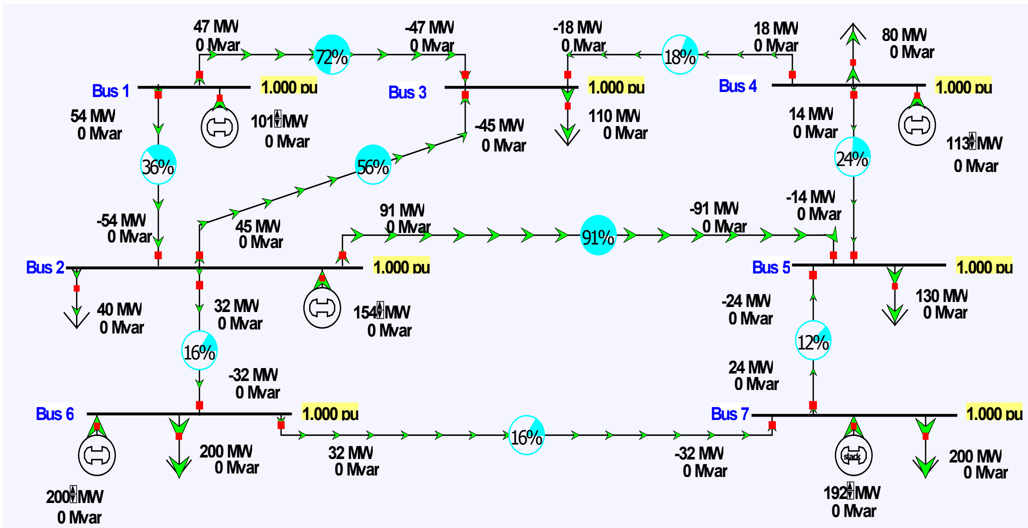

FIGURE 1.14 Seven-bus case from Figure 1.11 solved using the dc power flow. SOURCE: Courtesy of Thomas J. Overbye, University of Illinois, Urbana-Champaign.

Interconnected Power System Steady-State Operations

The basics of steady-state operations can be fairly well thought of as a slowly changing power flow solution. As the load slowly varies, the values of various controls are changed, either automatically or manually by the power system operator, and the line flows respond to them. The most crucial control is the modification of the real power outputs of the generators to match changes in the system load, a process known as automatic generation control (AGC).

While an interconnected grid is just one big electric circuit, many of them, including the North American Eastern and Western Interconnections, were once divided into “groups”; at first, each group corresponded to an electric utility. These groups are now known as load-balancing areas (or just “areas”). The transmission lines that join two areas are known as tie lines, and the algebraic sum of the real power flow on the tie lines for an area is known as its net interchange, with the usual sign convention that power flow out of an area is defined as positive.

The area control error (ACE) for area k is then defined as

$$\mathrm{ACE}_k = \mathrm{NI}_{A,k} - \mathrm{NI}_{S,k} - 10\beta_k\,(\mathrm{F}_A - \mathrm{F}_S) - \mathrm{I}_{ME,k} \tag{13}$$

where $\mathrm{NI}_{A,k}$ is the actual net interchange in MW, $\mathrm{NI}_{S,k}$ is the scheduled net interchange in MW, β_k is an area-specific bias term in MW/0.1 Hz (with a negative sign), F_A is the actual system frequency in hertz, F_S is the scheduled system frequency in hertz, and $\mathrm{I}_{ME,k}$ is the interchange metering error term that is usually small or zero (NERC, 2011). The scheduled system frequency is usually 60 Hz, but it can be either 59.98 or 60.02 Hz for time error correction. The ACE for each area and the system frequency are the most important numbers associated with the system’s operation; ACE is kept close to zero by using AGC to adjust the generation to match the changing load.

With this approach it is possible to easily implement power transactions between different areas. These are known as bilateral transactions, since they involve two players. The scheduled net interchange for each area is just equal to the sum of its transactions. Modifying the scheduled net interchange causes a change in the ACE, causing AGC to adjust the outputs of generators in the area. This is demonstrated in Figure 1.15, in which the original system is now subdivided into three areas: top (containing buses 1 to 5), left (containing bus 6), and right (containing bus 7). The system is modeled with a single 100-MW transaction going from right (the seller) to left (the buyer). The system is modeled with the ACE for each area equal to zero, so the net flow across each area’s

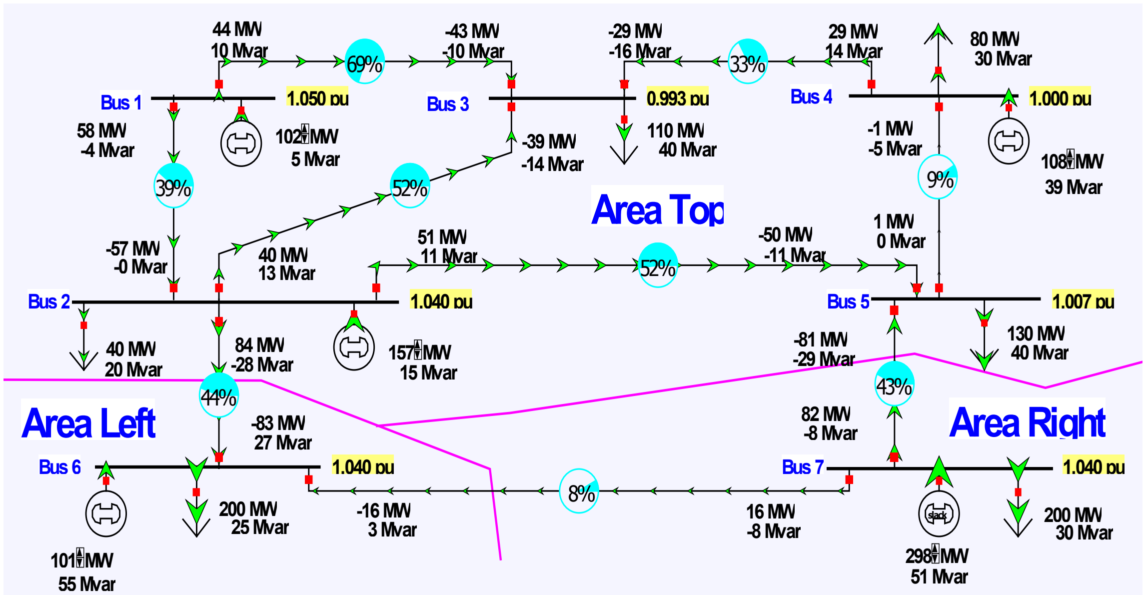

FIGURE 1.15 Original system with a 100-MW transaction from right to left. SOURCE: Courtesy of Thomas J. Overbye, University of Illinois, Urbana-Champaign.

tie lines is equal to its schedule values (with the power flow assumption that the actual and scheduled frequencies are 60 Hz). When the flows in Figure 1.15 are compared to those in the original Figure 1.11, it is seen that a single transaction can impact the flows across an interconnect.

Power transactions between different players (e.g., electric utilities, independent generators) in an interconnection can take from minutes to decades. In a large system such as the Eastern Interconnection, thousands of transactions can be taking place simultaneously, with many of them involving transaction distances of hundreds of miles, each potentially impacting the flows on a large number of transmission lines. This impact is known as loop flow, in that power transactions do not flow along a particular "contract path" but rather can loop through the entire grid.

With a power flow solution, the incremental impact of each transaction can be calculated from sensitivity analysis, with the sensitivities of how much a single transaction impacts the flows on each line known as power transfer distribution factors (PTDFs). The PTDFs for the Figure 1.15 transaction are visualized in Figure 1.16, with the pie charts and arrows now showing the percentage of the right to left transaction that flows on each line, with a total of 100 percent leaving the right and arriving at the left.

Because the electric grid is regularly subject to faults and other disturbances (e.g., lightning hitting a transmission line or a generator failing), a crucial aspect of power system operations is the need to continue operating with no limit violations even when subject to such contingencies. Examples of limits include keeping the transmission line and transformer flows below a specified MVA value and keeping the bus voltage magnitudes within a PU range (e.g., between 0.95 and 1.05 PU). The standard operating paradigm is to be at least N – 1 reliable, meaning that if any single credible contingency were to occur there would be no limit violations. N – 1 reliability is assessed using contingency analysis (CA), which in its simplest form consists of running potentially thousands of power flow solutions, each considering a different contingency. Online CA is commonly run in electric control centers on about a 5-minute interval. Since each contingency is independent, CA may be easily parallelized.

Another common online analysis tool is optimal power flow (OPF). The purpose of OPF is to minimize some scalar value, such as total operating cost, while satisfying various equality and inequality constraints:

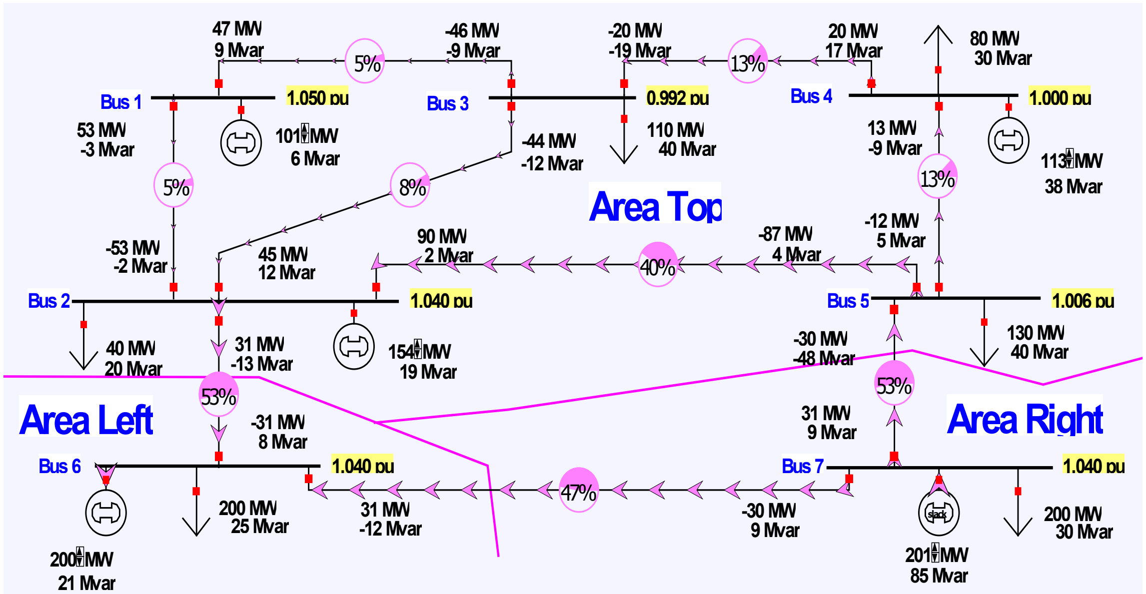

FIGURE 1.16 Original system PTDFs for a transaction from area right to area left. SOURCE: Courtesy of Thomas J. Overbye, University of Illinois, Urbana-Champaign.

$$
\begin{aligned}
\text{Min F} \quad & (\mathbf{y},\mathbf{u}) \\
\text{such that} \quad & \mathbf{0} = \mathbf{g}(\mathbf{y},\mathbf{u}) \\
& \mathbf{h}_{\min} \leq \mathbf{h}(\mathbf{y},\mathbf{u}) \leq \mathbf{h}_{\max} \\
& \mathbf{u}_{\min} \leq \mathbf{u} \leq \mathbf{u}_{\max}
\end{aligned} \tag{14}
$$

The key equality constraints are the power balance equations from the power flow shown in equations (11), while the key inequality constraints are the need to operate with the branch flows, bus voltage magnitudes, and generator reactive powers within their limits. The system controls may be either continuous (e.g., generator real power outputs) or discrete (e.g., transformer tap positions, switched shunt status). Figure 1.17 shows an example of an OPF solution for the seven-bus case. The one line has been modified to show the incremental cost of enforcing the real power constraint at each bus (in $/MWh), a value known as the locational marginal cost (LMP); LMPs are widely used in the operation of electric power markets (discussed in the next chapter). Also, as is common in actual power markets, a color contour is used to visualize the variation in the LMPs. In this example the system is segmented because of the MVA limit on the line between buses 2 and 5. OPF is commonly combined with CA to determine the optimal dispatch taking into account all the contingencies, so the final solution is N – 1 reliable (something that was not done in Figure 1.17). This is known as security-constrained optimal power flow, (SCOPF). As originally formulated, the OPF used the full ac power flow equations as given in (10). This is now often refereed to as the ACOPF. A new, more approximate approach uses the dc power flow equations given in (12). This is often referred to as the DCOPF. Likewise, the SCOPF can be formulated using either the ac power flow equations or the dc equations. Commonly, however, the terms OPF and SCOPF are used generically to refer to either the ac or the dc approach.

In order to run the previous analysis techniques online with an actual grid, it is first necessary to obtain a starting power flow solution that matches as closely as possible the actual grid conditions. This is done in a process known as state estimation (SE), in which a large number of imperfect measurements, such as bus voltage magnitudes and line-flow real and reactive flow values, are used to obtain the solution of equations (11) that best matches the measurements. Electric control centers typically run SE every few minutes. In contrast to power flow

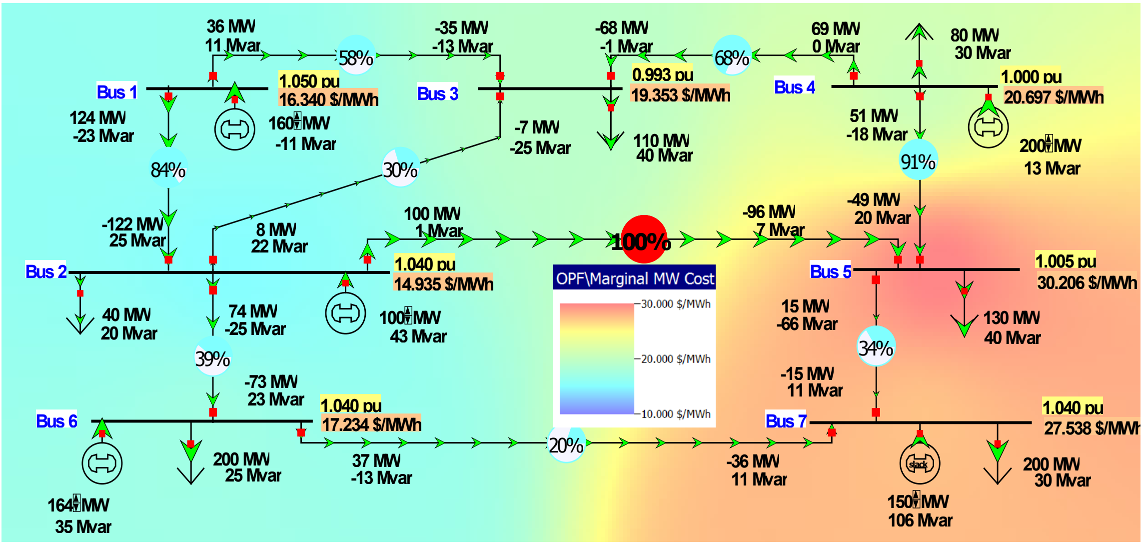

FIGURE 1.17 OPF solution of original seven-bus system. SOURCE: Courtesy of Thomas J. Overbye, University of Illinois, Urbana-Champaign.

in which the number of variables matches the number of equations, SE is an overdetermined problem. A discussion of the currently used algorithms for solving power flow—CA, OPF, and SE—is contained in Chapter 4.

Day-Ahead Planning and Unit Commitment

In order to operate in the steady state, a power system must have sufficient generation available to at least match the total load plus losses. Furthermore, to satisfy the N – 1 reliability requirement, there must also be sufficient generation reserves so that even if the largest generator in the system were unexpectedly lost, total available generation would still be greater than the load plus losses. However, because the power system load is varying, with strong daily, weekly, and seasonal cycles, except under the highest load conditions there is usually much more generation capacity potentially available than required to meet the load. To save money, unneeded generators are turned off.

The process of determining which generators to turn on is known as unit commitment. How quickly generators can be turned on depends on their technology. Some, such as solar PV and wind, would be used provided the sun is shining or the wind blowing, and these are usually operated at their available power output. Hydro and some gas turbines can be available within minutes. Others, such as large coal, combined-cycle, or nuclear plants, can take many hours to start up or shut down and can have large start-up and shutdown costs.

Unit commitment seeks to schedule the generators to minimize the total operating costs over a period of hours to days, using as inputs the forecasted future electric load and the costs associated with operating the generators. As will be considered in Chapter 2, unit commitment constraints are a key reason why there are day-ahead electricity markets. Complications include uncertainly associated with forecasting the electric load, coupled increasingly with uncertainty associated with the availability of renewable electric energy sources such as wind and solar.

The percentage of energy actually provided by a generator relative to the amount it could supply if it were operated continuously at its rated capacity is known as its capacity factor. Capacity factors, which are usually reported monthly or annually, can vary widely, both for individual generators and for different generation technologies. Approximate annual capacity factors are 90 percent for nuclear, 60 percent for coal, 48 percent for natural gas combined cycle, 38 percent for hydro, 33 percent for wind, and 27 percent for solar PV (EIA, 2015). For some technologies, such as wind and solar, there can be substantial variations in monthly capacity factors as well.

One issue associated with day-ahead planning is the need to ensure there is sufficient generation that can change (ramp) its output quickly in order to meet changes in the net load. As illustrated in Figure 1.5, ramping of generation to meet the changing load has long been a part of power system operations. However, with the growth in solar PV generation, ramping is becoming more of an issue as the net load rapidly decreases in the morning as the sun rises and falls in the evening as it sets. This impact of solar PV is illustrated in Figure 1.18, in what is known in the industry as the "duck" curve, because it resembles the aquatic bird.

Longer-Term Power System Planning

Much of the preceding discussion applies both to online operations and planning. However, planning has some unique aspects that deserve special consideration. Planning takes place on time scales ranging from perhaps hours in a control room setting, to more than a decade in the case of high-voltage transmission additions. The germane characteristic of the planning process is uncertainty. While the future is always uncertain, recent changes in the grid have made it even more so. Planning was simpler in the days when load growth was fairly predictable and vertically integrated utilities owned and operated their own generation, transmission, and distribution. Transmission and power plant additions could be coordinated with generation additions since both were controlled by the same utility.

As a result of the open transmission access that occurred in the 1990s, there needed to be a functional separation of transmission and generation, although there are still some vertically integrated utilities. Rather than being able to unilaterally plan new generation, a generation queue process is required in which requests for generation interconnections needed to be handled in a nondiscriminatory fashion. The large percentage of generation in the queue that will never actually get built adds uncertainty, since in order to determine the incremental impact of each new generator, an existing generation portfolio needs to be assumed. They cannot be considered independently. There

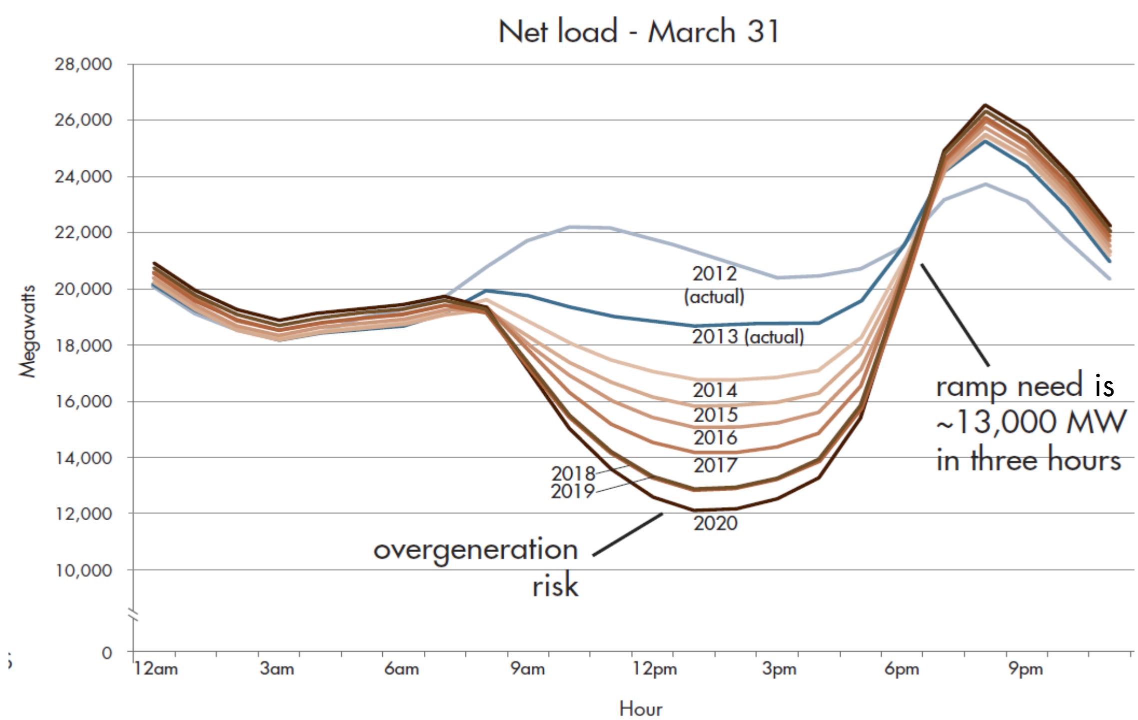

FIGURE 1.18 "Duck" curve. SOURCE: Courtesy of California Independent System Operator (California ISO, 2013). Licensed with permission from the California ISO. Any statements, conclusions, summaries or other commentaries expressed herein do not reflect the opinions or endorsement of the California ISO.

is also the question of who bears the risk associated with the construction of new generation. More recently, additional uncertainty is the growth in renewable generation such as wind and solar PV and in demand-responsive load.

Power System Stability

Switching attention to the faster time scales, transient stability is concerned with power system behavior on time frames ranging from about 0.01 sec to perhaps a few dozen seconds. In contrast to power flow, which seeks to determine a quasi-steady-state equilibrium point (EP), transient stability seeks to determine whether following a system contingency, such as a short circuit or loss of a generator, the system will return to an equilibrium point that may, however, often be different from the original EP. The general form of the problem is as a set of differential algebraic equations (DAEs):

$$\begin{aligned} \dot{\mathbf{x}} &= \mathbf{f}(\mathbf{x},\mathbf{y},\mathbf{u}) \\ \mathbf{0} &= \mathbf{g}(\mathbf{x},\mathbf{y},\mathbf{u}) \end{aligned} \tag{15}$$

in which $\mathbf{x}$ is a vector of state variables, $\mathbf{u}$ is the vector of system inputs, and $\mathbf{y}$ is the vector of algebraic variables, with many entries similar to the power flow variables such as the bus voltage magnitudes and angles. The starting point for a transient stability study is usually a power flow, and the initial values for $\mathbf{x}$ are determined by solving $\mathbf{f}(\mathbf{x},\mathbf{y},\mathbf{u}) = \mathbf{0}$.

Many of the differential equations contained in $\mathbf{f}$ are associated with modeling the behavior of the synchronous machines during this time frame. The most important of the synchronous generator differential equations is what is known as the generator swing equation, which can be expressed for generator k as two first-order differential equations,

$$\begin{aligned} \frac{d\delta_k}{dt} &= \Delta\omega_k \\ \frac{d\Delta\omega_k}{dt} &= \frac{1}{\mathrm{M}_k}\left[T_{mech,k}(\mathbf{x},\mathbf{y},\mathbf{u}) - T_{elec,k}(\mathbf{x},\mathbf{y},\mathbf{u}) - \mathrm{D}_k\Delta\omega_k\right] \end{aligned} \tag{16}$$

where δ_k and $\Delta\omega_k$ are state variables (elements of $\mathbf{x}$) that represent the generator's rotor torque angle and the generator's deviation from synchronous speed; $T_{mech,k}$ is the mechanical torque input to the generator; $T_{elec,k}$ is the electrical torque output from the generator; D_k is a damping coefficient; and M_k is a value that depends on the inertia of the electric generator. The generator swing equation is commonly written in terms of mechanical and electric power rather than torque, with the rationale that the machine's speed is usually quite close to synchronous speed.

Commonly generators are represented with additional differential equations for the electric machines, for their exciters (to control the terminal voltage), for their governors (to control the mechanical power input), and for their stabilizers (to reduce system oscillations). A block diagram of these relationships is shown in Figure 1.19, with often a dozen or more differential equations modeled per generator. Load dynamics, such as those of induction motors, can also be included.

For a large system a single transient stability solution might involve the integration of more than 100,000 differential equations with tens of thousands of algebraic constraints using a time step of perhaps ¼ cycle (0.004166 sec for 60 Hz). Traditionally, transient stability solutions involved just a few seconds of simulation looking at "first swing" instability, though now they can run for dozens of seconds, looking at the longer-term behavior of quantities such as frequency and bus voltage magnitudes.

Figure 1.20 shows an example of first swing stability, plotting the generator torque angles for the seven-bus system shown in Figure 1.11, which has been augmented to include generator dynamic models. In this example the contingency is a low impedance fault at 1.0 sec near bus 1 on the transmission line between buses 1 and 2, which is cleared after three cycles (0.05 sec) by opening this line. During the fault the voltage at bus 1 is quite depressed, which greatly reduces the power output from generator 1, causing the generator to accelerate, increasing

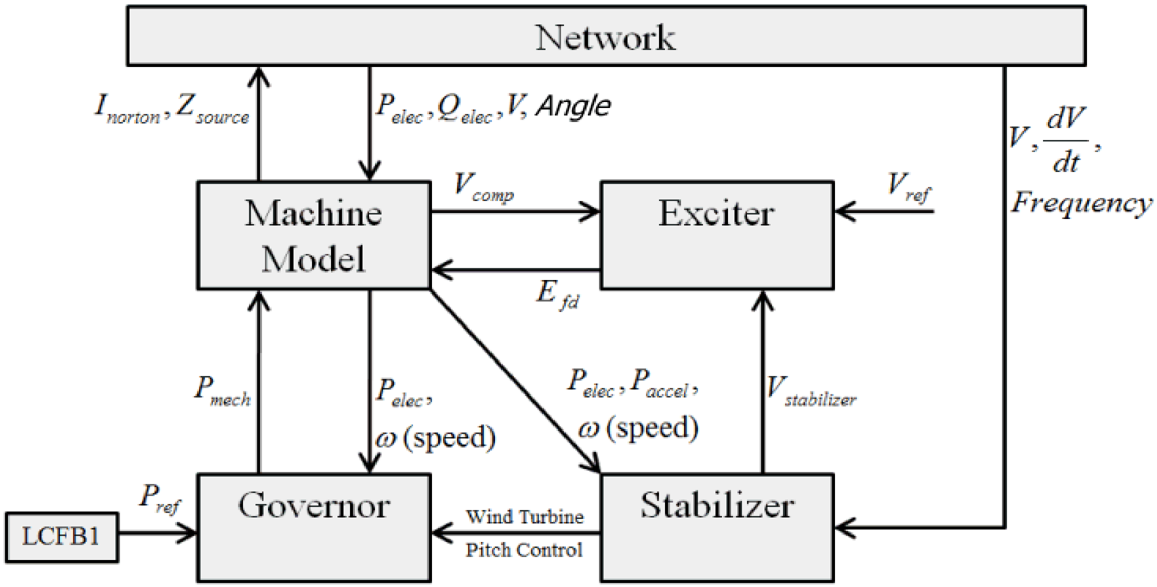

FIGURE 1.19 Transient stability generator model couplings. SOURCE: Courtesy of PowerWorld Corporation.

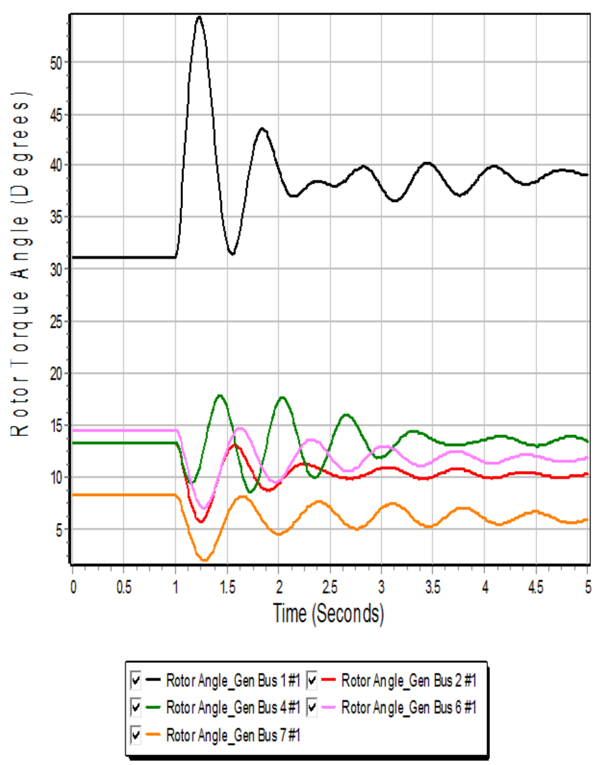

FIGURE 1.20 Seven-bus system generator torque angles for a bus fault. SOURCE: Courtesy of PowerWorld Corporation.

its torque angle with respect to the other generators. When the fault is cleared, the voltage is increased, with the simulation indicating that the system returns to a new quasi-steady-state equilibrium point.

In addition to generator torque angles, quantities of interest during a transient stability study include the generator speeds, the bus voltage magnitudes, and the bus frequencies. As a large case example, Figure 1.21 shows the generator speeds for an 18,000-bus case with a contingency modeling the opening of two large generators.

There has recently been increased interest in power system dynamics on the transient stability time frame. This is partially due to growing concerns about blackouts caused by transient stability issues, but also to greatly increased deployment of synchronized PMUs. By taking advantage of accurate time measurements available thanks to the Global Positioning System, PMUs can determine the power system voltages and current magnitudes and angles at typically 30 times per second. This is in contrast to the existing Supervisory Control and Data Acquisition systems, which return measurements every 4 to 12 seconds. Hence, transient stability time frame dynamics can now easily be viewed in real time at control centers, allowing for greatly improved modeling and analysis capabilities, and there is a growing desire to run transient stability studies in an online environment. Such an application would start from the SE solution (or even one directly observed by the PMUs) and then sequentially solve potentially thousands of contingencies. However, like traditional CA, this transient stability CA would be naturally parallelizable, since the transient stability for each contingency could be considered separately.

Situated between the power flow and transient stability time frames is voltage stability, defined as “the ability of a power system to maintain steady voltages at all buses in the system after being subjected to a disturbance from a given initial operating condition” (Kundur et al., 2004). The term “voltage collapse” is often used when voltage stability is lost, resulting in an uncontrolled decline in system voltages. According to a joint IEEE/CIGRE task force, voltage stability can be classified two ways—by the size of the disturbance and by the duration of the disturbance (Kundur et al., 2004). With respect to disturbance size, large-disturbance voltage stability considers the time domain response of a system after a large disturbance such as a generator outage, while small-disturbance voltage stability considers system response to small perturbations about a particular operating point. With respect to time, short-term voltage stability considers time frames on the order of several seconds, while long-term voltage stability extends the analysis to potentially many minutes. Transient stability analysis, augmented with appropriate additional models such as generator overexcitation limiters and LTC transformer dynamics, can be used to assess many aspects of voltage stability. Figure 1.22 shows an example of a short-term voltage collapse scenario, using the 18,000-bus case and contingency from Figure 1.21, augmented with some additional contingencies. The thick red lines show the decline in the PU voltage magnitude at several 500-kV buses, the green lines three 230-kV buses, and the blue lines three 115-kV buses.

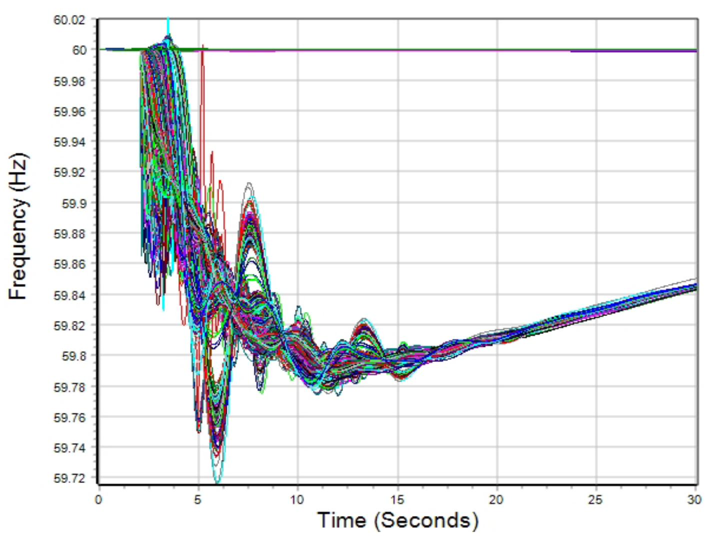

FIGURE 1.21 Frequency variation for generators following a generator outage contingency in a large grid. SOURCE: Courtesy of Thomas J. Overbye, University of Illinois, Urbana-Champaign.

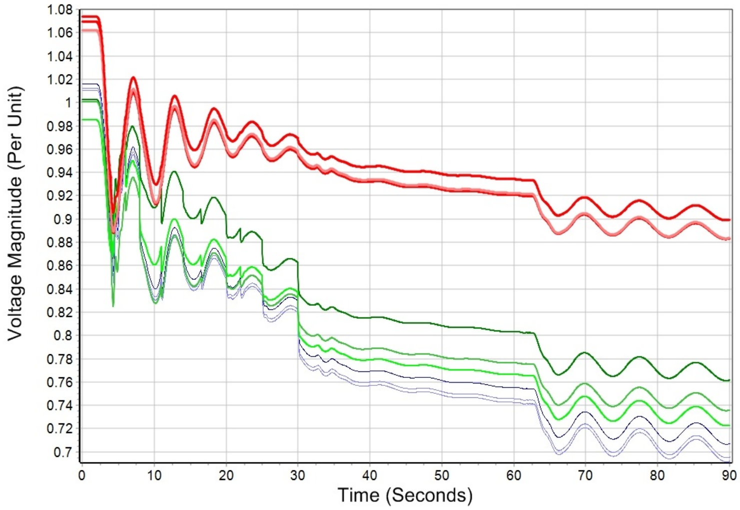

FIGURE 1.22 Short-term voltage collapse example modeling an 18,000-bus case using transient stability. SOURCE: Courtesy of Thomas J. Overbye, University of Illinois, Urbana-Champaign.

Usually power flow analysis is used to determine the risk to long-term, small-disturbance voltage stability by gradually increasing the system load (or some other parameter set) until the power flow equations no longer have a solution. This analysis is known as PV and/or QV analysis, since the bus voltage magnitudes are typically plotted with respect to either the real power (P) or the reactive power (Q) variation. An example of the PV curve for the two-bus case shown in Figure 1.23 is given in Figure 1.24. Here, provided the load is less than 500 MW, the power flow equations have two solutions, with the higher-voltage solution corresponding to the system operating point. A bifurcation point occurs at 500 MW, where the two solutions coalesce. This is the point of maximum loadabilility and would correspond to the long-term, small-disturbance voltage stability limit; the algebraic power flow equations have no solution for higher values of P. Having adequate reactive power is crucial to avoiding voltage instability. Several sources for reactive power are synchronous generators, static var compensators, and switched shunt capacitor banks. While the reactive power provided by the capacitors helps to support the voltage, a capacitor's reactive power output varies with the square of its voltage, meaning that as the voltage starts to fall its reactive power also rapidly decreases, resulting in a potential instability. Hence the reactive power supplied by the other devices can be crucial.

In referring back to Figure 1.4, the time scales of transient stability and voltage stability fall between the quasi-steady-state power flow and the faster switching surges, harmonics, and subsynchronous resonance. Like power flow, the assumption is that speed of light effects in the transmission network can be ignored, though this assumption is certainly less valid given that it can take 10 msec for light to transit a 2,000-mile grid, so the coupled algebraic power balance equations assumed in equation (15) cannot be fully true with a 4.16-msec time step. Even though the power system frequency is varying, the assumption is that the variation is small relative to the nominal frequency, with the branch impedances assumed fixed. So the transmission system is actually modeled assuming a fixed frequency. The impact of higher-frequency harmonics (e.g., 120, 180, . . . Hz) are not considered. Also, subsynchronous resonance, which might occur at frequencies between 10 and 30 Hz, an area of growing concern with wind farm installations, cannot be considered. Some generator and control system dynamics are included in transient stability, but faster ones (such as generator stator transients) are ignored.

Faster power system phenomena are usually studied by setting up a full three-phase model of the grid and then representing the transmission lines with the differential equations associated with the voltage and current relationships in inductors and capacitors. By using trapezoidal integration techniques, the models reduce to a network of coupled current sources and shunt resistances in which transmission line propagation delays can be considered

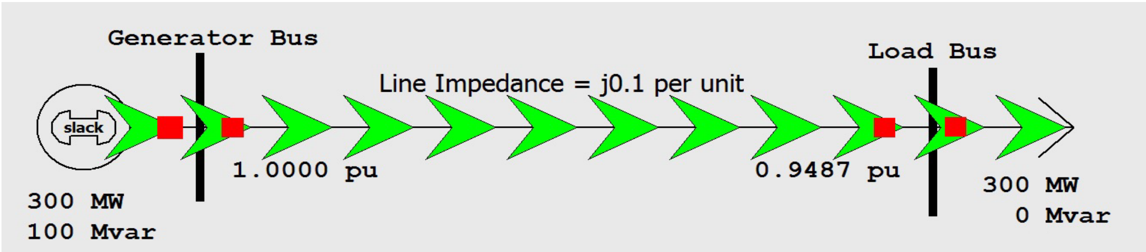

FIGURE 1.23 Two-bus case with generator supplying a constant power load. SOURCE: Courtesy of Thomas J. Overbye, University of Illinois, Urbana-Champaign.

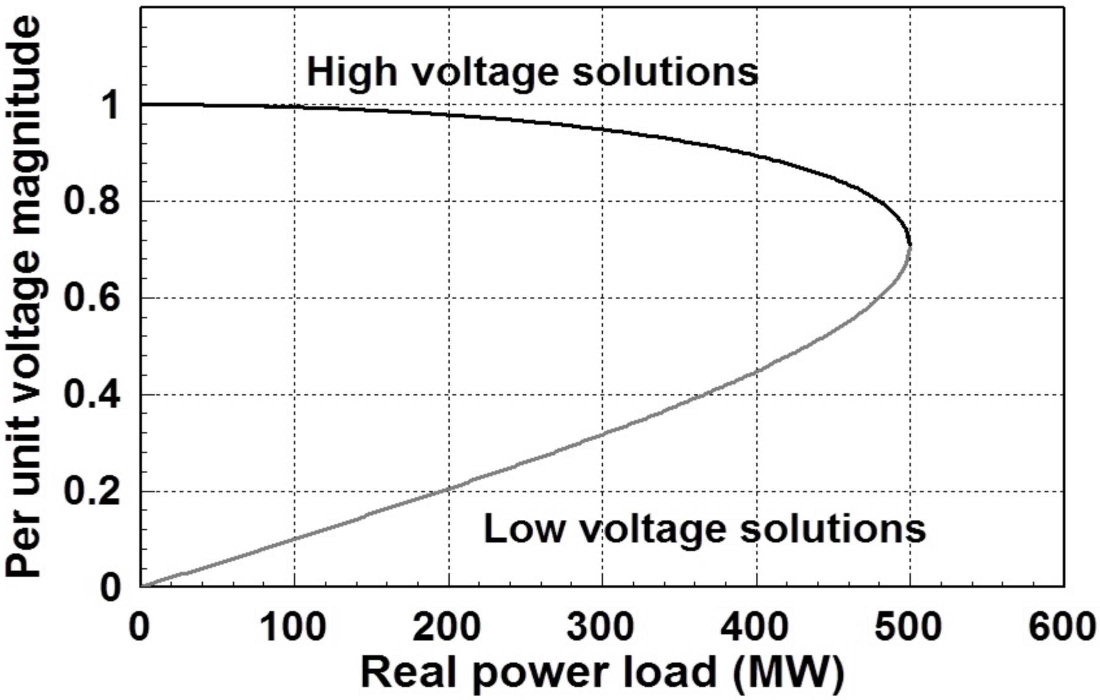

FIGURE 1.24 PV curve for the Figure 1.23 two-bus system. SOURCE: Courtesy of Thomas J. Overbye, University of Illinois, Urbana-Champaign.

explicitly. Integration time steps can then be as small as necessary to represent the phenomena of interest, such as 50 μsec for a switching transients study. An advantage of this approach is that a portion of a system can be solved in parallel since there is a natural decoupling due to the transmission line propagation delays. However, because of the small time constants, this approach requires large amounts of hardware for studies of even small systems.

Distribution Systems

As was mentioned earlier, the portion of the system that ultimately delivers electricity to most customers is known as the distribution system. This section provides a brief background on the distribution system as context for the rest of the report; further details are available in books such as Kersting (2012), Willis (2004), or Glover et al. (2012).

Sometimes the distribution system is directly connected to the transmission system, which operates at voltages above, say, 100 kV, and sometimes it is connected to a subtransmission system, operating at voltages of perhaps 69 or 46 kV. At the electrical substation, transformers are used to step down the voltage to the distribution level, with 12.47 kV being the most common in North America (Willis, 2004). These transformers vary greatly in size, from a few MWs in rural locations to more than 100 MW for a large urban substation.

The electricity leaves the substation on three-phase "primary trunk" feeders. While the distribution system can be networked, mostly it is radial. Hence on most feeders the flow of power has been one-way, from the substation to the customers. The number of feeders varies by substation size, from one to two up to more than a dozen. Feeder maximum power capacity can also vary widely from a few MVA to about 30 MVA. Industrial or large commercial customers may be served by dedicated feeders. In other cases smaller "laterals" branch off from the main feeder. Laterals may be either three phase or single phase (such as in rural locations). Most of the main feeders and laterals use overhead conductors on wooden poles, but in urban areas and some residential neighborhoods they are underground. At the customer location the voltage is further reduced by service transformers to the ultimate supply voltage (120/240 for residential customers). Service transformers can be either pole mounted, pad mounted on the ground, or in underground vaults. Typical sizes range from 5 to 5,000 kVA.

A key concern with the distribution system is maintaining adequate voltage levels to the customers. Because the voltage drop along a feeder varies with the power flow on the feeder, various control mechanisms are used. There include LTC transformers at the substation to change the supply voltage to all the substation feeders supplied by the transformer, voltage regulators that can be used to change the voltage for individual feeders (and sometimes even the individual phases), and switched capacitors to provide reactive power compensation.

Another key concern is protection against short circuits. For radial feeders, protection is simpler if the power is always flowing to the customers. Simple protection can be provided by fuses, but a disadvantage of a fuse is that a crew must be called in the event of it tripping. More complex designs using circuit breakers and re-closers allow for remote control, helping to reduce outage times for many customers.

While distribution systems certainly require substantial initial design (Willis, 2004, provides a good overview of planning considerations), the distribution system has traditionally "been characterized as the most unglamorous component" of an electric power system (Kersting, 2012). Most distribution systems are either unmetered or have customer meters that might be read only monthly, so that distribution systems are often overdesigned.

However, this is rapidly changing. With reduced costs for metering, communication, and control, the distribution system is rapidly being transformed. Distributed generation sources on the feeders, such as PV, mean that power flow may no longer be just one-way. Widely deployed advanced metering infrastructure systems are allowing near-real-time information about customer usage. Automated switching devices are now being widely deployed, allowing the distribution system to be dynamically reconfigured to reduce outage times for many customers. Advanced analytics are now being developed to utilize this information to help improve the distribution reliability and efficiency. Hence the distribution system is now an equal partner with the rest of the grid, with its challenges equally in need of the fundamental research in mathematical and computational sciences being considered in this report.

ORGANIZATION OF THE REPORT

Chapter 1, "Physical Structure of the Existing Grid and Current Trends," Chapter 2, "Organizations and Markets in the Electric Power Industry," and Chapter 3, "Existing Analytic Methods and Tools," lay out the current structure of the power grid, the economic markets involved in ultra-short-term decision making to long-term planning, and the analytic techniques that are currently used to study the behavior of the grid. Chapter 4, "Background: Mathematical Research Areas Important for the Grid," narrows the focus to the mathematics needed and currently used for these analyses.

Chapter 5, "Preparing for the Future," discusses the sources of uncertainty inherent in predicting the structure of the future grid and some of the general mathematical tools that may be needed. Chapter 6, "Mathematical Research Priorities Arising from the Electric Grid," examines research challenges for mathematics where progress will enable new technologies. Although a wide range of research areas have potential importance, this report will not discuss all of them in detail. Rather, it concentrates on the two areas that the committee felt were most relevant to the grid: optimization and dynamical systems. Advances in these areas could have widespread impacts, regardless of how the overall grid evolves.

In Chapter 7, "Case Studies," the report illustrates the diversity of mathematical problems, along with solutions, where they are now being solved. The problem of coordinating bid-based expenses with offer-based costs while satisfying regulatory, physical, operating, and business constraints—the unit commitment problem first described in Chapter 2—is given a representation as a mathematical programming problem. Other case studies illustrate the difficulty and importance of predicting low-frequency high-impact events, improving the resilience of the grid, and increasing the capability for handling anticipated massive amounts of data.

In Chapter 8, "Building a Multidisciplinary Research Community," the committee presents recommendations that will enlarge the community of researchers to include power engineers, mathematicians, and, potentially, other scientists—for example, statisticians and economists.

REFERENCES

Box, G.E.P., and N.R. Draper. 1987. *Empirical Model-Building and Response Surfaces*. Wiley, New York, N.Y.

California ISO. 2013. "What the Duck Curve Tells Us about Managing a Green Grid." http://www.caiso.com/documents/flexibleresourceshelp renewables_fastfacts.pdf.

EIA (U.S. Energy Information Administration). 2015. *Electric Power Monthly*. July. Washington, D.C.

Glover, J.D., M.S. Sarma, and T.J. Overbye. 2012. *Power System Analysis and Design,* 5th ed. Cengage Learning, Boston, Mass.

Huang, D., Y. Shu, J. Ruan, and Y. Hu. 2009. Ultra high voltage transmission in China: Developments, current status and future prospects. *Proceedings of the IEEE* 97(3):555-583.

Kersting, W.H. 2012. *Distribution System Modeling and Analysis,* 3rd ed. CRC Press, Boca Raton, Fla.

Kundur, P. 1994. *Power System Stability and Control*. McGraw-Hill Education, New York, N.Y.

Kundur, P., J. Paserba, V. Ajjarapu, G. Andersson, A. Bose, C. Canizares, and N. Hatziargyriou. 2004. Definitions and classification of power system stability. *IEEE Transactions on Power Systems* 19(3):1387-1401.

NERC (North American Electric Reliability Corporation). 2011. "Balancing and Frequency Control," January 26. http://www.nerc.com/docs/oc/rs/NERC%20Balancing%20and%20Frequency%20Control%20040520111.pdf.

Pai, M.A., P.W. Sauer and Thomas J. Overbye. 2012. Data-driven power system operations. Pp. 448-455 in *Computational Science—ICCS 2006* (V.N. Alexandrov, G.D. van Albada, and P.M.A. Sloot, eds.), Proceedings, Part III, Lecture Notes in Computer Science 3993. Springer-Verlag, Berlin, Heidelberg. Sauer, P.W., and M.A. Pai. 2007. *Power System Dynamics and Stability*. Stipes Publishing, Champaign, Ill. January 18.

Van Cutsem, T., and C. Vournas. 2007. *Voltage Stability of Electric Power Systems*. Springer, New York, N.Y.

Willis, H.L. 2004. *Power Distribution Planning Reference Book*, 2nd ed. CRC Press, Boca Raton, Fla.

Wood, A.J., B.F. Wollenberg, and G.B. Sheble. 2013. *Power Generation, Operation and Control*, 3rd ed. Wiley-Interscience, Hoboken, N.J.

2

Organizations and Markets in the Electric Power Industry

INTRODUCTION

In order to move forward with the mathematical and computational research needed to develop the next-generation electric grid, it is crucial to understand not just the physical structure of the grid but also its organizational structures.

Physically, a large-scale grid is ultimately an electrical circuit, joining the loads to the generators. However, it is a shared electrical circuit with many different players utilizing that circuit to meet the diverse needs of electricity consumers. This circuit has a large physical footprint, with transmission lines crisscrossing the continent and having significant economic and societal impacts. Because the grid plays a key role in powering American society, there is a long history of regulating it in the United States at both the state and federal levels.

Widespread recognition that reliability of the grid is paramount led to the development of organizational structures playing major roles in how electricity is produced and delivered. Key among these structures is the Federal Energy Regulatory Commission (FERC), the North American Electric Reliability Corporation (NERC), and federal, regional, and state agencies that establish criteria, standards, and constraints.

In addition to regulatory hurdles, rapidly evolving structural elements within the industry, such as demand response, load diversity, different fuel mixes (including huge growth in the amount of renewable generation), and markets that help to determine whether new capacity is needed, all present challenges to building new transmission infrastructure. With these and many other levels of complexity affecting the planning and operation of a reliable power system, the need for strong, comprehensive, and accurate computational systems to analyze vast quantities of data has never been greater.

HISTORY OF FEDERAL AND STATE REGULATION WITH REGIONAL STANDARDS DEVELOPMENT

Since the creation of Edison's Pearl Street Station in 1882, electric utilities have been highly regulated. This initially occurred at the municipal level, since utilities needed to use city streets to route their wires, necessitating a franchise from the city. In the late 1800s, many states within the United States formed public utility regulatory agencies to regulate railroad, steamboat, and telegraph companies. With the advent of larger electric power utility companies in the early 1900s, state regulatory organizations expanded their scopes to regulate electric power companies.

Some challenges related to large size and scalability have been offset by increasing computer processing capabilities. Nevertheless, to make a significant improvement in the efficiency and accuracy of the unit commitment process (the process of finding the least-cost dispatch of available generation resources to meet the electrical load) made possible by modeling and solving an ac power flow formulation, there must be either a substantial increase in processing capability or a reformulation of the problem. The ability to solve an ac unit commitment problem, compared to the linearized dc approximation solution in place today, could significantly improve the modeling and efficient dispatch of resources during the commitment, dispatching, and pricing processes.

Out-of-market actions taken by system operators have countered many of the shortcomings of the dc approximation currently in use. These actions typically are not well captured in the dc model, and the side effects create market inefficiencies, such as uplift payments and underfunded transmission rights.

Regulatory Development

Almost from their inception, electric utilities were viewed as a natural monopoly. Because of the high cost of building distribution systems and the social impacts associated with the need to use public space for the wires, it did not make sense to have multiple companies with multiple sets of wires competing to provide electric service in the same territory. Electric utilities were franchised initially by cities and later (in the United States) by state agencies. An electric utility within a franchised service territory "did it all." This included owning the increasingly larger generators and the transmission and distribution system wires, and continued all the way to reading the customer's meters. Customers did not have a choice of electric supplier (many still do not). Local and state regulators were charged with keeping electric service rates just and reasonable within these franchised service territories.

As electric utilities grew and expanded, holding companies formed that allowed for the rapid growth of the electric utility industry. This growth created regulatory challenges for local and state utility commissions. Many holding companies engaged in interstate commerce, which went beyond local and state commissions' regulatory authority and capacity.

Following the crash of the stock market in 1929, the U.S. Congress passed the Public Utility Holding Company Act of 1935, which increased the regulation of electric utilities by limiting their operations to a single state or forcing divestiture so that each public utility company served only a limited geographic area (EIA, 1993).

In 1920, Congress had created the Federal Power Commission (FPC) to coordinate hydroelectric projects under federal control. The Federal Power Act of 1935 and the Natural Gas Act of 1938 gave the FPC the power to regulate the sale and transportation of electricity and natural gas. This power subsequently expanded to include the regulation, sale, and transportation of interstate electricity and natural gas.

In 1967, the FPC recommended the formation of a council on power coordination made up of representatives from each of the nation's regional coordinating organizations to assist in resolving interregional coordination matters (FERC, 2015d).

Reliability Organization Development

On June 1, 1968, the electricity industry formed NERC in response to the FPC recommendation and the 1965 blackout, when 30 million people lost power in the northeastern United States and southeastern Canada.

In 1973, the utility industry formed the Electric Power Research Institute to pool research and improve reliability.

After another blackout occurred in New York City in July 1977, Congress reorganized the FPC into the Federal Energy Regulatory Commission (FERC, 2015d) and expanded the organization's responsibilities to include the enactment of a limited liability provision in federal legislation, allowing the federal government to propose voluntary standards (NERC, 2013a).

In 1980, the North American Power Systems Interconnection Committee (known as NAPSIC) became the Operating Committee for NERC, putting the reliability of both planning and operation of the interconnected grid under one organization.

In 1996, two major blackouts in the western United States led the members of the Western System Coordinating Council to develop the Reliability Management System. Members voluntarily entered into agreements with the council to pay fines if they violated certain reliability standards (NERC, 2013a).

In response to the same two western blackout events, NERC formed a blue-ribbon panel and the Department of Energy formed the Electric System Reliability Task Force. These independent investigations led the two groups to recommend separately the creation of an independent, audited self-regulatory electric reliability organization to develop and enforce reliability standards throughout North America.

Both groups concluded that federal regulation was necessary to ensure the reliability of the North American electric power grid. Following those conclusions, NERC began converting its planning policies, criteria, and guides into reliability standards (NERC, 2013a).

On August 14, 2003, North America experienced its worst blackout to that date, with 50 million people losing power in the midwestern and northeastern United States and in Ontario, Canada (NERC, 2013).

On August 8, 2005, the Energy Policy Act of 2005 authorized the creation of an electric reliability organization and made reliability standards mandatory and enforceable. On July 20, 2006, FERC certified NERC as the electric reliability organization for the United States. From September through December 2006, NERC signed memoranda of understanding with Ontario, Quebec, Nova Scotia, and the National Energy Board of Canada (NERC, 2013a).

Following the execution of these agreements, on January 1, 2007, the North American Electric Reliability Council was renamed the North American Electric Reliability Corporation. Following the establishment of NERC as the electric reliability organization for North America, FERC approved 83 NERC Reliability Standards, representing the first set of legally enforceable standards for the bulk electric power system in the United States.

On April 19, 2007, FERC approved agreements delegating its authority to monitor and enforce compliance with NERC reliability standards in the United States to eight regional entities, with NERC continuing in an oversight role (NERC, 2013b).

North American Regional Entities

There are many characteristic differences in the design and construction of electric power systems across North America that make a one-size-fits-all approach to reliability standards across all of North America difficult to achieve. A key driver for these differences is the diversity of population densities within North America, which affects the electric utility design and construction principles needed to reliably and efficiently provide electric service in each different area.

There are eight regional reliability organizations covering the United States, Canada, and a portion of Baja California Norte Mexico (Figure 2.1). The members of these regional entities represent virtually all segments of the electric power industry and work together to develop and enforce reliability standards, while addressing reliability needs specific to each organization (NERC, 2013b).

Changes in Regulation

Starting in the late 1970s, FERC began to discuss and explore ways to deregulate the electric power industry to comply with reform of the Public Utility Holding Company Act. A series of orders set in motion the legal framework to provide consumers with a greater choice of suppliers for their electricity.

The Energy Policy Act of 1992 established these changes. FERC followed up on this effort in 1996 with Orders 888 and 889, which required transmission owners to provide "equal and open access" to others seeking to transmit energy over transmission owners' facilities (FERC, 2015d).

FERC issued Order 2000 in 1999 to ensure equal access to transmission within the United States. Order 2000 created the framework for the formation of independent system operators (ISOs) and regional transmission organizations (RTOs) and provided the specific criteria that transmission entities must meet in order to qualify as an ISO or RTO. The formation of these organizations was deemed necessary in order to create transparent electricity trading markets where independent market operators calculate and post prices to facilitate the efficient sale of electricity (FERC, 2015d).

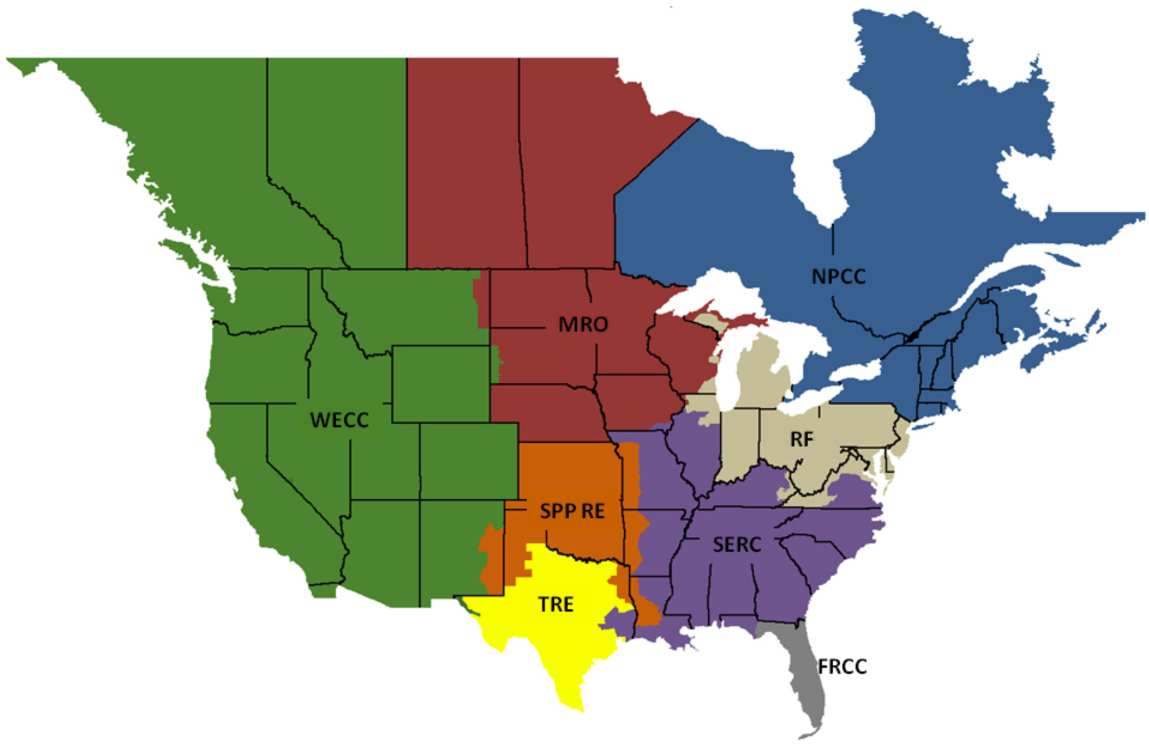

FIGURE 2.1 North American regional reliability organizations. SOURCE: NERC (2013b). This information from the North American Electric Reliability Corporation's website is the property of the North American Electric Reliability Corporation and is available at http://www.nerc.com/AboutNERC/keyplayers/Pages/default.aspx. This content may not be reproduced in whole or any part without the prior express written permission of the North American Electric Reliability Corporation.

U.S. WHOLESALE POWER MARKETS

As mentioned in Chapter 1, one of the many benefits an interconnected electric grid provides is the ability for different players to enter into power transactions. Initially these were bilateral transactions between neighboring electric utilities sharing transmission tie lines. This exchange of power initially took place by the utilities adjusting their area control error calculations (the difference between generation and load scheduled between geographic areas). These calculations enabled operators to determine how the exchanges of power would take place.

As communication technologies improved, utility members developed power pools to allow sharing of generation resources. For example, PJM started in 1927 with three electric utilities sharing generation resources and jointly planning transmission. Both the New York power pool and the New England power pool were formed to take advantage of economy of scale. By the 1980s, improved communication and computer technologies enabled the automatic exchange of buy-and-sell offers for electricity to take place on an hourly basis. Thus, markets were established so electricity providers and users could benefit from competition and transparent pricing.

FERC's Order 2000 led to the eventual formation of nine organized wholesale power markets in the United States and Canada (as of May 2015), with general geographic locations shown in Figure 2.2. Many areas of the United States are not presently covered by organized markets, however. Utilities in those areas continue to engage in power transfers through other means such as bilateral transactions. Different practices are employed in different areas, and the institutional structure may change. Currently Duke Power, Southern Company, and the Tennessee Valley Authority in the Southeast, and the Bonneville Power Administration in the Pacific Northwest are among those that largely utilize bilateral transactions. The organized wholesale power markets in the United States are these:

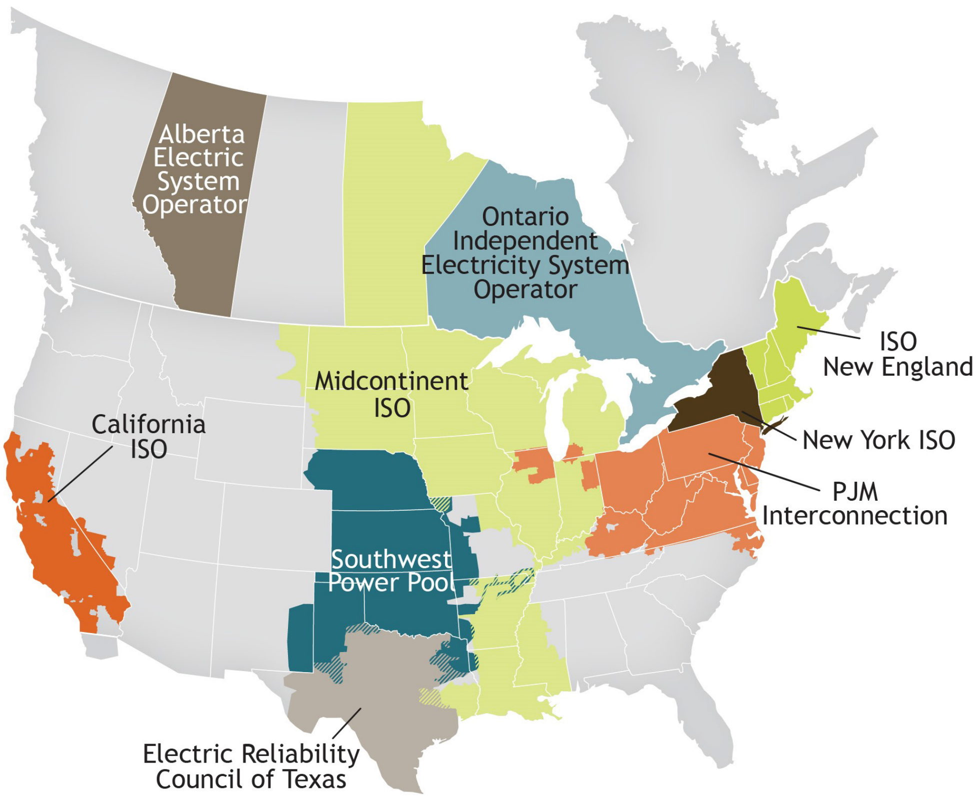

FIGURE 2.2 North American regional transmission organizations. SOURCE: Courtesy of California Independent System Operator. Licensed with permission from the California ISO. Any statements, conclusions, summaries or other commentaries expressed herein do not reflect the opinions or endorsement of the California ISO.

- *California Independent System Operator (California ISO, or CAISO)*. A single-state independent system operator that serves approximately 30 million customers across 26,000 miles of transmission lines with an installed generation capacity of 65,000 MW.
- *Electric Reliability Council of Texas (ERCOT*). A single-state independent system operator that serves approximately 24 million customers across more than 43,000 miles of transmission lines with an installed generation capacity of 86,000 MW.
- *ISO New England (ISO-NE)*. A multistate independent system operator that serves approximately 14 million customers across 8,130 miles of transmission lines with an installed generation capacity of approximately 32,000 MW.
- *Midcontinent ISO (MISO)*. A multistate independent system operator that serves approximately 48 million customers across 65,250 miles of transmission lines with an installed generation capacity of 205,759 MW.
- *New York Independent System Operator (NYISO)*. A single-state independent system operator that serves approximately 19.5 million customers across 11,056 miles of transmission lines with an installed generation capacity of 37,978 MW.
- *PJM Interconnection (PJM)*. A multistate independent system operator that serves approximately 61 million customers across 62,556 miles of transmission lines with an installed generation capacity of 183,604 MW.
- *Southwest Power Pool (SPP)*. A multistate independent system operator that serves approximately 6.2 million customers across 48,930 miles of transmission lines with an installed generation capacity of 77,366 MW.
- *Alberta Electric System Operator* and *Ontario Independent Electricity System Operator* are also shown in Figure 2-2.

Common Features of Electric Markets

While there are certainly differences among these markets, their physical power systems and regulatory constraints require them to provide certain common basic services. For example, because of the unit commitment and start-up constraints, all have a day-ahead market in which generator commitment decisions are made the previous day to ensure that there is sufficient generation available to match the forecast electric load.

However, since day-ahead forecasts are never perfect and unexpected events can occur, all of these markets also have a real-time energy market, sometimes referred to as a balancing market, to ensure that actual generation matches the load on a continuous basis. Other commonalities are financial transmission rights markets, regulation markets, reserve markets, and synchronized reserve markets.

These markets currently also offer demand response service—programs that encourage consumers to reduce their use of electricity during certain high-demand periods in return for a reduction in their power bills. In addition, all U.S. wholesale power markets have binding must-offer requirements in their day-ahead markets and also allow virtual transaction[1] participation (IRC, 2014).

Electricity Market Co-optimization

The energy markets operated by each of these ISO/RTOs allow agreements to be set 1 day ahead and in real time. Energy prices are co-optimized with some ancillary service products—which are the additional requirements, beyond just electric energy, for operating the electric grid. Examples include reactive power needed to maintain transmission voltages within acceptable limits, load frequency control to continually match the total generation to the total load (plus losses), power supply reserves to bring generation and load back into balance following an unexpected loss of generation, and black-start services to restart critical generation after a large-scale system blackout.

One natural co-optimization is between the generation capacity that is used to produce energy and the capacity available for reserve use, because the same capacity resources cannot provide both energy and reserves simultaneously. Co-optimization is typically simultaneous with market solutions and occurs every 5 min in near real time with the objective function to minimize the cost of production. For example, CAISO co-optimizes energy and ancillary services in its 15-min real-time unit commitment interval; however, its dispatch at the 5-min interval is not co-optimized with ancillary services (IRC, 2014).

PRICING

At present, all U.S. energy markets that calculate their locational marginal prices (LMPs) use a variation of the generic security-constrained optimal power flow (SCOPF) algorithm covered in Chapter 1. LMPs set wholesale electric energy market prices based on actual operating conditions at a specific time and place. When transmission system limits prevent the lowest-priced electricity from flowing into a location, grid operators dispatch more expensive generation to meet the demand. As a result, the wholesale price of electricity will likely be higher in that location than elsewhere.

In this approach offers are received from the resources (e.g., generators), indicating the price at which they will provide a certain quantity of electric energy to the market (e.g., 50 MW for 1 hr at $100/MWh). In the day-ahead market these offers are matched with the forecast electric load in the SCOPF to determine which generators will actually be committed. Resources that are priced too high are not accepted and hence do not run.

The SCOPF then calculates the LMPs. As covered in Chapter 1, "Physical Structure of the Existing Grid and Current Trends," the LMPs indicate the marginal cost to provide electricity to a particular bus in the grid. In the absence of any transmission system constraints (congestion) or any incremental transmission losses, the LMPs at all the buses would be identical. This is seldom the case, however, so each LMP usually includes a congestion component. In addition, six U.S. energy market LMP values include a marginal transmission loss component. The

[1] Virtual transactions are a fundamental component of two-settlement markets; they are used to arbitrage price differences between the day-ahead market and the real-time market.

lone exception is ERCOT, where losses are charged to load based on a load ratio share. Figure 2.3 shows LMP superimposed on the PJM footprint. Interestingly, sometimes the LMPs can actually be negative. In such locations generators would have to pay to produce electricity, and customers would get paid for their usage.

Energy Pricing Example

As an illustration of LMP pricing and the impact of energy offer caps, Figure 2.4 shows the seven-bus system from Chapter 1 with all generators offering to provide power at their previous prices, except the generator at bus 4, offering in at the maximum of $1,000/MWh. Not unsurprisingly, these units are not dispatched and receive no revenue. Their high offer has changed the LMPs throughout the system and altered the system congestion, so that two lines are now loaded to their maximums. However, if the system load is slightly different with a modest increase in the bus 3 load, for example, the LMPs change drastically, with the bus 4 generator now being dispatched at 17 MW, having an LMP of $1,000/MWh, and receiving $17,000/hr in revenue. This example is illustrated in Figure 2.5 and illustrates the potentially high sensitivity of the LMPs to small system changes.

Day-Ahead Market

CAISO, ERCOT, MISO, NYISO, PJM, and SPP energy and ancillary services are simultaneously co-optimized in their day-ahead markets. ERCOT's day-ahead market also simultaneously co-optimizes congestion-hedging products with energy and ancillary services by maximizing bid-based revenues and minimizing offer-based costs (subject to resource and network constraints). ISO-NE's day-ahead market respects operating reserve requirements but does not co-optimize energy and operating reserves (IRC, 2014).

Unit Commitment

In the day-ahead time frame, the CAISO, ERCOT, ISO-NE, MISO, PJM, and SPP markets employ a day-ahead reliability unit commitment process. CAISO's process physically commits resources for reliability based

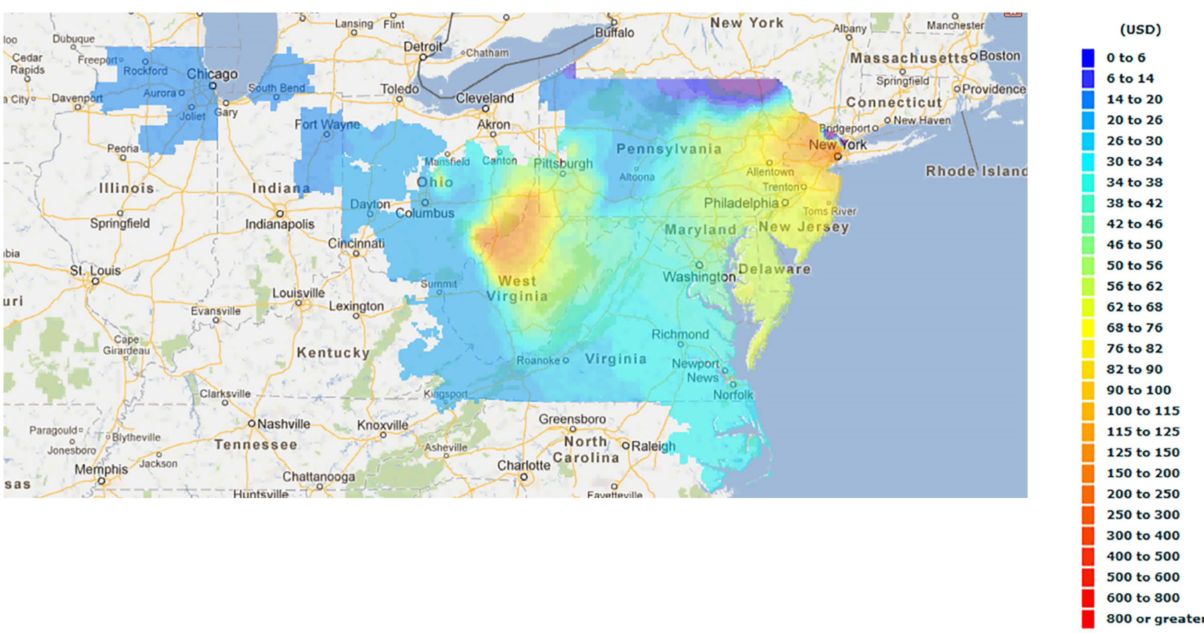

FIGURE 2.3 LMP superimposed on the PJM footprint. Units are $/MWh. SOURCE: Courtesy of PJM Interconnection.

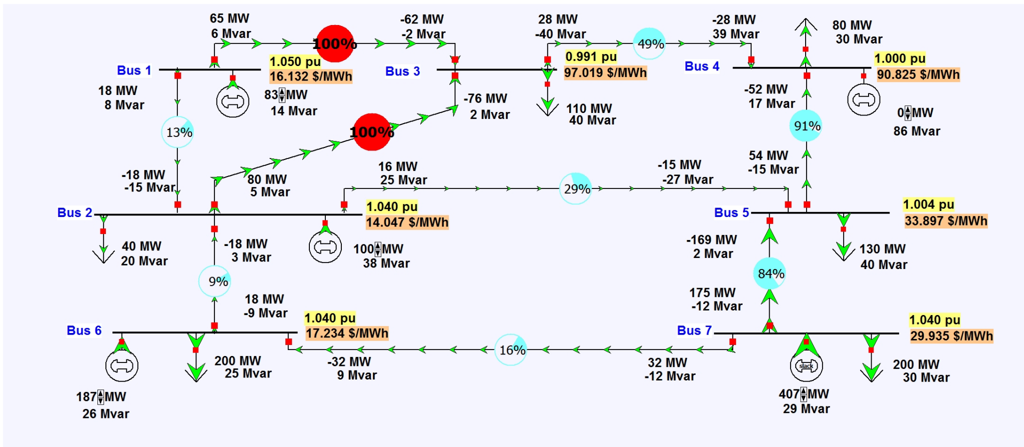

FIGURE 2.4 OPF solution of original seven-bus system with generator at bus 4 offering high.

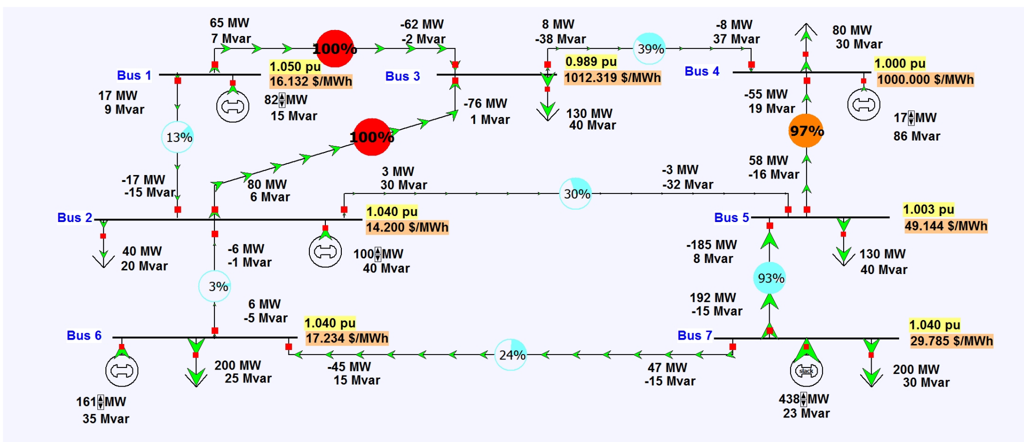

FIGURE 2.5 OPF solution of modified seven-bus system with generator at bus 4 offering high. SOURCE: Courtesy of Thomas J. Overbye, University of Illinois, Urbana-Champaign.

on forecast demand. Results are published simultaneously upon completion of the day-ahead market process. ERCOT runs a daily and an hourly process to ensure enough generators are online to meet the load forecast. MISO performs its process at 4 p.m. after the day-ahead market closes. MISO also performs additional evaluations as needed throughout the operating day.

PJM performs its commitment at 6 p.m. after the day-ahead market closes and as needed until the operating day begins. SPP's commitment runs prior to, and commits resources for, the operating day to minimize the cost of capacity as computed based on resource offers. SPP also performs an intraday reliability unit commitment process that runs at least every 4 hours to maintain system balance.

The optimization for the day-ahead market uses a dc power flow and a mixed integer program for optimization. The dc power flow does not represent voltage constraints well. With faster supercomputers and improved analytics, a solution with a full ac power flow may now be possible.

PJM also operates a day-ahead scheduling reserve market used to ensure energy reserves are available for up to 30 minutes to deal with unexpected system conditions during the operating day (IRC, 2014).

Scarcity Pricing

LMP markets are designed to send transparent price signals to market participants to make both short-term decisions about how to use existing resources (either generation or load), as well as to make longer-term decisions about where to locate new resources to reduce system congestion.

Most of the time the markets work as planned. However, scarcity can occur when there are not enough resources to balance the generation and the load, a potential system emergency. At such times, special scarcity pricing mechanisms are used to both incentivize the existing resources to help maintain grid reliability and allow longer-term entry of new resources at the best locations on the grid. Currently, CAISO, ERCOT, MISO, NYISO, PJM, and SPP have scarcity pricing mechanisms in place.

In CAISO, both ancillary services and shortage pricing are based on the scarcity demand curve, which takes into account both relevant regulations and power supply reserves. CAISO's highest scarcity price is $1,000/MWh, which is the energy bid cap.

ERCOT uses an operating reserve demand curve to set scarcity pricing and has a cap of $9,000/MWh under scarcity conditions. MISO's scarcity pricing is based on the regulation and operating reserve demand curve, with the regulation offer cap at $500/MWh and the supplemental offer cap at $100/MWh.

NYISO shortage pricing is included within its optimization tools for energy and ancillary services. PJM sets shortage prices using an operating reserve demand curve that sets the price to serve as a "penalty factor" for being unable to meet the reserve requirement. The penalty factor cap is $850/MWh as of June 2015 (IRC, 2014).

Capacity Markets

Capacity markets are a mechanism to help ensure that there will be sufficient generation capacity to match the anticipated load. This is needed because construction of new generation facilities can take a long time—usually well over a year. Currently, only ISO-NE, MISO, NYISO, and PJM contain capacity market constructs, albeit these constructs differ considerably. Here is the high-level description of each capacity market:

- *ISO-NE*. The ISO-NE forward capacity market is designed to acquire qualified resources 3 years in advance of the commitment period. This meets the installed capacity requirement established by ISO-NE system planning and the stakeholder process, and encourages the retention and development of resources to maintain adequate operating reserves. There is a forward capacity auction each year, where market participants obtain a capacity supply obligation to deliver physical capacity by the start of a commitment period. Capacity supply obligations can be acquired or shed bilaterally or in reconfiguration auctions.
- *MISO*. MISO has a resource adequacy capacity market that voluntarily procures capacity in an annual auction to supply energy based on aligned regulations.
- *NYISO*. The NYISO conducts three types of installed capacity (ICAP) auctions: the capability period auction, the monthly auction, and the ICAP spot market auction. On any day for which it supplies unforced capacity, a provider is obligated to schedule or bid into the day-ahead market (or declare itself to be unavailable) an amount of energy that is not less than the installed capacity equivalent of the amount of unforced capacity it is supplying to the New York control area.
- *PJM*. The PJM capacity market—the reliability pricing model—procures long-term capacity resources 3 years ahead, whereby committed, dispatchable resources are obligated to offer into the day-ahead market.

U.S. BILATERAL MARKETS

Markets in the U.S. Southeast and Northwest are made up predominantly of vertically integrated utilities engaging in bilateral transactions between two market participants.

The Northwest electric region covers Idaho, Montana, Nevada, Oregon, Utah, Washington, Wyoming, and a small portion of northern California. Approximately two-thirds of the electric power production comes from hydroelectric sources in the northwestern United States. The surplus power production has historically been sold into California and the U.S. Southwest (FERC, 2015a).

The Southeast electric region covers all or portions of Alabama, Florida, Georgia, Mississippi, North Carolina, South Carolina, and Tennessee. Coal and natural gas have been the predominant marginal fuel types in this region (FERC, 2015b).

The Southwest electric market encompasses the Arizona, New Mexico, southern Nevada, and the Rocky Mountain Power Area (RMPA) subregions of the Western Electric Coordinating Council. Peak demand is approximately 42 GW in summer. There are approximately 50 GW of generation capacity, composed mostly of gas and coal units.

The Southwest relies on nuclear and coal generators for baseload electricity, with gas units generally used as peaking resources. The coal generators are generally close to coal mines, resulting in low delivered fuel costs. Some generation is jointly owned among multiple nearby utilities, including the Palo Verde nuclear plant, the nation's largest nuclear plant, with three units totaling approximately 4,000 MW, which has owners in California and the Southwest (FERC, 2015c).

EXAMPLES OF INTERNATIONAL MARKETS

Open access to transmission and competition in wholesale markets has been developing around the world for over 25 years. The United Kingdom was first. Australia has the highest price cap on its "energy only" market. Japan is the latest country to move to markets to encourage investment to replace the nuclear capacity that was shuttered after the March 11, 2011, Tōhoku magnitude 9.0 earthquake and tsunami. It can be useful to understand how these and other foreign markets are evolving.

Australian National Electricity Market

The Australian National Electricity Market (NEM) has five trading regions that geographically cover the eastern and southern portion of Australia: Queensland, New South Wales, Victoria, South Australia, and the island state of Tasmania, which is connected underwater to Victoria. The NEM produces approximately 200 TWh (terawatt-hour) of electricity annually, which makes up about 80 percent of Australia's total energy consumption from an installed generation capacity of about 45,900 MW (AEMC, 2015).

The NEM is an energy-only pool, where all generators are required to sell all output into the market. It matches generation bids with load requirements every 5 minutes in the most cost-efficient manner and provides dispatch instruction to generators whose bids clear. The NEM averages the 5-minute prices and posts spot prices every 30 minutes for each of the five trading regions. The market price cap is set by the NEM rules and is currently \$12,500/MWh. Additionally, the NEM has a price floor of –\$1,000/MWh (EEX, 2015).

The Australian Energy Market Operator (AEMO) operates the NEM as well as gas markets in Australia. The two markets are cleared at the same time but are not co-optimized. The gas markets include the Declared Wholesale Gas Market in Victoria, the Short-Term Trading Markets, and the Wallumbilla Gas Supply Hub. AEMO's corporate brochure states "AEMO oversees the vital system operations and security of the National Electricity Market Power System and the Victorian Gas Declared Transmission System" (AEMO, 2014).

AEMO notes that electricity consumption in the NEM began to decline in 2010. This was due to a combination of factors, including commercial and residential customers reducing consumption in response to higher electricity retail prices; increasing energy efficiency measures; increases in solar photovoltaic installations (thanks to government incentives); and weaker energy demand from industrial manufacturing facilities (AEMO, 2014).

German Electricity Market

The German electricity market is the largest in Europe, with an installed generation capacity of approximately 125 GW. The current annual energy consumption within the German electricity market is around 550 TWh per year. Four transmission system operators run the transmission system in Germany. They procure primary, secondary, and tertiary control reserves via the German Control Reserve Market, which allows participation by plant operators and electricity consumers (German Transmission System Operators, 2015).

Following the Fukushima nuclear plant accident in Japan, the German government made the decision to shut down the country's eight oldest nuclear power plants. The country plans to shut down the remaining nine nuclear power plants by 2022. Germany has also set aggressive climate change initiatives to reduce greenhouse gas emissions. Its targets exceed those set by the Kyoto Protocol, which called for a reduction (compared to 1990 levels) in greenhouse gas emissions of 20 percent by 2020. Germany has established a national target for reductions of 40 percent by 2020 and 80 to 95 percent by 2050.

To make up for the changing energy resource mix, Germany has put in place an aggressive program of increasing supplies of energy from renewable resources. It aims to have between 40 and 45 percent of the energy consumed to be sourced from renewable resources by the year 2025. Despite the increase in renewables, coal has made a comeback in Germany because of the country's goal to shutter nuclear plants.

In order to fund this investment in renewable resources, Germany introduced feed-in tariffs, which have increased retail prices while they have lowered the wholesale price of energy in Germany, creating revenue challenges for existing resources. As a result of the feed-in tariffs, the average price of electricity for industrial customers in Germany is above the European Union average and significantly above the average price of power for U.S. industrial customers (German Federal Ministry for Economic Affairs and Energy, 2015).

CONCLUSIONS

The day-ahead market unit commitment problem is the most complex math problem solved by most of the ISO/RTOs that operate power markets. There are many other optimization programs used by ISO/RTOs to schedule and dispatch the electric power system and clear power markets, but they are simplifications of the unit commitment problems.

The most common challenge ISO/RTOs encounter with market-clearing engines, and specifically the unit commitment problem, relates to problem size and scalability. Several factors, including significant increases in bid and offer volumes, large numbers of transmission constraints, and a large number of continuous and binary variables, can slow solution time. The interaction between them can increase problem complexity exponentially.

Progress toward better modeling of ac constraints on the system would significantly advance the accuracy and efficiency of the unit commitment solution and the overall market. The committee envisions that the next-generation grid will only increase in complexity.

Recommendation 1: The Department of Energy should develop and test a full ac optimal power flow (ACOPF) model with an optimization algorithm using all nodes in the market area, taking advantage of supercomputers and parallel processing, and respecting all thermal and voltage constraints.

The committee believes that this research should include several versions of the ACOPF. It should include modeling all nodes in the market area reflecting different degrees of nonlinearity, size, and connectedness, and respecting all thermal and voltage constraints. In order to enable a solution of the much larger models necessitated by the next-generation grid, computational testing of these models should be undertaken with the flexibility to allow advantage to be taken of advances in high-end computing including parallel and distributed systems. The optimization methods used should reflect a careful examination of the validity and effectiveness of emerging techniques for nonlinear, nonconvex optimization that are based on applying techniques for convex and mixed-integer optimization and reduced-order modeling, as called for in Recommendation 8.

REFERENCES

AEMC (Australian Energy Market Commission). 2015. "National Electricity Market." http://www.aemc.gov.au/Australias-Energy-Market/Markets-Overview/National-electricity-market. Accessed March 2015.

AEMO (Australian Energy Market Operator). 2014. "AEMO Pivotal to Australia's Energy Future." Brochure. http://www.aemo.com.au/~/media/Files/Other/corporate/AEMO_Corp_Brochure14.2.ashx. July.

EEX (Energy Efficiency Exchange). 2015. "How the Energy Market Operates." http://eex.gov.au/energy-management/energy-procurement/energy-pricing/how-the-energy-market-operates. Accessed March 2015.

EIA (Energy Information Administration). 1993. Public Utility Holding Company Act of 1935: 1935-1992. U.S. Department of Energy. http://www.eia.gov/electricity/archive/0563.pdf. Accessed March 2015.

FERC (Federal Energy Regulatory Commission). 2015a. "Electric Power Markets: Northwest." http://www.ferc.gov/market-oversight/mkt-electric/northwest.asp. Accessed March 2015.

FERC. 2015b. "Electric Power Markets: Southeast." http://www.ferc.gov/market-oversight/mkt-electric/southeast.asp. Accessed March 2015.

FERC. 2015c. "Electric Power Markets: Southwest." http://www.ferc.gov/market-oversight/mkt-electric/southwest.asp. Accessed December 15, 2015.

FERC. 2015d. "History of FERC." http://www.ferc.gov/students/ferc/history.asp. Accessed March 2015.

German Federal Ministry for Economic Affairs and Energy. 2015. "Energy of the Future: An Overall Strategy for the Energy Transition." http://www.bmwi.de/EN/Topics/Energy/Energy-Transition/overall-strategy.html. Accessed March 2015.

German Transmission System Operators. 2015. "Market for Control Reserve in Germany." https://www.regelleistung.net/ip/action/static/market-information. Accessed March 2015.

IRC. 2014. "Market Design Executive Summary." IRC Markets Committee. October 9. http://ircweb.businesscatalyst.com/Documents/Report/20150312_2014IRCMarketDesignExecutiveSummary.pdf.

IRC (Independent System Operator/Regional Transmission Organization (ISO/RTO) Council). 2015. "About the IRC." http://www.isorto.org/about/default. Accessed March 2015.

NERC (North American Electric Reliability Corporation). 2013a. "History of NERC." http://www.nerc.com/AboutNERC/Documents/History%AUG13.pdf. August.

NERC. 2013b. "Regional Entities." http://www.nerc.com/AboutNERC/keyplayers/Pages/Regional-Entities.aspx. Accessed March 2015.

3

Existing Analytic Methods and Tools

INTRODUCTION

Since the genesis of electric grids in the early 1880s the amount of expertise that has been gained in their operation and planning is enormous. And over the last 60 years, much of that expertise has been embedded in software that is of increasing complexity. In order to estimate the mathematical and computational research needed for the next-generation electric grid it will be important to first consider some of the existing analytic methods that are used in this already quite sophisticated software to plan and operate the grid of today. Building on the foundation described in Chapters 1 and 2, such consideration is the purpose of this chapter. Chapter 4 then covers the mathematical underpinnings for these algorithms.

An important caveat is that the algorithms presented here were developed to meet the needs of a grid that was initially operated by vertically integrated utilities that obtained most of their electricity from centrally controlled, large-scale generators. Scenarios for how the grid of the future might evolve are discussed in Chapter 5, with subsequent chapters discussing how the techniques presented here will need to evolve and research will need to move forward. Another important caveat: The committee focuses here on common algorithms that are widely used by the power industry. Given the wide-ranging scope of the literature in the field, it would be impossible to comment on all of the many approaches that have been proposed for solving the cited problems. And even with this caveat the chapter can only scratch the surface of power grid applications, providing brief coverage for some of the key algorithms, their underlying assumptions, and their approximations.

POWER FLOW (LOAD FLOW)

As was described in the first chapter, power flow is the key to solving quasi-steady-state power balance equations, allowing calculation of the per-unit voltage magnitude and angle at every bus (node) in the system. Usually the power flow also involves calculating the values for a host of continuous and discrete power system controllers as well, such as tap positions for load-tap-changing (LTC) transformers and the status of switched reactive control devices such as capacitors. Once all of these values have been determined, the power flowing on all of the system transmission lines and transformers (branches) can be determined. The power flow solution can then be used to check whether system quantities are within their limits (e.g., transmission line and transmission flows are less than the limits, and voltage magnitudes are between their minimum and maximum limits). Power flow is probably the most common power system analysis technique.

Historically the first power flows were done using analog computers, known as network analyzers. As digital computers started to become common in the 1950s and 1960s the use of network analyzers was replaced by numerical techniques. Since the power flow equations are nonlinear, their solution required an iterative approach. The Gauss-Seidel approach was initially the most common technique since its solution worked well on computers with limited memory, with 50-bus systems solved in the 1950s on computers with 2 kilobytes (kB) of memory (Brown and Tinney, 1957). However, convergence could be slow. While still taught and occasionally used, it has mostly been replaced by algorithms based on the Newton-Raphson (NR) approach.

Power flow solution by the NR method, or some variation of it, is currently the most common technique, having been introduced in 1967 (Tinney and Hart, 1967). The NR method takes advantage of the fact that each bus in a power system is joined to only a handful of other buses, making the network incidence matrix sparse and leading to a sparse Jacobian matrix. Taking advantage of the development of sparse matrix methods in the 1960s, including improved ordering algorithms pioneered to a large extent by power engineers (Sato and Tinney, 1963), the NR method could be applied to larger systems on computers with limited memory. In 1967 (Tinney and Hart, 1967) a 949-bus system was solved, with each iteration taking 10 seconds on an IBM 7040 and leading to speculation that a 2,000-bus network could be solved on a computer with 32 kB of memory. The quadratic convergence of Newton's method allows solutions of even large systems in just a few iterations. However, Newton's method-based algorithms can fail to converge for ill-conditioned problems, which can occur when the system voltage is outside a "normal" range. Improving the convergence of the power flow has been an area of active research for many years (Iwamoto and Tamura, 1981) and continues to require research. This is partially due to the fractal domains of attraction for power flow solutions (Demarco and Overbye, 1988; Thorp and Naqavi, 1989).

Over the years several enhancements were proposed, including the fast decoupled power flow (Stott and Alsac, 1974), which eliminated the need for updating the Jacobian inverse at every iteration, and the transformation of the Jacobian/network matrix to orthogonal eigenvectors, which has the potential for speedier solutions and improved iterations. The dc power flow model is also common, particularly in power market analysis, in which the nonlinear equations are approximated by a set of linear equations, eliminating the need for the iterative NR algorithm.[1]

The largest power flow cases routinely solved now contain at most 100,000 buses.[2] While this is a significant increase from the hundreds of buses modeled 50 years ago, the size of the largest power flow problems (with on the order of a million nonzeros in the Jacobian) is now quite modest compared to both the growth in computer memory and problems in other domains, which may have billions of nonzeros. Interconnected system power flow models are not expected to grow in size significantly because the radial distribution system can be effectively aggregated for interconnect-wide studies.

There are actually two types of power flow models. The one that is most often used, and almost exclusively employed in large-scale system studies, assumes that the underlying three-phase system is balanced. This allows the use of per-phase models in which the actual three-phase system can be treated as though it were an equivalent single-phase system. This is also known as the positive sequence model. The second type, known as the three-phase power flow, explicitly models the three phases including their mutual couplings, allowing it to handle unbalanced conditions. Presently, the three-phase power flow is mostly used to model distribution systems or microgrids in which significance imbalance could occur, as well as detailed models of distribution circuits where some branches are only one or two phases. If the three-phase power flow were used on large system models, the number of buses would increase by a factor of three and the number of Jacobian elements by a factor of nine. However, getting and using the data needed to set up such models would be a challenge. The remainder of this section refers to the more common balanced, three-phase approach.

As a stand-alone application, the power flow is used in situations ranging from real-time studies in a control center to planning studies looking at system conditions decades in the future. For near-real-time studies the power flow case would be derived from the output of the real-time state estimator (discussed later), which combines

[1] As mentioned in Chapter 1, the dc power flow is a linearized solution method that is used to give solutions to the ac power flow; it has nothing to do with the solution of actual dc systems.

[2] The term case is used to refer to the power system model parameters needed to solve a power flow. At a minimum a case would include a list of the bus parameters (such as their numbers, names, nominal kilovolts), generator and load parameters (including their net power injection values), and transmission line and transformer parameters.

models of the system components (such as the impedances of the transmission lines) with actual system data. In this situation it might be used to determine the impact of an anticipated control action on the real-time system. For longer-term planning studies (e.g., years to decades into the future) the power flow could be used to determine the consequences of proposed system changes, such as new transmission and generation, coupled with changes in the load.

Historically there has been a difference between how the power system is represented in real-time power flows and the longer-term planning power flows. The real-time approach has used what is known as a "node-breaker" full-topology model in which the many actual power system switching devices (such as circuit breakers and disconnects) are represented, with their terminals designated as nodes. Since the switches have essentially zero impedance, when the power flow is solved these nodes are dynamically consolidated, through what is known as topology processing, into a much smaller number of buses. Each bus then corresponds to a set of nodes. So a 100,000-node system might be reduced to perhaps 20,000 buses. This is needed for real-time analysis in which detailed models of the substation topology are available and the status of the switches is known. In the planning context, in which this information is usually not available or may not be fully known for future substations, a more aggregated model is used in which most switches are not explicitly modeled (see Figure 3.1). This difference in modeling assumptions has resulted in a divide between software designed for real-time usage and that for planning applications. However, newer planning software is increasingly able to bridge this divide by directly supporting the node-breaker models.

There are several sources of uncertainty in power flow cases. Usually the reactance and susceptance terms used in the models of the branches are known with good accuracy. There is, however, some uncertainty associated with the resistance term because the resistivity of aluminum and copper has a temperature sensitivity of 0.4 percent per degree Celsius and the actual conductor temperature is seldom known. (The conductor temperature varies in some locations by more than 100°C over the course of a year.) Usually the system solution is relatively insensitive to

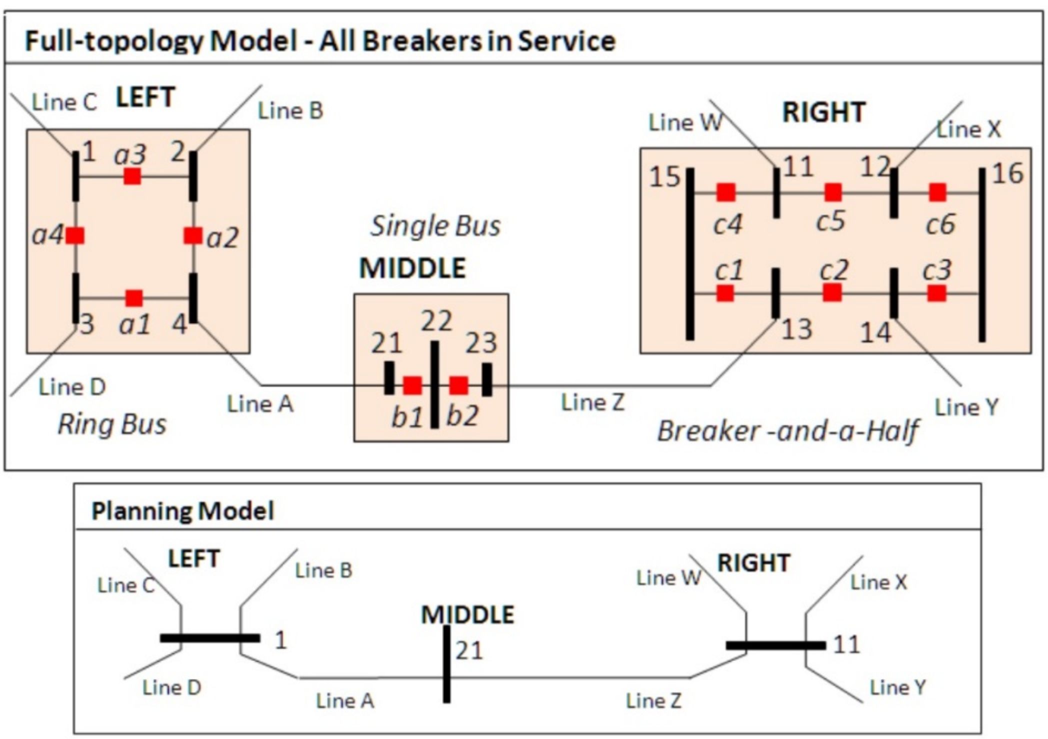

FIGURE 3.1 Node-breaker full-topology models versus bus-branch planning models. SOURCE: Courtesy of PowerWorld Corporation.

resistance errors since the reactances of the high-voltage lines can be many times their resistive values. As resistance heating changes the effective length of the conductor due to thermal expansion effects, the inductance also changes slightly owing to changed geometry. Understanding line-sag effects in real time is a crucial operational issue as excessive sag can lead to conductor-to-tree faults.

The uncertainty in the assumed load and generation depends on the application. In real time these values are provided from the state estimator and would have little error. In day-ahead and longer studies the load uncertainty would depend on the accuracy of the forecast. In longer-term studies there is also uncertainty surrounding which generation and transmission resources would actually be available. The voltage magnitude dependence of the load can also be a source of uncertainty since it ultimately depends on the actual composition of the load, which is continually varying—for example, air-conditioning usage through the course of a summer day or patterns of use for lighting through the day. For an uncontrolled resistive load (such as incandescent lighting) the power varies with the square of the voltage, whereas for electronics there is often essentially no voltage dependence (unless the voltage gets too low). Almost always a constant load power model is assumed.

Another issue with power flow algorithms is assumptions on the response of the embedded power system controls to a change in the system state. These assumptions can cause different software packages to have potentially quite different results; yet each result is valid in that it satisfies all of the specified solution criteria. For example, opening a transmission line would cause a change in the bus voltage magnitudes. A number of continuous and discrete control devices would respond to changes in the bus voltages, including generator reactive power outputs, static var compensators, transformer LTC taps, and automatically switched shunts. For an actual system, how these devices would respond and how they are coordinated depends on their dynamics, as well as on the actions a human operator might take, such as manually changing switched capacitors to keep generator reactive power outputs within the middle of their range to provide reactive reserves. While the basic power flow algorithm is rather simple, much of the sophistication of commercial software packages lies in their handling of these control responses.

The complete set of model parameters necessary to solve the power flow is referred to as a "case." A power flow case would at least have parameters associated with the buses, the generators, the load, and the transmission lines and transformers. Different commercial power flow packages may support different numbers of parameters for the different object types. For example, some packages allow the specification of the latitude and longitude for the buses, and they allow buses to be grouped into substations, while others do not. Cases can represent either entirely fictitious (synthetic) power systems—for example, the seven-bus case from Chapters 1 and 2—or an actual power system.

Within North America, power flow cases are available for all four major interconnects. Until 2001 some of these cases were publicly available through the U.S. Federal Energy Regulatory Commission's FERC715 filings. Subsequent to an October 11, 2001, FERC order, these power flow cases have been treated as confidential for national security reasons, but interested persons may submit a Critical Energy Infrastructure Information request to obtain access. Power flow cases can also sometimes be obtained for research purposes through nondisclosure agreements with electric utilities.

The cases can be interchanged between different software packages using several text file formats, with the formats changing slightly from software version to version. One current impediment to research is that some of the power flow text file formats used in the exchange of power flow data, such as with the FERC715 filings, are proprietary and not fully available to the public. While many researchers in the power discipline have access to these formats since they have purchased the commercial software, this would be less true for researchers in other disciplines such as mathematics. The Institute of Electrical and Electronics Engineers (IEEE) did publish a public format in 1973 (Working Group on a Common Format for Exchange of Solved Load Flow Data, 1973), but it was never updated and is now essentially obsolete. The text file formats permitted for the FERC cases are listed in FERC (2010). Another impediment to research is that there are few publicly available large-scale cases. Since security concerns limit the distribution of actual system cases, an alternative would be the creation of synthetic cases with characteristics like those of the actual system. But even the distribution of synthetic data is limited unless the formats used for their exchange are publicly available—this issue is further discussed in Chapter 8.

Recommendation 2: The Federal Energy Regulatory Commission (FERC) should require that all text file formats used for the exchange of FERC715 power flow cases be fully publicly available.

While power flow analysis is widely used, there are certainly still open research issues. The first is power flow convergence. Being nonlinear, the power flow equations can have multiple solutions or even no solution. Whether a solution is found and, if it is, whether it is the desired solution depends on the initial guess. Commercial practice for large systems is to start a new power flow from the previous solution; few commercial software packages can actually reliably determine a solution if they are not provided with a reasonable initial guess, usually a previous solution. Determining if the desired solution has been found is also a challenge. Complicating convergence is the presence of many additional system automatic controls that must be considered, including discrete controls with either narrow regulation ranges that might not allow for a solution, or wide ranges that allow for a range of solutions. The use of proprietary or specific control algorithms for particular apparatus, as opposed to generic models, creates issues when exchanging cases between different software packages. The increased penetration of inverter-based resources creates needs for rapid development of accurate models and controls and incorporation of these into public models as well as codes. High-voltage dc transmission links, static-var compensators, and now wind farm and solar interconnection are among the many examples of inverter-based resources. As four-quadrant inverters capable of volt/var support become required, this question becomes more and more important. Inverters also cause novel problems in dynamic analyses and harmonic analyses, as discussed below.

A second research issue is dealing with the stochastic nature of the loads and generation. Current practice is to assume a deterministic model, which can then be precisely solved (convergence issues aside). However, even a small amount of uncertainty in some of the loads or generation can result in an almost unmanageable potential range in solutions. This is becoming more of an issue, in particular because of the growth in stochastic renewable generation resources.

STEADY-STATE CONTINGENCY ANALYSIS

An important application of the power flow is known as steady-state contingency analysis (CA), where the adjective steady-state is used to distinguish it from the dynamic contingency analysis discussed later in this chapter in the section "Transient Stability and Longer-Term Dynamics." As touched on in Chapter 1, reliable grid operation requires that the grid be able to operate even with the loss of any single device (e.g., a branch or a generator). This is known as N – 1 reliability. CA refers to the automated process of doing such calculations, in which a set of contingencies is first defined, with each contingency modifying the system in some way, such as the removal of one or more devices. CA then solves each contingency in the set, either sequentially or in parallel, to determine if there are any postcontingency violations.

The earliest automated CA procedures date to the early 1960s (El-Abiad and Stagg, 1963), in which the power flow was just sequentially solved for all the branches in the case. Over the next several decades a number of improvements were introduced, including the recognition that since most contingencies will not cause post-contingency violations they can be processed quite quickly by more approximate contingency screening methods. Since the impact of most contingencies is local, a common screening technique is to apply the contingency to a much smaller equivalent system. The much faster dc power flow algorithm is also used for linear screening. Matrix compensation methods such as the matrix inversion lemma can be used to avoid continually fully factoring the sparse matrix for the incremental changes caused by a single device outage. Bounding algorithms are also used to limit the extent of the solutions required by quickly determining that the final solution will not exceed limits based on initial iterative results.

Online CA has been used in control centers since the 1980s (Subramanian and Wilbur, 1983), running now as often as every minute. For a control center for a large area, many thousands of individual contingencies can be simulated. However, CA is a naturally parallel application because the solution of each contingency is independent of the solutions of the other contingencies. Accordingly, the solution times can be greatly reduced with the use of parallel processing. Sometimes screening techniques are used, and sometimes a full solution is done for each contingency.

As computer speeds have increased, the N – 1 criterion is giving way to what is known as N – 1 – 1. In this approach the initial N – 1 contingencies are solved. During this solution for each N - 1 contingency the postcontingent solution is adjusted using criteria to mimic the actions that would likely be taken immediately after the contingency has occurred (such as generation re-dispatch, phase shifter adjustment, voltage optimization, and load shifting). Then the entire N – 1 set is again applied to each of the originally N – 1 contingencies. Computationally this is of order N^2. The North American Electric Reliability Corporation (NERC) has standards for which contingencies need to be considered (NERC, 2015).

Another newer development is more detailed consideration of contingencies that are likely to initiate cascading failures. Such contingencies are those that would probably cause other devices to fail, resulting in a cascading sequence and, potentially, a large-scale blackout. This is known as N – k analysis, with newer algorithms having been developed to identify such situations.

The CA solution can also require special modeling of automatic actions that would quickly take place following a contingency. Such actions go by different names in different interconnections, including remedial action scheme (RAS), special protection system, and, sometimes, operating procedures. One such action might be to automatically trip a set of generators or lines immediately following a line outage, or to do automatic switching to transfer the supply for a load.

In designing a CA application, special attention must be paid to avoid the "data overwhelm" situation that would occur if a precontingent system already had existing violations. Without such attention, running thousands of contingencies would result in an unmanageable number of contingent violations. New statistical and uncertainty-analysis-based methods need to be developed to provide probabilistic or robust guarantees accounting for uncertainty and fluctuations in loads, renewables, and other components of the system. One issue with CA is that the probability of a given contingency occurring varies widely. Therefore risk-based CA needs to be investigated more thoroughly. NERC is developing criteria for risk-based CA, or stochastic CA, where cases to be studied would be weighted based on some assessment of their relative probability of occurrence and where N – 2 or N – k cases with multiple outages would come into play if the risk assessment so indicated. This is an important area for future research and should be incorporated into the ACOPF research from Recommendation 1.

CA has traditionally been a power flow (steady-state) based problem. Dynamic studies of outages were done manually using transient stability (TS) simulations (see below), which studied if the system could reach a new steady state without physical instability. But for large disturbances, especially ones that cause generators to shift their outputs, there is a growing trend to integrating CA with TS, in which rather than just being solved using power flow the contingencies are solving using TS. This allows determination of the longer-term dynamics (generation response) and whether the system remains secure as it moves to a new steady state.

OPTIMAL POWER FLOW AND SECURITY-CONSTRAINED OPTIMAL POWER FLOW

In its simplest form, the purpose of optimal power flow (OPF) is to optimize an objective function (usually energy costs, but—potentially—losses or other metrics) subject to security and operational limits on voltages and branch flows and subject to the overarching electrical equations as represented in the power flow model. The OPF has been a topic of research and development since the 1950s, when the first digital power flow algorithms were introduced for economic dispatch of the power grid. The problem becomes one of finding the best set of control variables (generation, voltage set points, and so on) such that when a set of these is selected for best objective function and the other injections (load) are input, the power flow solution will be feasible and no constraints violated. It was formulated in the 1960s using a Lagrangian function to minimize generation cost subject to the equality constraints from the power flow and the inequality constraints on both the controllers (e.g., the generator outputs) and other values such as the voltage magnitudes and transmission line flows (Dommel and Tinney, 1968).

Early research utilized the generalized reduced gradient approach, utilizing penalty functions to enforce the binding constraints. The advent of cheap and efficient linear programming (LP) solutions led to LP-based OPF algorithms, with compromises made in the representation of cost functions to accommodate LP formulations.

In the early 1970s the OPF was augmented to include not just constraints associated with violations in the base case power flow but also constraints that could arise because of CA violations (Alsac and Stott, 1974). This is now known as the security-constrained OPF (SCOPF). The goal of the SCOPF is to determine the optimal control settings for the base system, such that the objective function is minimized while simultaneously ensuring there are no violations in the base case or in any of what could be a large number of contingency cases.

While easy to describe, the SCOPF can be quite difficult to fully solve. This is because it is a large-scale, nonlinear optimization that includes a mix of discrete and continuous controls wrapped around the power flow and CA problems. The problem is nonconvex and may admit multiple locally optimal solutions. Today no widely used commercial SCOPF algorithms include methodology to determine if multiple local solutions exist, and most algorithms rely on restrictive formulations to ensure adequate convexity. Additionally, stopping criteria in terms of the change in the objective function as well as changes in the solved state vector can be somewhat ad hoc. A distinction must also be made between the control actions that need to take place precontingency (i.e., need to be applied to the actual solution) and those that would only need to be taken postcontingency. Postcontingency actions, which would include the RASs discussed previously, would only need to occur in the unlikely event the contingency actually occurs. The presence of control algorithms within the solution can lead to discrete variables and discontinuities.

Today's solution approaches take advantage of the several-orders-of-magnitude improvements that have taken place over the last two decades in mixed-integer LP algorithms and in nonlinear programing algorithms. As mentioned in Chapter 2, the SCOPF that is used for determining LMPs in most markets is based on the simplified dc SCOPF that utilizes a dc (as opposed to full ac) power flow model. Typical LMP comparisons between the ac and dc SCOPF methods are given in Overbye et al. (2004). Application of the full ac SCOPF therefore remains an open research issue, with a survey of current approaches covered in Castillo and O'Neill (2013) and more details given in Chapter 6. Future issues around the transmission–distribution interface may well arise if resources on the distribution system are to be included in the SCOPF solution of the transmission grid.

STATE ESTIMATION

In order to work in real time, power flow, CA, and SCOPF require a real-time model of the current system's operating condition. This is provided by what is known as state estimation (SE). Originally formulated in 1970 by Schweppe and Wildes, SE combines a model of fixed system parameters such as the branch impedances, with actual system measurements of quantities such as circuit breaker statuses, bus voltage magnitudes, branch flows, and generator/load injections. The output is the estimated state variables (e.g., the voltage magnitudes and angles), which can then be used to determine the real-time power flow solution that best matches the system measurements. To function, SE requires that the number of measurements be greater than the number of estimated system states (i.e., the system needs to be overdetermined) and that all parts of the system be observable from the measurements.

As normally formulated, SE is a maximum likelihood estimator, where error statistics are assumed for the measurements but no a priori information is assumed about the state variables. The estimated states are those that maximize the likelihood of the set of observations used as inputs. The process is only modeled to the extent of the power flow equations—no process dynamics are modeled. Thus, as opposed to a Kalman filter, power system state estimation is a nonlinear static state estimation problem as formulated. SE is commonly solved using an iterative, weighted-least-squares approach.

If all the measurement errors are assumed to be uncorrelated, the problem is only slightly more difficult than the power flow problem. In the real world, measurements are correlated, so measurement bias is an issue (as opposed to having normally distributed zero mean error, as is commonly assumed). While research has been done on estimating measurement bias, the only enhanced application in use is "bad data detection," where the statistics of the estimated residuals (i.e., the difference between the measured and estimated values) can be used to detect the presence of one or more "bad" data points and then successive estimates without those points are used to validate the bad data identification. This in turn can be used to detect the presence of measurement error, which would also affect basic Supervisory Control and Data Acquisition operation. It is also possible in theory, and as demonstrated in research projects, to use state estimation results over time to estimate network parameters. Practically speaking, this is not widely used because considerable manual intervention is required to pick the parameters to estimate. There are also research-level algorithms to detect and correct topology errors that are due to incorrect switch and breaker status information. However, to date these algorithms have not been widely deployed commercially. These can be formulated to be performed on a substation basis without the need for solution of the ac network equations—Kirchhoff's current law and logic are sufficient.

In the control room, shown in Figure 3.2, the online model provided by the SE is the input to all the other advanced tools (power flow, CA, SCOPF). The best-developed control room state estimators can solve large networks with on the order of 250,000 measurements every minute. However, if for some reason the SE solution does not converge, then an up-to-date model would not be available for these tools. Currently the best large-scale SEs converge well over 98 percent of the time (PJM, 2015). However, convergence alone may not be a sufficient test of the validity of the input assumptions on measurement accuracy and the solution. Rather, a chi-squared test of the measurement residuals (Schweppe and Masiello, 1971) should be applied to see if the solution is "reasonable" given the assumed error statistics.

The deployment of phasor measurement units (synchrophasors, or just PMUs, discussed in Chapter 1), which measure the system voltages and currents (both magnitudes and phase angles) at about 30 times per second, has allowed the development of "direct" state estimation, meaning that explicit, direct noniterative solutions of the power flow equations using phase angles (as measured) as inputs are possible; this is also known as a linear SE. These solutions can be performed very rapidly, at data acquisition scan rates of 10 seconds or less, to provide calculated network conditions as though they were measured. PMU data can also be used to directly estimate the parameters of a branch between two PMU measurement points (in conjunction with measured branch flows)—something that is useful in detecting potential transmission line sag as a function of line loading and ambient conditions.

FIGURE 3.2 Power system tools in use in the PJM control center. SOURCE: Courtesy of PJM Interconnection.

The use of SE on the distribution system is not a well-developed methodology, nor is it in widespread use as typically there are not sufficient redundant measurements to warrant its application.

Including control algorithms for locally autonomous apparatus (as with the power flow problem) is not a well-developed methodology in SE either, and the advent of more and more inverter-based resources capable of local volt/var control will likely make this a more important question.

TRANSIENT STABILITY AND LONGER-TERM DYNAMICS

The previous techniques (power flow, CA, SCOPF, and SE) are focused on determining characteristics associated with power system quasi-steady-state equilibrium points. Hence the results for each contingency in CA determine characteristics of a potential new equilibrium point but do not tell whether the power system, which includes numerous dynamics, will actually be able to reach that equilibrium point. This is determined by TS, which is concerned with the power system dynamic response for time frames from about 0.01 sec to a few dozen seconds or more.

When a contingency occurs, such as a fault on a transmission line or the loss of a generator, the system experiences a "jolt" that results in a mismatch between the mechanical power delivered by the generators and the electric power consumed by the load. The phase angles of the generators relative to one another change owing to power imbalance. If the contingency is sufficiently large it can result in generators losing synchronism with the rest of the system, or in the protection system responding by removing other devices from service, perhaps starting a cascading blackout.

Stability issues have been a part of the power grid since its inception, with Edison having had to deal with hunting oscillations[3] on his steam turbines in 1882, when he first connected them in parallel (Hughes, 1983). Digital computer simulations date to the late 1950s (Dyrkacz et al., 1960), with a wide variety of different solution techniques presented in the literature by the 1970s (Stott, 1979). As described in Chapter 1, TS involves the time domain simulation of differential algebraic equations (DAEs). Initially the differential equations were used to represent just the electromechanical effects on the synchronous machines and their prime movers, including the excitation systems and the prime mover governor controls and excitation controls. They have been expanded to include a host of other devices, such as those simulating the frequency response characteristics of the large inertial loads on the system (e.g., pumping motors) as well as wind turbine generators. The dynamic simulations are interleaved with the algebraic power balance equations representing the network power balance constraints (similar to the power flow equations). Hence TS studies require the data from a power flow, augmented by a representation of the system dynamics.

The DAEs can be solved using either explicit or implicit methods. Several widely used commercial TS packages use the explicit approach; others use an implicit approach. While numerical instability can be an issue with explicit methods, in practice it is seldom a concern. This is partially due to the use in commercial packages of multirate methods, in which different variables are integrated with different time steps. Such an approach uses smaller time steps for fast varying variables and larger time steps for the slowly varying ones. Time steps of half or quarter cycle are common in TS studies.

Within North America, TS cases are available for all four major interconnects.[4] Large-scale studies done for the Western Interconnection and the Electric Reliability Corporation of Texas interconnection usually include a representation of the entire interconnect, whereas for the large Eastern Interconnection (which includes the Quebec Interconnection) an equivalenced representation is often used. Typical system sizes might be up to 20,000 buses and more than 100,000 state variables. As an example, Figure 3.3 shows the bus frequencies at six locations in a large system 2 sec after a large generator outage contingency.

A TS study can be similar to CA in that a variety of different contingencies might be considered. Also like CA, such a TS study is naturally parallel since the solution of each contingency is independent of the others. On

[3] "Hunting oscillation" is a self-oscillation about an equilibrium describing how a system "hunts" for equilibrium.

[4] As with power flow, "case" is used here to refer to the power system model parameters needed to solve a transient stability. Usually the transient stability case information supplements the information provided by the power flow case, with both needed to do a transient stability study.

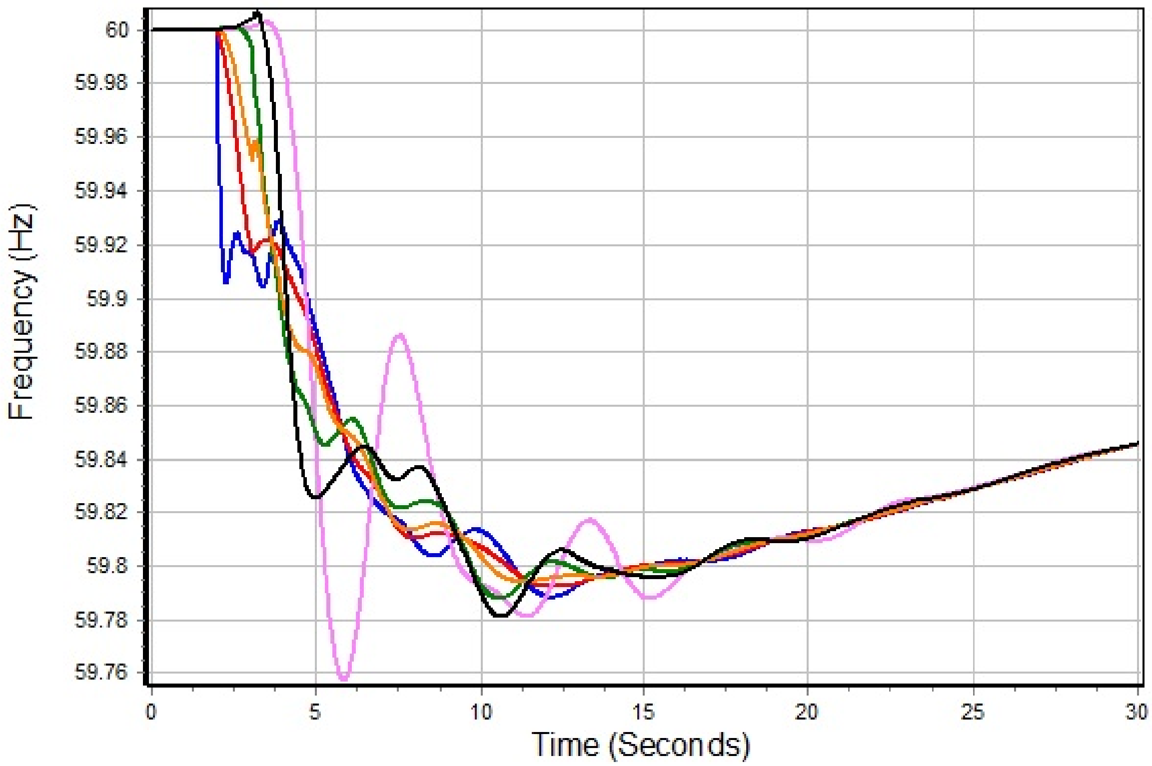

FIGURE 3.3 Bus frequencies at six locations in a large interconnection after a generator outage contingency. SOURCE: Overbye et al. (2012). Courtesy of Thomas J. Overbye, University of Illinois, Urbana-Champaign.

a reasonably fast PC, a single TS solution for a single contingency on a 20,000-bus system might run an order of magnitude slower than in real time.

Transient stability studies are routinely performed for planning purposes and for next-day operational studies. Their use in real-time operations, in which the starting point would be the SE power flow case augmented with system dynamics, is in its infancy. The increasing penetration of inverter-based resources leads to concerns among some grid operators about managing inertial response and TS better operationally, so there is a need to increase such near-real-time TS assessment. While a great deal of research has been done to explore the use of direct methods for assessing stability (such as with Lyapunov functions), to date there has not been a successful formulation that can be used for realistic problems; Lyapunov-based techniques are sometimes used for screening to determine contingencies that are likely to have problems.

While TS data are considered more confidential than power flow data, cases can sometimes be obtained for research purposes through nondisclosure agreements. The cases can be interchanged between different software packages using several text file formats, with the formats changing slightly from software version to version. Commercial packages support on the order of several hundred different dynamic models. As an example, Figure 3.4 shows the block diagram for a nine-state hydro governor model. Models have been growing in complexity, particularly those used to represent the load. For example, now one load model requires over 100 parameters. However, there is a continued need for improved dynamic load models. This is an area in which new data-driven models based on machine learning could play a significant role.

As was the case with the power flow, one difficulty potentially impeding research is that some of the transient stability models used in systemwide studies are not publicly available. This includes proprietary user-defined models distributed only in a machine-readable format. Hence these models cannot be evaluated by the broad

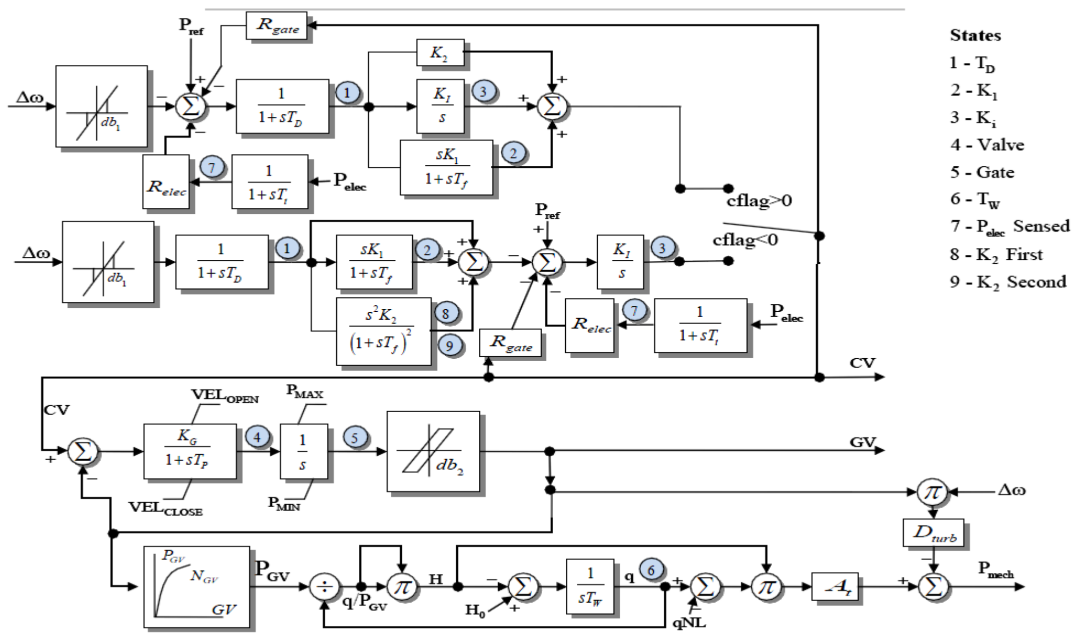

FIGURE 3.4 Block diagram for a hydro governor. SOURCE: Overbye et al. (2012). Courtesy of PowerWorld Corporation.

power community including researchers. While the text file formats and the widely used models can be obtained by purchasing the commercial packages, the use of proprietary models and text file formats can limit access to researchers outside the power community. An example of good public availability of widely used hydro turbine-governor models and their text file descriptions for hydro generators is Koritarov et al. (2013).

Recommendation 3: The Federal Energy Regulatory Commission should require that descriptions of all models used in systemwide transient stability studies be fully public, including descriptions of any associated text file formats.

Another issue impeding research is that there are few publicly available large-scale cases. Since security concerns limit the distribution of actual system cases, an alternative would be the creation of synthetic cases with characteristics like those of the actual system; as with the power flow synthetic cases, this topic is covered more in Chapter 8.

Over the years, little work had been published on comparing the results of TS studies on the same system using different commercial packages. Such comparisons have been hindered in part because the different packages can use slightly different models. Recent work has demonstrated that quite close results could be obtained on an 18,000-bus case using two common packages (Shetye et al., 2016). However, certainly more work is needed on expanding the existing comparisons to other cases which would include different model parameters and different models, and on expanding the comparisons to other widely used packages.

Validation of the simulated results with respect to the actual system is also an important area for additional research. The models and their parameters are validated for the individual generators by disconnecting the generator from the system and then subjecting the generator to various tests to exercise its dynamics. Such testing is obviously not possible for the system itself. Rather there is a need to utilize results from the periodic disturbances that occur on the grid. Whole-system validation has also been hindered both by the lack of fast real-time measurements and the absence of dynamic models integrated with SE results. However, this is now beginning to change, driven in part by the availability of PMU results. Load model validation differs from whole-system validation since the composition of the load is continually changing.

The advent of inverter-based resources, especially wind farms and photovoltaic farms, has greatly complicated the TS problem. The dynamics of inverter response to changed frequency and voltage can become important, especially as interconnection standards for low voltage and fault ride-through (the ability of a device to operate and remain connected to the grid while the fault is cleared) are more and more prevalent. In the case of wind farms, the dynamics of the turbine and turbine controls behind the inverter are also important. Because these technologies are developing rapidly and in some cases are manufacturers' proprietary models, industry standard models with sufficient fidelity for TS lag behind the real-world developments (Yaramasu et al., 2015). The development of inverter-based synthetic inertia and synthetic governor response from wind farms, photovoltaic farms, and grid-connected storage systems will create additional modeling complexity.

The simulation of the system including time dynamics over longer periods of time than TS has been called mid-term stability simulation, and long-term stability simulation as the time period is extended (Kundur, 1994). Dynamic stability (DS) simulations typically simulate the generator prime mover (steam firing and steam turbine, governor, and controls) but not the electromechanical effects. In effect, DS simulations consider the TS problem with the assumption of a common single-frequency throughout, solutions over longer time periods (from hours up to a day), and with longer-term dynamics (such as prime mover fuel/combustion side effects) modeled. DS simulations can embed a power flow for a network solution or can rely on simpler representations to compute intercontrol-area interchange flows without the intra-area network representation. DS simulations are typically the basis of Operator Training Simulator real-time simulations for power system dynamics (Latimer and Masiello, 1978; Podmore et al., 1982; Prais et al., 1989) with TS solutions introduced when switching events occur.

DS solutions have become more important in recent years as a result of the increased use of renewable sources, which causes concerns about system dynamic performance in terms of frequency and area control error—control area dynamic performance. DS solutions typically rely on IEEE standard models for generator dynamics and simpler models for assumed load dynamics. As with TS solutions, providing accurate models for wind farm dynamics and for proposed synthetic inertial response and governor response is a challenge. DS solutions have also been used to investigate algorithms for incorporating fast storage into control area automatic generation control (AGC) and similar questions (Masiello and Katzenstein, 2012). Another recent trend is the incorporation of the longer-term DS dynamics into TS packages, blurring or eliminating the differences between the two.

SHORT-CIRCUIT ANALYSIS

Short-circuit calculations are a calculation of the short-circuit current and impedance visible when a fault to ground or from phase to phase is introduced in the network. Faults can be introduced at a node or along a branch. Short circuits can be single, two, or three phase, and can be phase to phase, so three-phase network models are normally utilized. While the primary purpose is to calculate the worst-case currents that will occur during a fault, it is also important to know the phase voltages during the fault. Short-circuit calculations are used in sizing circuit breakers, in analyzing the short-circuit duties of apparatus (especially transformers and generators), and in setting protection. A short-circuit calculation is a specialized load flow with a zero impedance to ground or from phase to phase inserted at the hypothesized fault location. Short-circuit analysis is generally used in transmission planning for design and protection setting; it is rarely used online, as the operating assumption is that the system is built and protected to be safe against short-circuit conditions.

ELECTROMAGNETIC TRANSIENTS

The transmission lines can be represented as a form of waveguide or transmission pipe for the purpose of assessing wave propagation down the line. At every change of impedance along the line in a network, reflections are created. As discussed in Chapter 1, switching transients, faults, and, especially, lightning strikes cause waves of voltage change to propagate along the network at near the speed of light in the medium. These step-function waves deteriorate over time and distance due to losses, but the reflections can combine and in some cases produce voltage transients of double the nominal line voltage or more. Lightning protection such as arresters is designed to constrain the maximum overvoltage that can occur. The first digital solutions for solving electromagnetic transients date from the late 1960s with the introduction of a technique using trapezoidal integration coupled with sparsity techniques (Dommel, 1969). Electromagnetic transient analysis is used in power system engineering and planning applications.

The advent of high penetrations of inverter-based renewable generation (wind farms, solar farms) has led to a requirement for interconnection studies for each new renewable resource to ensure that the new wind farm will not create problems for the transmission system. These interconnection studies begin with load-flow analyses to ensure that the transmission system can accommodate the increased local generation, but then broaden to address issues specific to inverter-based generation, such as analyzing harmonic content and its impact on the balanced three-phase system.

HARMONIC ANALYSIS

The models described in all sections of this report are based on the 60-Hz waveform and the assumption that the waveform is "perfect," meaning that there are no higher-order harmonics caused by nonlinearities, switching, imperfect machines and transformers, and so on. However, inverters are switching a dc voltage at high frequencies to approximate a sine wave, and this inevitably introduces third, fifth, and higher-order harmonics or non-sine waveforms into the system. The increased use of renewables and also increased inverter-based loads make harmonic analysis—study of the behavior of the higher harmonics—more and more important. While interconnection standards tightly limit the harmonic content that individual inverters may introduce into the system, the presence of multiple inverter-based resources in close proximity (as with a new transmission line to a region having many wind farms) can cause interference effects among the multiple harmonic sources.

GENERATION ANALYTICS

For more than 50 years, the problem of balancing generation to load has been addressed with a series of control and decision-support tools operating at different time frames. The primary response is the autonomous operation of generator governor control in response to frequency deviation from 60 Hz so as to control frequency deviation in a coordinated way. Governor "droop" on a uniform basis across many generators ensures that each governor responds to frequency changes proportional to its size. Secondary control operating at a 2- or 4-second interval responds to the net frequency change (residual of governor action) to restore frequency to nominal. (This is called load frequency control, or LFC.) In interconnected systems, which is the norm everywhere except for systems that constitute "electrical islands," the deviation of inter-tie flows from scheduled flows is adjusted in a coordinated way using tie-line-bias control, which uses the natural aggregate frequency response of each control area in conjunction with the tie flow deviation to allow each control area to adjust generation to meet its own load and restore frequency. The principles of tie line bias control have not changed since the 1950s, and NERC standards today govern the operation of AGC. The actual control algorithm in use in almost all control centers is a proportional integral derivative controller with more or less sophisticated logic for allocating the control signal to the generators participating in secondary control (called "regulation" in most markets/control areas). Model predictive control (MPC) has been developed extensively in the literature for the AGC problem but has rarely been applied in the field. The minor improvements in the system (which are not required by NERC standards today) do not justify the increased cost and complexity of the software and the models needed. However, high penetration by renewables, decreased conventional generation available for regulation, the advent of new technologies such as fast short-term storage (flywheels, batteries), and short-term renewable production forecasting may reopen the investigation of MPC for AGC (Masiello and Katzenstein, 2012).

Tertiary control, or real-time dispatch, occurs at slower intervals—5 minutes in most market systems, or on demand as total load changes in vertically integrated utility operations. Tertiary control, historically called economic dispatch, reallocates the total generation among online units so as to minimize production cost. This occurs when all the unconstrained generators have the same incremental cost, with this value referred to as the system lambda and when the total generation is equal to the total load plus losses. Originally economic dispatch was solved on special analog computers (Kirchmayer, 1959). Refinements on "lambda dispatch" added a second Lagrange multiplier μ as the cost of an aggregate constraint such as total system reserve (Stadlin, 1971). This economic dispatch paradigm is still in widespread use in vertically integrated or small control areas.

In a market environment hour-ahead bids/offers for incremental/decremental generation are used as a proxy for the unit incremental cost curves used in economic dispatch. Increasingly, market operators are using a variation of the mixed integer programming scheduling solution to solve the 5-min dispatch problem as an integrated solution of current state and near-term forecast (5, 10, 15, . . . , 60 min ahead) conditions as "trajectories" for optimal dispatch. These solutions can accommodate quick-start units that can be started in near real time, short-term loads, and renewable production forecasts.

It has become apparent in recent work on the United Kingdom's national grid that there are trade-offs among the three control domains—primary, secondary, and tertiary—in that altered performance in one causes altered requirements in the others. For example, less primary response entails more secondary response; better performing secondary response can mitigate the efforts required in the tertiary controls; and better forecasting in the "trajectory" solution will keep dispatch closer to load and place less demand on secondary response. What has not been done as yet is to unify the mathematics for these three "products" so as to enable rigorous analysis of the best portfolio for cost and reliability, as opposed to standards-based determination, especially of the first two.

The hour-ahead and day-ahead schedules, as well as simulations of annual production costing on an hourly basis, can be lumped under the domain of security-constrained unit commitment (SCUC). Dynamic programming (Larson, 1967) was first used to perform the unit commitment analysis (identifying which generators should be running at what level each hour) in the 1970s. Today, commercial SCUC solutions have migrated to the mixed integer programming (MIP) formulation for vastly increased computational performance (as described in Chapter 4). MIP also allows more flexible modeling of complex unit behavior such as multistate combined cycle plants (which have more than one combustion turbine and more than one steam turbine capable of operating in different configurations). When integrated with an OPF (usually a semilinearized or "dc" network model) the SCUC can produce nodal prices. Applications of MIP for unit scheduling include market operations; generator operator simulation of markets for bidding support; annual production costing for studying future generation portfolios; renewable penetrations; impacts of transmission planning; and generation interconnection studies, including the probability of wind curtailment for transmission constraints. Even as MIP algorithms have enabled larger and larger networks and generation fleets to be studied, the industry appetite grows faster. To perform one interconnection study, it is typical to simulate annual production cost for an entire interconnection, which can take several hours to perform.

One noteworthy point is that the commercial SCUC software tools in use today will formulate the problem in proprietary code/databases but then interface to third-party MIP engines using industry standard integration layers. The YALMIP suite is an example of this. Users can then select the MIP engine (CPLEX, for example) that best suits their problem at hand. Developers of the MIP engines focus on the performance and robustness of their particular algorithm and code.

Dynamic programming (DP) is still the algorithm of choice for generation scheduling where energy levels are a state variable linking the solution at each time step. This is typical in hydrothermal (H-T) coordination. H-T coordination has become less of an issue in the United States, where market regimes eliminate vertical decision making, but is very much an issue in other regions (Brazil, for example). An interesting question today is whether the integration of large numbers of independent storage resources with their markets will tax MIP engines and cause a revisit of the DP versus MIP decision or the development of new algorithms to adapt to large numbers of storage resources. Current market rules that force storage to participate on the same basis as generators make this question moot, but market rules may evolve to raise the issue. Conversely, the issue may itself limit the evolution of market rules around storage. Incorporating stochastic characterizations of renewable production into the SCUC formulation can lead to brute force Monte Carlo simulations or to stochastic DP formulations (LLNL, 2014). With the possible exception of long-term hydrothermal coordination codes, these are still being studied and no solutions are commercially available today.

In planning studies, generation capacity and contingency analysis studies have been focused on a probabilistic analysis of the likelihood of insufficient capacity at a given moment owing to multiple unit outages. This loss-of-load probability (LOLP) (Billington, 1996) has become a reliability standard in long-term planning. For capacity planning today, renewable resources and demand response resources are assigned a capacity factor or de-rating for use in capacity adequacy studies and LOLP calculations. There may be a need to consider how these capacity factors can be made stochastic and integrated into the LOLP along with stochastic generator outage statistics.

MODELING HIGH-IMPACT, LOW-FREQUENCY EVENTS

An emerging area for which some analytic tools and methods are now becoming available is the modeling of what are often referred to as high-impact, low-frequency (HILF) events (NERC, 2010)—that is, events that are statistically unlikely but still plausible and, if they were to occur, could have catastrophic consequences. These include large-scale cyber or physical attacks, pandemics, electromagnetic pulses (EMPs), and geomagnetic disturbances (GMDs). This section focuses on GMDs since over the last several years there has been intense effort in North America to develop standards for assessing the impact of GMDs on the grid. Associated with the effort has been the emergence of commercial tools to help utilities carry out such assessments.

GMDs, which are caused by coronal mass ejections from the Sun, can impact the power grid by causing low-frequency (less than 0.1 Hz) changes in Earth's magnetic field. These magnetic field changes then cause quasi-dc electric fields, which in turn cause what are known as geomagnetically induced currents (GICs) to flow in the high-voltage transmission system. The GICs impact the grid by causing saturation in the high-voltage transformers, leading to potentially large harmonics, which in turn result in both greater reactive power consumption and increased heating. It has been known since the 1940s that GMDs have the potential to impact the power grid; a key paper in the early 1980s showed how GMD impacts could be modeled in the power flow (Alberston et al., 1981). The two key concerns associated with large GMDs are that (1) the increased reactive power consumption could result in a large-scale blackout and (2) the increased heating could permanently damage a large number of hard-to-replace high-voltage transformers (NERC, 2012).

The magnitudes of GMDs are expressed in nanotesla (nT)-per-minute variations in Earth's magnetic field.[5] Large GMDs are quite rare but could have catastrophic impact. For example, a 500 nT/min storm blacked out Quebec in 1989. Larger storms, with values of up to 5,000 nT/min, occurred in 1859 and 1921, both before the existence of large-scale grids. Since such GMDs can be continental in size, their impact on the grid could be significant, and tools are therefore needed to predict them and to allow utilities to develop mitigation methods.

As a result of the recent effort led by NERC over the last 3 years, GMD assessment has been integrated into several commercial power analysis tools (Overbye et al., 2012). For example, GICs can be calculated for assumed uniform or nonuniform electric field variations and simultaneously their transformer impact integrated into the power flow calculations. Figure 3.5 shows the GICs calculated for an assumed uniform 2 V/km electric field over the Eastern Interconnection. More recently such calculations have also been integrated into the TS calculations, paving the way for the modeling of the much larger but shorter-time-frame GICs that could be caused by an EMP. While good progress has been made, the power system modeling of HILFs has only just begun.

[5] A tesla (T) is a unit of magnetic induction equal to one weber per square meter; nT abbreviates nanotesla.

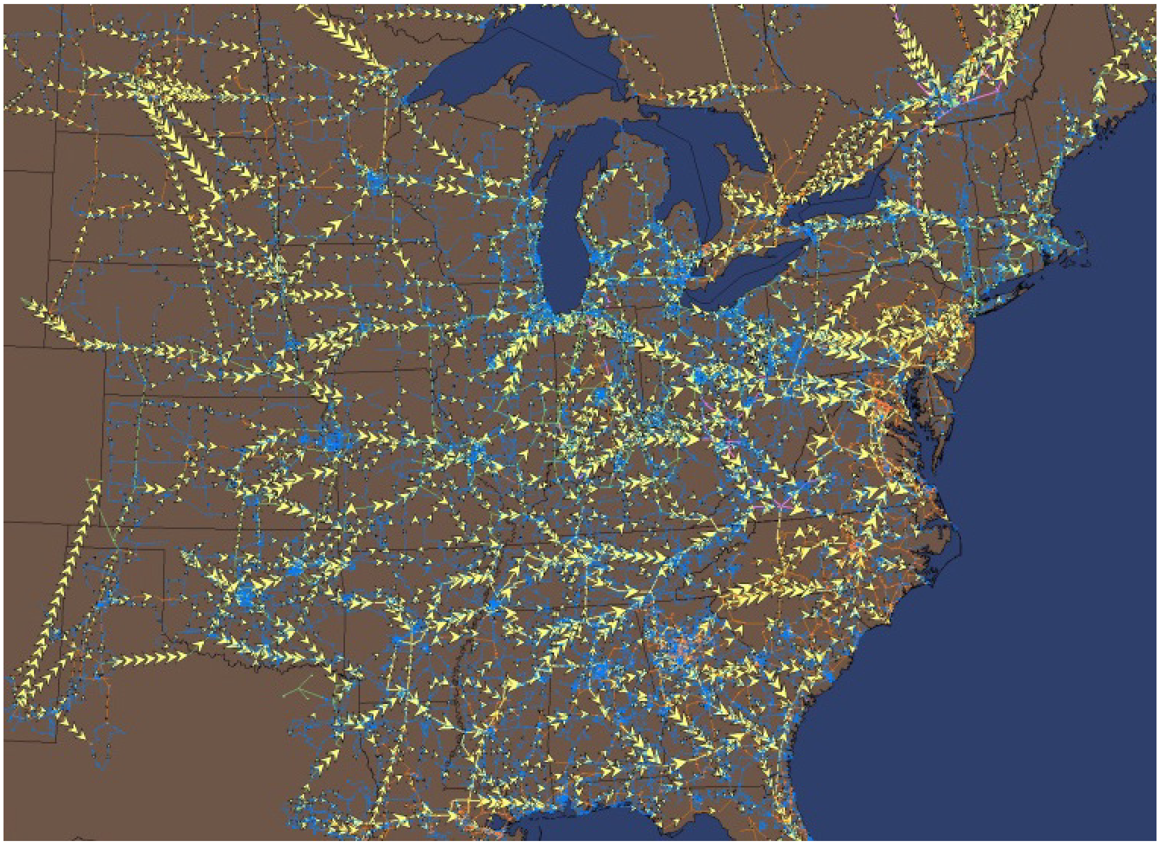

FIGURE 3.5 Visualization with yellow arrows showing the geomagnetically induced currents (GICs) in the Eastern Interconnection for a uniform 2 V/km eastward electric field. SOURCE: Overbye et al. (2012). Courtesy of Thomas J. Overbye, University of Illinois, Urbana-Champaign.

REFERENCES

Albertson, V.D., J.G. Kappenman, N. Mohan, and G.A. Skarbakka. 1981. Load-flow studies in the presence of geomagnetically-induced currents. *IEEE Transactions on Power Apparatus and Systems* 100(2):594-607.

Alsac, O., and B. Stott. 1974. Optimal load flow with steady-state security. *IEEE Transactions on Power Apparatus and Systems* PAS-93(3):745-751.

Billington, R.A. 1996. *Reliability Evaluation of Power Systems*, 2nd ed. Springer, New York, N.Y.

Brown, R.J., and W.F. Tinney. 1957. Digital solutions for large power networks. *Transactions of the American Institute of Electrical Engineers, Part III: Power Apparatus and Systems* 76: 347-351.

Castillo, A., and R.P. O'Neill. 2013. "Survey of Approaches to Solving the ACOPF." U.S. Federal Energy Regulatory Commission. March. http://www.ferc.gov/industries/electric/indus-act/market-planning/opf-papers/acopf-4-solution-techniques-survey.pdf.

DeMarco, C.L., and T.J. Overbye. 1988. Low voltage power flow solutions and their role in exit time bases security measure for voltage collapse. Pp. 2127-2131 in *Proceedings of the 27th IEEE Conference on Decision and Control*. doi:10.1109/CDC.1988.194709.

Dommel, H.W. 1969. Digital computer solution of electromagnetic transients in single- and multiphase networks. *IEEE Transactions on Power Apparatus and Systems* PAS-88(4):388-399.

Dommel, H.W., and W.F. Tinney. 1968. Optimal power flow solutions. *IEEE Transactions on Power Apparatus and Systems* PAS-87(10):1866-1876.

Dyrkacz, M.S., C.C. Young, and F.J. Maginniss. 1960. A digital transient stability program including the effects of regulator, exciter and governor response. *Transactions of the American Institute of Electrical Engineers, Part III: Power Apparatus and Systems* 79(3):1245-1254.

El-Abiad, A.H., and G.W. Stagg. 1962. Automatic evaluation of power system performance—Effects of line and transformer outages. *Transactions of the American Institute of Electrical Engineers, Part III: Power Apparatus and Systems* 81(3):712-715.

FERC (Federal Energy Regulatory Commission). 2010. Form No. 715 Instructions. January. http://www.ferc.gov/docs-filing/forms/form-715/instructions.asp#Part 2.

Hughes, T.P. 1983. *Networks of Power*. John Hopkins University Press, Baltimore, Md. p. 43.

Iwamoto, S., and Y. Tamura. 1981. A load flow calculation for ill-conditioned power systems. *IEEE Transactions on Power Apparatus and Systems* PAS-100(4):1736-1743.

Kirchmayer, L.K. 1959. Differential analyzer aids design of electric utility automatic dispatching system. *Transactions of the American Institute of Electrical Engineers, Part II: Applications and Industry* 77(6):572-579.

Koritarov, V., L. Guzowski, J. Feltes, Y. Kazachkov, B. Lam, C. Grande-Moran, G. Thomann, L. Eng, B. Trouille, and P. Donalek. 2013. *Review of Existing Hydroelectric Turbine-Governor Simulation Models*. Argonne National Laboratory. ANL/DIS-13/05. August. http://ceeesa.es.anl.gov/projects/psh/ANL_DIS-3_05_Review_of_Existing_Hydro_and_PSH_Models.pdf.

Kundur, P. 1994. *Power System Stability and Control*. McGraw-Hill Education, New York, N.Y.

Larson, R.E. 1967. A survey of dynamic programming computational procedures. *IEEE Transactions on Automatic Control* 12(6):767-774.

Latimer, J.R., and R. Masiello. 1978. Design of a dispatcher training system. Abstract, 1977 Power Industry Computer Applications Conference. *IEEE Transactions on Power Apparatus and Systems* PAS-97(2):315.

LANL (Los Alamos National Laboratory). 2014. "Automated Demand Response and Storage for Renewable Integration." California Energy Commission Public Workshop. June 16. Unpublished.

Masiello, R.D., and W.P. Katzenstein. 2012. "Adapting AGC to Manage High Renewable Resource Penetrations." *Power and Energy Society General Meeting*. IEEE. pp. 1-6. doi:10.1109/Pesgm.2012.6345486.

NERC (North American Electric Reliability Corporation). 2010. *High-Impact, Low-Frequency Event Risk to the North American Bulk Power System*. June. http://www.nerc.com/pa/CI/Resources/Documents/HILF%20Report.pdf.

NERC. 2012. *2012 Special Reliability Assessment Interim Report: Effects of Geomagnetic Disturbances on the Bulk Power System*. February. http://www.nerc.com/files/2012GMD.pdf.

NERC. 2015. "Standard TLP-001-4—Transmission System Planning Performance Requirements." http://www.nerc.com/files/TPL-001-4.pdf. Accessed September 15, 2015.

Overbye, T.J., X. Cheng, and Y. Sun. 2004. "A Comparison of the ac and dc Power Flow Models for LMP Calculations." *Proceedings of the 37th Hawaiian International Conference on System Sciences*. https://www.computer.org/csdl/proceedings/hicss/2004/2056/02/205620047a.pdf.

Overbye, T.J., T.R. Hutchins, K.S. Shetye, J. Weber, and S. Dahman. 2012. "Integration of Geomagnetic Disturbance Modeling into the Power Flow: A Methodology for Large-Scale System Studies." *Proceedings of the 2012 North American Power Symposium*. file:///C:/Users/cgruber/Downloads/Overbye_NAPS_GIC_2012_June1-2012.pdf.

PJM Operations Support Division. 2015. *Energy Management System (EMS) Model Updates and Quality Assurance (QA)*. Manual M-03A. January. http://www.pjm.com/documents/manuals.aspx.

Podmore, R., J.C. Giri, M.P. Gorenberg, J.P. Britton, and N.M. Peterson. 1982. An advanced dispatcher training simulator. *IEEE Transactions on Power Apparatus and Systems* PAS-101:17-25.

Prais, M., G. Zhang, Y. Chen, A. Bose, and D. Curtice. 1989. Operator training simulator algorithms and test result. *IEEE Transactions on Power Systems* 4(3).

Sato, N., and W.F. Tinney. 1970. Techniques for exploiting the sparsity of the network admittance matrix. *IEEE Transactions on Power Apparatus and Systems* PAS-89:120-125.

Schweppe, F.C., and R.D. Masiello. 1971. A tracking static state estimator. *IEEE Transactions on Power Apparatus and Systems* 90(3):1025-1033.

Schweppe, F.C., and J. Wildes. 1974. Power system static-state estimation, Part 1: Exact model. *IEEE Transactions on Power Apparatus and Systems* PAS-93:859-869.

Shetye, K.S., T.J. Overbye, S. Mohapatra, Ran Xu, J.F. Gronquist, and T.L. Doern. 2016. Systematic determination of discrepancies across transient stability packages. *IEEE Transactions on Power Systems* 31(1):432-411.

Stadlin, W.O. 1971. Economic allocation of regulating margin. *IEEE Transactions on Power Apparatus and Systems* PAS-90(4):1776-1781.

Stott, B. 1979. Power system dynamic response calculations. *Proceedings of the IEEE* 67(2):219-241.

Stott, B., and O. Alsac. 1974. Fast decoupled load flow. *IEEE Transactions on Power Apparatus and Systems* PAS-93(3):859-869.

Subramanian, A.K., and J. Wilbur. 1983. Power system security function as of the energy control center at the Orange and Rockland utilities. *IEEE Transactions on Power Apparatus and Systems* PAS-102(12):3825-3834.

Thorp, J.S., and S.A. Naqavi. 1989. "Load Flow Fractals." *Proceedings of the 28th IEEE Conference on Decision and Control*. doi:10.1109/CDC.1989.70472.

Tinney, W.F., and C.E. Hart. 1967. Power flow solution by Newton's method. *IEEE Transactions on Power Apparatus and Systems* PAS-86(11):1449-1460.

Working Group on a Common Format for Exchange of Solved Load Flow Data. 1973. Common format for exchange of solved load flow data. *IEEE Transactions on Power Apparatus and Systems* PAS-92(6):1916-1925.

Yaramasu, V., B. Wu, P. Sen, S. Kouro, and M. Naraimani. 2015. High-power wind energy conversion systems: State-of-the-art and emerging technologies. *Proceedings of the IEEE* 103(5):740-788.

4

Background: Mathematical Research Areas Important for the Grid

INTRODUCTION

Building on the electric grid basics presented in Chapters 1 and 2 and the existing analytic methods from Chapter 3, this chapter covers the key mathematical research areas associated with the electric grid. As was the case in the previous chapters, the scope of the electric power industry and its wide variety of challenges mean that only the key mathematical techniques can be touched on.

The mathematical sciences provide essential technology for the design and operation of the power grid. Viewed as an enormous electrical network, the grid's purpose is to deliver electrical energy from producers to consumers. The physical laws of electricity yield systems of differential equations that describe the time-varying currents and voltages within the system. As described in Chapter 1, the North American grid is operated in regimes that maintain the system close to a balanced three-phase, 60-Hz ideal. Conservation of energy is a fundamental constraint: Loads and generation must always balance. This balance is maintained in today's network primarily by adjusting generation. Generators are switched on and off while their output is regulated continuously to match power demand. Additional constraints come from the limited capacity of transmission lines to deliver power from one location to another.

The character, size, and scope of power flow equations are daunting, but (approximate) solutions must be found to maintain network reliability. From a mathematical perspective, the design and operation of the grid is a two-step process. The first step is to design the system so that it will operate reliably. Here, differential equations models are formulated, numerical methods are used for solving them, and geometric methods are used for interpreting the solutions. The next section, "Dynamical Systems," briefly introduces dynamical systems theory, a branch of mathematics that guides this geometric analysis. Stability is essential, and much of the engineering of the system is directed at ensuring stability and reliability in the face of fluctuating loads, equipment failures, and changing weather conditions. For example, lightning strikes create large, unavoidable disturbances with the potential to abruptly move the system state outside its desired operating regime and to permanently damage parts of the system. Control theory, introduced in a later section, "Control," is a field that develops devices and algorithms to ensure stability of a system using feedback. In that section the committee describes some of the basic types of control that are currently used on the grid.

More generation capacity is needed than is required to meet demand, for two reasons: (1) loads fluctuate and can be difficult to accurately predict and (2) the network should be robust in the face of failures of network components. The organizations and processes used to regulate which generation sources will be used at any given time were covered in earlier chapters. Key to this operation and its associated design is extensive use of optimization algorithms. The next section, "Optimization," describes some of the mathematics and computational methods for optimization that are key aspects of this process. Because these algorithms sit at the center of wholesale electricity markets, they influence financial transactions of hundreds of millions of dollars daily.

The electrical grid operates 24/7, but its physical equipment has a finite lifetime and occasionally fails. Although occasional outages in electric service are expected, an industry goal is to minimize these and limit their extent. Cascading failures that produce widespread blackouts are disruptive and costly. Systematic approaches to risk analysis, described in the section "Risk Analysis, Reliability, Machine Learning, and Statistics," augment physical monitoring devices to anticipate where failures are likely and to estimate the value of preventive maintenance.

The American Recovery and Reinvestment Act of 2009 funded the construction and deployment of many of the phasor measurement units (PMUs) discussed in Chapter 1, so that by 2015 there are approximately 2,000 production-grade PMUs just in North America that are sampling the grid 30 to 60 times per second (NASPI, 2015). This is producing an unprecedented stream of data, reporting currents and voltages across the power system with far greater temporal resolution (once every 4 to 6 seconds) than was available previously from the existing Supervisory Control and Data Acquisition (SCADA) systems. The subsections in "Complexity and Model Reduction in the Time of Big Data" discuss evolving areas of mathematics that seem likely to contribute to effective utilization of these new data. The last subsection discusses data assimilation, which has become an important tool for weather forecasting and may have application to the grid. Data assimilation enables the ongoing aggregation of weather observations into large numerical models that simulate the global evolution of the atmosphere to produce weather forecasts. Initialized with typical observational data, these computer models have very fast transients that in the real atmosphere have already decayed. One of the primary goals of data assimilation is to avoid these transients while still using the observational data. Data assimilation has yet to be tried on simulations of the electric grid, but the prospect of doing so is attractive. Methods for determining the extent to which a system with a large number of degrees of freedom behaves like a system with many fewer degrees of freedom are also discussed in the section "Complexity and Model Reduction in the Time of Big Data." The prospect of applying these methods to the data from the PMUs is also very attractive.

The final section, "Uncertainty Quantification," introduces mathematical methods for quantifying uncertainty. This area of mathematics is largely new, and the committee thinks that it has much to contribute to electric grid operations and planning. There are several kinds of uncertainty that affect efforts to begin merging real-time simulations with real-time measurements. These include the effects of modeling errors and approximations as well as the intrinsic uncertainty inherent in the intermittency of wind and solar generation and unpredictable fluctuations of loads. Efforts to create smart grids in which loads are subject to grid control and to generation introduce additional uncertainty. Hopefully, further development of smart grids will be able to exploit mathematical methods that quantify this uncertainty.

Some of the uncertainty associated with the next-generation grid is quite deep, in the sense that there is fundamental disagreement over how to characterize or parameterize uncertainty. This can be the case in situations such as predictions associated with solar or wind power, or risk assessments for high-impact, low-frequency events; with economic models that can be used to evaluate electricity markets or set rational pricing schemes for grid-connected distributed resources; or methods of assimilating peta-scale (or larger) data sets from distribution grids to help inform utility or consumer decision making and manage stochastic resources. The committee hopes that research into uncertainty quantification will extend to the development of new mathematical models to evaluate decisions made in the face of deep uncertainties. Existing decision models for power system planning, and existing economic models of electricity markets, have a very difficult time incorporating relevant time dimensions of any order.

The material in this chapter is intended to present sufficient background about these mathematical areas to understand important issues raised in their application to the power grid, building on the power grid material presented in Chapters 1 to 3. Chapter 5, Preparing for the Future, discusses challenges that the next-generation grid will present requiring new mathematical analysis, and Chapter 6, Mathematical Research Priorities Arising from the Electric Grid, discusses the new mathematical capabilities required to meet these challenges. Chapter 7's goal is to illustrate some current mathematical and computational techniques in greater detail than could be captured in earlier chapters in the form of case studies of the main challenges for mathematical and computational sciences created by ongoing changes in the power industry. Readers may wish to go directly to these descriptions, referring back to this chapter and earlier chapters when additional background material is needed. Finally, the committee notes here that differences in terminology between mathematics and power systems engineering sometimes cause confusion when the same word is given different meanings or when the same concept is described using different terminology. A simple example is the imaginary unit, which is i in mathematics and j in power engineering. This report attempts to give appropriate translations between the two fields.

DYNAMICAL SYSTEMS

Dynamical systems model the changes in time of interacting quantities. For electrical circuits, these quantities are the voltages and currents associated with the components of the network. As discussed in Chapter 1, power systems have many different time scales that need to be modeled as dynamical systems. Differential equations derived from the physical laws of electricity describe the rates of change of the variables. Kirchhoff's laws for currents and voltages in an electrical network, together with models for network devices like generators, transformers, and motors, are the core ingredients for models of the electric grid. System dynamics can be simulated by solving these equations with analytical or numerical methods to obtain the solution that begins with a specified initial condition. Usually, this process is repeated for many different initial conditions of the variables and values of the parameters that appear in the equations defining the system. The transient stability (TS) analysis for power systems, covered in Chapters 1 and 3, is one example that utilizes these types of simulations. Following a contingency, which might be a fault due to an equipment failure or a discrete control action that switches equipment off or on, the system has a transient response that (hopefully) leads to a new stable operating point. Another example is the faster electromagnetic transients analysis covered in Chapter 3. For both there are many levels of models that can be simulated, and an ever present question is whether the dynamical properties of coarser and finer models are consistent with one another.

Dynamical systems theory goes further and provides a language for interpreting and understanding simulation results. This theory emphasizes qualitative properties of solutions and long-time behavior. Guckenheimer and Holmes (1983) offer one among many introductions to this subject. They view a solution as a point moving in an abstract phase space of all possible values. The path traced along the solution is a trajectory. Where do trajectories ultimately go? Different patterns are possible. For example, the system may approach an equilibrium where the variables remain constant in time or a periodic orbit where they regularly return to values they have had previously. Both types of behavior are immediately relevant to power systems. More complicated quasiperiodic or chaotic asymptotic states are found in many systems, and the study of their statistical properties has been another focus of the theory. Stability is a central concern: If the asymptotic state is perturbed, does the system return to its previous behavior? One goal of the theory is to produce a phase portrait that depicts which trajectories have the same asymptotic states.

Structural stability is a further question of interest: If system parameters are varied, does the phase portrait have the same topology? For equilibrium points and periodic orbits of a structurally stable system, linearization produces eigenvalues that determine their stability. An equilibrium is linearly stable if all of its eigenvalues have negative real parts. In the power systems literature, this is referred to as dynamic stability. The basin of attraction of a stable equilibrium determines which perturbations of the state of the system produce trajectories that return to that equilibrium. The stable manifold theorem gives further information about equilibria and periodic orbits that are not attractors, which is useful in computing phase portraits and basin boundaries. Thus the basic question investigated by TS analysis is whether postfault states of a system lie in the basin of attraction of a desired attracting state. If not, control action that steers the system into this region is needed.

The ac design of the electrical grid presents a modeling challenge. The transmission grid has important time scales of minutes to hours that are much slower than the 60-Hz ac oscillations. Models that explicitly represent the ac oscillations have no equilibria; their simplest attractors are periodic orbits with a period of 1/60 sec. In averaged systems liked those used in power flow they become equilibria. This makes analysis significantly more difficult because finding equilibria requires the solution of only algebraic rather than differential equations. Simulation is also more difficult because numerical methods must use small time steps that are able to track the rapid oscillations. Consequently, models that average these oscillations are commonly employed. The averaging process is only valid for systems that are almost balanced and synchronized. Departures from these conditions require that more detailed models be used. The handling of the dynamics covered in electromagnetic transients analysis, with time scales of milliseconds or microseconds and hence significantly faster than fundamental 60 Hz, is another area for additional research.

Bifurcation theory gives a qualitative classification of the changes that occur in phase portraits when structural stability fails. Many properties of the bifurcations are universal and have been used as landmarks in the analysis of systems from diverse fields. For the averaged models of power systems, normal operating points are equilibria. As parameters of the model are changed to represent slow changes within the system, bifurcations of the equilibria may occur. With variations of a single parameter, there are only two kinds of generic bifurcation of an equilibrium in the space of all smooth vector fields: saddle-node bifurcations and Hopf bifurcations. In a saddle-node bifurcation, an equilibrium becomes unstable by merging with another equilibrium that has a single unstable mode (eigenvalue/eigenvector). The result for the power system is the small disturbance voltage collapse mentioned in Chapter 1 in which a blackout occurs without an apparent precipitating event. As the system parameters move closer to the bifurcation, fluctuations in the direction of the critical mode will be damped increasingly slowly. Real-time measurement of this slowing rate is one strategy for anticipating and preventing an incipient voltage collapse (Dobson, 1992). Hopf bifurcations of an equilibrium point initiate oscillations of the averaged system. These are manifested as rhythmic changes of voltages and currents. Large-scale oscillations of the power system with frequencies between approximately 0.1 and 5 Hz are occasionally seen. The widespread deployment of PMUs is providing the raw data needed to drive this research.

A caveat for the use of bifurcation analysis is that it describes generic behaviors. When applied in a context where systems have structure that limits the allowable perturbations, then bifurcation analysis can still be used *within the framework of perturbations that retain the structure*. One example is conservation of energy in conservative mechanical systems. Most dynamical systems do not have global functions that remain constant on trajectories, but conservative mechanical systems do. This prevents such systems from having asymptotically stable equilibria or periodic orbits. Bifurcations in this restricted class of systems can be investigated, but the possibilities are very different. Thus, identification of the setting within which systems of interest are generic is an important aspect of the qualitative analysis of their dynamics. Models of the electrical grid have a network structure inherited from the physical network. An important theoretical challenge is to determine how the network structure of coupled systems of oscillators (like the power grid) constrains their dynamics. Progress in this area could lead to new design principles for the grid.

Numerical algorithms that compute approximate solutions of differential equations are an essential tool for simulating trajectories of dynamical systems. Development of these algorithms is mature, and the algorithms are one of the most frequent types of numerical computation used today. Their performance is limited by the characteristics of the system being studied. Hairer and Wanner (2009) give a comprehensive survey of numerical methods for solving initial valuation problems of ordinary differential equations. Multiple time scales are frequently an issue and must be confronted by numerical methods. Initial value solvers advance approximate solutions of a system in time steps that are constrained by the fastest time scales in the system. Consequently, very large numbers of steps may need to be used to determine the behavior of the system on slower time scales. Specialized stiff methods use step sizes commensurate with slow time scales when trajectories move along attracting slow manifolds (see Hairer and Wanner, 2004). Multiple time scales are a prominent feature of power systems: Fast transients may occur in microseconds, while growing instabilities that lead to blackouts may happen on scales of minutes and hours. Another issue for the power system is that there are many discrete events that occur as loads and generators are turned off and on, equipment fails, lightning strikes, or protective devices like circuit breakers trip. Accurate simulation of a system must determine precisely when state-based events occur, introducing an additional layer of complexity to models of the power grid and to simulation software.

Maintaining reliability in the face of equipment failures, accidents, and acts of nature is a fundamental goal of grid operations. The N – 1 reliability introduced in Chapter 1 for contingency analysis mandates that a power system should operate with no limit violations following the failure of any single component. For steady-state considerations, power flow is used to assess whether the postcontingency equilibrium point (i.e., the power flow solution) has limit violations. However, as covered in Chapter 3, TS analysis is increasingly being used to assess whether the system can reach this new equilibrium point when system dynamics are considered. The many components in a large network necessitate large amounts of simulation, prompting the development of methods that directly locate bifurcations of a system without simulation. The example of saddle-node bifurcations and voltage collapse illustrates such algorithms. Root-finding algorithms can locate equilibria of a dynamical system, typically with only a few iterative steps rather than the large numbers of steps used by initial value solvers. When an equilibrium has been located, its eigenvalues determine its stability. The presence of a zero eigenvalue is a defining equation for saddle-node bifurcation. Linearization of the model equations and inclusion of a varying system parameter lead to an augmented system of equations whose solutions locate the saddle-node bifurcations. Continuation methods incorporate systematic procedures to determine how equilibria and their bifurcations depend upon variation of additional system parameters. Similar strategies are used to locate periodic orbits of a dynamical system with boundary value solvers replacing the root-finding algorithms for locating equilibria. Bifurcation analysis is still an area of evolving research, but it provides tools that go beyond simulation for studying stability of a system. See Kuznetsov (2004) for a description of these methods and AUTO[1] and MatCont[2] for relevant software.

OPTIMIZATION

The goal of optimization algorithms is to minimize an objective function, subject to both equality and inequality constraints. One way to classify optimization problems refers to permitted properties of the variables. In continuous optimization, variables are allowed to assume real values (e.g., the amount of electric current or power). Discrete optimization problems, also called integer programming problems, require variables to be integers. This restriction is appropriate, for example, when a variable signifies whether a generator is on or off, such as in power system unit commitment. In addition, there are "mixed" problems, in which only some of the variables must be integers.

General-Purpose Optimization Methods and Software

Continuous optimization, including both theoretical analysis and numerical methods, has been an active research area since the late 1940s. During the decades since then, there has been consistent and significant progress, punctuated by bursts of activity when a new, or apparently new, idea becomes known. In addition, researchers constantly revisit "old" methods that may have been abandoned or deprecated, not necessarily for good reasons. In particular, changes outside optimization (for example, the wide availability of parallel computing and the growing demand for solving machine learning problems involving big data) have led to changes in perspective about optimization methods. The issue is even more complicated because there is, in general, no unarguably best method, even for relatively narrow problem classes. So, today's state of the art in generic continuous optimization includes both new and old ideas.

The case of linear programming (LP), which is the optimization of a linear function subject to linear constraints, illustrates the swings in opinion about solution methods. The simplex method, invented by Dantzig in 1947, was the workhorse of LP for more than 35 years but was seen by some as unreliable because of its worst-case exponential complexity. In 1984, Karmarkar began the "interior-point revolution" with his announcement of a polynomial-time algorithm that was faster in practice than the simplex method. Because of their polynomial-time complexity, it was predicted by some researchers that interior-point methods would quickly replace the simplex method, but this has not happened. Thirty years later, connections are known between a wide family of interior-point methods and classical methods, and a variety of interior-point methods are used with great success to solve new problem classes (such as semidefinite programming, to be described later). In addition, algorithmic improvements have

[1] Computational Mathematics and Visualization Laboratory, "AUTO: Software for Continuation and Bifurcation Problems in Ordinary Differential Equations," http://indy.cs.concordia.ca/auto/. Accessed September 15, 2015.

[2] MatCont is available at https://sourceforge.net/projects/matcont/, last updated November 27, 2015. Accessed December 1, 2015.

continued to be made in software that implements the simplex method. As described by Bixby in 2015, there is no clear favorite for LP problems: Interior-point methods are faster than the simplex method for some linear programs, and the simplex method is faster for others, and there is no guaranteed technique for judging in advance which will be better. (This is an open research problem.)

The best known general-purpose software packages today for nonlinearly constrained optimization including the codes CONOPT,[3] IPOPT, KNITRO, MINOS,[4] and SNOPT,[5] are founded on plain vanilla versions of several techniques, including generalized reduced-gradient methods, augmented Lagrangian methods, penalty methods, sequential quadratic programming methods, interior-point/barrier methods, trust-region methods, and line-search methods (see Nocedal and Wright, 2006, for definitions and motivation).

The underlying methods in all these codes are widely taught in generic forms, and the associated software is updated frequently, often borrowing ideas from other methods. One reason for changes in software is the need for reliability when presented with a situation that is impossible, such as inconsistent constraints, even in an idealized world.

A further reason that optimization software must be updated to retain maximum efficiency is that optimization methods for large-scale problems rely on linear algebra, a research area that itself is progressing. High-quality optimization software invariably uses linear algebraic techniques whose speed depends on the size, structure, and sparsity of relevant matrices, and these properties may change during the course of iteration toward a solution. For example, some optimization methods factorize or update a matrix whose dimension increases as the number of currently active constraints increases, while other methods factorize a matrix whose dimension decreases in the same circumstances. These properties vary widely from problem to problem and cannot, in general, be deduced in advance.

The desirable property of convexity and the undesirable property of nonconvexity are often mentioned in describing the state of the art in nonlinear optimization. In the view of an eminent optimization researcher,

> The great watershed in optimization isn't between linearity and nonlinearity, but convexity and nonconvexity. (Rockafellar, 1993)

Broadly speaking, optimization problems involving convex functions tend to be *nice* in several precise senses. (For example, any minimizer is the unique global minimizer; convex optimization problems can often be solved rapidly, with theoretical guarantees of convergence.) In contrast, the presence of even a single nonconvex function can cause an optimization problem to become highly difficult, even impossible, to solve.

Some optimization methods, such as quasi-Newton methods, explicitly control the nature of matrices used to represent second derivatives. But this is problematic when exact second derivatives are provided for the objective function and constraints, since the method is, in effect, changing the problem. Methods that use second derivatives therefore include a variety of sophisticated strategies, often based on criteria from a linear algebraic subproblem, when indefiniteness is detected, as is always a possibility with nonconvex or even highly ill-conditioned convex problems. (See below for further comments on nonconvexity.) However, by definition, it is, in general, not possible to guarantee convergence to the global optimum if the problem being solved is nonconvex. Even for a quadratic program (minimizing a quadratic function subject to linear constraints), finding the global minimizer is NP-hard.

A ubiquitous feature of general-purpose optimization software is the presence of numerous parameters that can be chosen by the user, as well as a set of default values and guidance about choosing them. Documentation for these codes always stresses the importance, for difficult or delicate problems, of setting these parameters with great care, since they can have a huge impact on the performance of the software.

[3] See the CONOPT website at http://www.conopt.com/. Accessed September 15, 2015.

[4] Systems Optimization Laboratory, "User Guide for MINOS 5.5: Fortran Package for Large-Scale Optimization," http://web.stanford.edu/group/SOL/minos.htm. Accessed December 1, 2015.

[5] University of California, San Diego, "SNOPT," last updated May 12, 2015, https://ccom.ucsd.edu/~optimizers/software.html. Accessed September 15, 2015.

Grid-Related Continuous Optimization

As in general-purpose optimization, there has been consistent progress in optimization for grid-related problems. For example, consider the ac optimal power flow (ACOPF) problem introduced in the early chapters, whose basic concept was formulated by Carpentier in 1962:

> . . . minimize a certain function, such as total power loss, generation cost, or user disutility, subject to Kirchhoff's laws, as well as capacity, stability, and security constraints (Low, 2014).

Broadly speaking, an ACOPF is typically used for determining the settings for the power system controllers, such as generator real power outputs, so that the total generation is equal to load plus losses, all the controllers are within their limits, and there are no power system limit violations. A standard version of the ACOPF involves minimization of a generic quadratic function subject to quadratic constraints, and this problem is known to be NP-hard (see the section "Optimization" in Chapter 6). As was discussed in Chapter 1, the ACOPF is an optimization applied to the standard ac power flow equations.[6] Hence the power flow can be thought of as determining a feasible solution for the ACOPF. However, as was noted in Chapter 1, the power flow not only involves solving a set of nonlinear equations but is also augmented to model the automatic response of various continuous and discrete power system controllers. For background material, see Andersson (2004) and Glover et al. (2012). Bienstock (2013) is a recent survey.

General-purpose optimization software has been applied for many years to some ACOPF problems. In particular, starting in December 2012, a series of reports about the ACOPF was produced by the Federal Energy Regulatory Commission (FERC) addressing numerous aspects of the state of the art—including history, modeling alternatives, and numerical testing. (See, for example, Cain et al., 2013; Castillo and O'Neill, 2013a,b.) Despite the known difficulty of the problem, application of the best available general-purpose software for nonlinear optimization (see the subsection on general-purpose optimization methods and software in this chapter) to a variety of formulations of the ACOPF produced, in many cases, acceptable solutions (Castillo and O'Neill, 2013b).

Even so, a familiar scenario has arisen in which practitioner expectations rise as problems previously viewed as intractable change from being "challenging" to being "easy." At the November 2014 ARPA-E workshop, Richard O'Neill, the Chief Economic Adviser at FERC, said that "ac optimality has been an unachievable goal for 50+ years" (Heidel, 2014). Not surprisingly, incarnations of ACOPF submitted for numerical solution have become much harder to solve because the details of the problem formulation have not remained the same: Not only has the mathematical form of the problem become much more complicated, but also the associated dimensions have increased (Ferris et al., 2013). Even with today's highest-end computing, some important versions of ACOPF are too large to be solved within an acceptable time frame, say in a real-time environment.

The precise reasons for this unsatisfactory situation are not fully understood, but one cause is widely perceived to be related to nonconvexity. Recent approaches that attempt to finesse nonconvexity in the ACOPF and other grid-related problems are discussed in Chapter 6.

Mixed-Integer Linear Programs

Integer programming concerns the solution of optimization problems where some of the variables are explicitly integer valued. The most common case arises with binary variables, and there are several settings in which such variables naturally arise. In a power engineering context, the unit commitment problem (see Sheble and Fahd, 1994, for background) decides which generators will operate over a certain time window. Hedman et al. (2011) discuss whether it is advantageous to switch off some transmission lines: It is noteworthy that power systems effectively can exhibit nonconvexities that make line switching an attractive option.

[6] As covered in Chapter 1, generically the term power flow refers to the solution of the nonlinear power flow equations. It is occasionally called the ac power flow; the term dc power flow refers to a solution of a set of linear equations that approximate the nonlinear power flow equations. Both the ac power flow and the dc power flow usually determine an equilibrium point for an assumed balanced, three-phase 50- or 60-Hz system.

The line-switching problem proves a useful reference point to describe the capabilities (and limits) of integer programming technology. In the dc power flow the approximation of the per unit, real power flow on the transmission line between buses k to m with phase angles of θ_k and θ_m is

$$P_{km} = \frac{1}{X_{km}}(\theta_k - \theta_m) \tag{1}$$

where X_{km} is the per unit reactance for the transmission line between the buses. In the line switching problem, a binary variable w_{km} would be added, which is given the value of one when line km is switched off. Equation (1) needs to be modified in order to reflect this relationship, and there are several ways to do so—for example, by replacing (1) with the system

$$\frac{1}{X_{km}}(\theta_k - \theta_m) - MW_{km} \le P_{km} \le \frac{1}{X_{km}}(\theta_k - \theta_m) + MW_{km} \tag{2}$$

$$| P_{km} | \le M'(1 - W_{km}) \tag{3}$$

where M and M′ are appropriately large positive constants. When $w_{km} = 0$ (line not switched off), constraint (2) is equivalent to (1), and (3) is inactive if M′ is large enough. When $w_{km} = 1$ (line switched off), constraint (3) enforces $P_{km} = 0$, while (2) is inactive if M is large enough, allowing θ_k and θ_m to assume any values. Thus, subject to the stipulation that an optimization engine capable of forcing the w_{km} variables to take binary values is employed, system (2) correctly models the line switching paradigm.

Of course, how to solve the resulting optimization problem in the case of a transmission system with thousands of lines is a nontrivial task. Typical optimization methodologies will solve a sequence of convex (and thus continuous) optimization problems that progressively approximate the discrete optimization problem of interest and that are gradually adjusted so as to eventually converge. This is a highly technical field with many pitfalls—for example, the arguments that led to system (2) may produce large values of M and M′ (a "big M" method), a strategy that is known not to be ideal, and indeed one seeks to choose such values as small as possible while still producing a valid formulation.

The field of integer programming has rapidly progressed in recent decades to the extent that many problem classes that were considered unsolvable are now routinely solved. For basic background, see Nemhauser and Wolsey (1988), Schrijver (1998), and Wolsey (1998). In the case of the unit commitment problem, the measurable improvements, from an industry standpoint, have been remarkable and it is safe to say that mixed-integer programming technology is now the default choice (see Bixby, 2015).

Binary mixed-integer programming also arises in other settings—for example, in bilevel programming (Bard and Moore, 1990), often used to model adversarial settings. It can also arise when modeling logical conditions that do not reflect a preexisting or straightforward operational decision. In a power engineering setting an example could be the operational option to set the output of a given generator to a certain range if estimated wind turbine output falls in another range. In such a setting the ability to model (and act on) this decision is captured by a binary variable; however, that binary variable is not one that would directly arise in the power engineering context.

In recent years a new field has emerged that is arguably even more compelling from an engineering perspective: mixed-integer nonlinear programming. It is easy to argue for the importance of this field in a power engineering context. Both problems discussed above, the unit commitment and line switching problems, arise in an ac power flow setting, in which case the problems have both continuous and binary variables but where the underlying equations are nonlinear. This apparently simple change radically increases the complexity of the computational task: The successes cited above were all in the case of linear integer programming, which has heavily benefited from much improved LP engines and a much deeper intellectual understanding of linear mixed-integer optimization. In the nonlinear setting, by contrast, both the basic computation (of nonlinear, nonconvex systems of equations) and the deep mathematics of discrete optimization are significantly more challenging. Fortunately, this field is a hot one in the optimization community, and a large number of talented researchers are now contributing interesting work (see Belotti et al., 2013, for a recent survey).

Stochastic Optimization

The operation of a power transmission system requires the frequent solution of complex mathematical problems that take as input a large amount of data, much of which may be estimated or forecast. In fact, even the underlying mechanism that causes uncertainty in data may itself be poorly understood. The optimal power flow (OPF) is a good example of a data-rich optimization problem that is exposed to data uncertainty. In an online context it can be run as frequently as every 5 min in order to set the output of generators over the next time window. In a time period of a day or two, the OPF would also be combined with unit comment to determine which generators will actually be committed over, say, a 24-hour time span. This is a mixed-integer (i.e., noncontinuous) optimization problem. In a longer-term planning context the OPF might be used to look at assumed system conditions many years into the future.

These examples involve mathematical problems of significant intrinsic complexity, calling for sophisticated algorithms that must be implemented with care and should run accurately and quickly. But, as the committee indicated above, the data inputs for such algorithms may not be precisely known. A current and increasingly compelling example is provided by wind turbine output in the context of the OPF problem: In a small time window the standard deviation of the wind speed can be of the same order of magnitude as the expectation; managing this uncertainty is clearly important and nontrivial (Bienstock et al., 2014). Another example is the forecast of loads (demands) in the context of unit commitment. Uncertainty of loads over a 24-hour period can be significant (if, say, weather conditions are uncertain), so it is important to handle this uncertainty in a cost-effective manner that does not place the grid into shortfall conditions (Sheble and Fahd, 1994; Shiina and Birge, 2004). Trying to forecast loads years into the future is even more uncertain.

Uncertainty can be classified very broadly into a number of categories, all to achieve the general goal of computing controls and policies that are robust as well as "optimal":

- *Pure noise*. Any optimization mechanism that is presented with a single point estimate of data parameters may produce nonrobust policies—the mechanism will optimize assuming the data are precise, and the resulting mechanism may fail if the actual data deviate, even by a small amount. The challenge is to produce policies that are robust with respect to such small data deviations while remaining cost effective.
- *Model uncertainty*. It may be the case that the source of data uncertainty is poorly understood. A decision-making tool that assumes a particular model for uncertainty (a particular stochastic data distribution, say, or a causal relationship) may produce unreliable outcomes. The field of robust optimization (Ben-Tal and Nemirovski, 2000; Bertsimas et al., 2011) seeks to produce methodologies that yield good solutions while remaining agnostic as to the cause of uncertainty. A challenge is to provide the flexibility to adjust conservatism in one's outlook.
- *Scenario uncertainty*. It can be that an important source of uncertainty is a particular, well-understood behavior or even parameter. For example in the unit commitment setting, one might be concerned that a generator may need to go off-line in the next 24 hours owing to a mechanical condition. In that case the decision may be made to trip the generator now or to postpone that decision until 12 hours from now, when more information will become available and when the relative likelihoods may be well understood—that is, the probabilities—of the events that will cause the various realizations of that information. This gives rise to a decision under fairly well understood alternative data scenarios. Such problem settings are the domain of stochastic programming (see Birge and Louveaux, 2011; Prékopa, 1995; and Shapiro et al., 2009).

The optimization community has developed a diverse set of methodologies for handling uncertainty that are effective, computationally fast, well grounded, and suited for analyzing the uncertain behavior of power systems and for producing robust operating schemes. Stochastic programming, described next, is an example of such a methodology. A subsection in Chapter 6 describes robust and chance-constrained optimization, a less mature set of methods deserving further development. These two approaches need not be (and should not be) deployed exclusively of one another; however, they are presented separately to highlight specific modeling strengths and weaknesses.

Stochastic programming is perhaps the oldest and most mature form of optimization incorporating stochastics. It is a scenario-driven methodology, and, more precisely, it aims to optimize expectation over a given and possibly very large family of scenarios with known probabilities. Here the committee presents an example that is abstracted from the unit commitment problem so as to highlight the role of recourse, an important modeling element. (See Higle and Sen, 1991; Oliveira et al., 2011; Papavasiliou and Oren, 2013; Wang et al., 2012, and citations therein.)

Consider a power system whose generators, G, are partitioned into two sets: G^L and G^F (slow and fast, respectively, according to their ramp-up speeds). At time $t = 0$ one needs to plan for generation over two consecutive stages or time periods, the first one beginning right now and the second starting at time Δ. The loads (demands) for the first stage are known. However, the loads during the second stage are not precisely known, and instead it is assumed that one of a fixed family S of known scenarios, each specifying a set of loads for the second stage, will be realized starting at time Δ; further, the probability of each scenario is known at $t = 0$. The actions available to the power grid operator are these:

(1) At $t = 0$ (here and now), choose which generators from the set G^L to start up and their respective output. Each such generator will incur two costs: a start-up cost and a cost depending on the output level.
(2) Additionally, resulting power flows must meet the loads in the first stage.
(3) At $t = \Delta$, the planner is assumed to observe which demand scenario has been realized. The planner can choose additional generators to start up from among the fast set G^F and can set their output so as to meet the demand in the given scenario.

Step (3) embodies the recourse—the model allows the planner to delay committing generators until the demand uncertainty in the second stage is resolved. To cast this problem in (summarized) mathematical form, the following variables are used: for each generator g, let $y_g = 1$ if g is to be started at $t = 0$, and write $y_g = 0$ otherwise. Furthermore, set p_g as the output of generator g during the first stage. Likewise, for each $g \in G^s$ and each scenario s, let $w_g^s = 1$ if, under scenario s, g is to be started at $t = \Delta$, and write $w_g^s = 0$ otherwise. The output of any generator g in scenario s is denoted by p_g^s.

Using these conventions, the optimization problem is written as follows:

$$\min\left[\sum_{g\varepsilon G^L} K_g^{(1)} y_g + f_g^{(1)}\left(p_g\right)\right] + \sum_{s\varepsilon S} \pi_s \left[\sum_{g\varepsilon G^F} K_g^{(2)} w_s^s + \sum_{g\varepsilon G^L} f_g^{(2)}\left(p_g^s\right)\right] \tag{4}$$

so that p is feasible during the first stage; $L_g y_g\, p_g \leq U_g y_g$ for all g; (p, p^s) is feasible during the second stage under scenario s; $\mathrm{L}_g w_g^s \leq p_g^s \leq U_g w_g^s$ for all $g \in G^F$ and scenario s; and all variables y_g, w_g^s take value 0 or 1.

In this formulation, π_s is the probability that scenario s is realized and $K_g^{(k)}$ and $f_g^{(k)}$ are, respectively, the start-up cost and the cost of operating generator g at output level x during stage k (= 1, 2). The constraints are used to indicate in shortened form that the generator outputs can feasibly meet demands; these constraints will normally require additional variables (e.g., phase angles) and constraints (e.g., line limit constraints) to fully represent feasibility. The quantities L_g and U_g indicate, respectively, the lower and upper limits for output of generator g. Note the binary variables, which could make the problem difficult, especially if the number of scenarios is large. In practice, the functions f_g, f_g^s would be convex quadratic. The solution to this problem would provide one with the correct actions to take at time 0, plus a recipe to follow when the appropriate scenario has been identified at time Δ. This approach is attractive in that it allows hedging without committing the most per-unit expensive generators at time 0. All that is needed is a good representation of the demand scenarios and, of course, an adequately fast solution methodology. Regarding this last point, the formulation will generally be a prohibitively large mixed-integer program, a common feature of realistic stochastic programming models. A number of mature methodologies have been developed to address this issue with some significant computational successes, sometimes requiring significant parallel computational resources. See, for example, Birge and Louveaux (2001) and Higle and Sen (2012). A common technique involves decomposition—for example, iterative methods that progressively approximate the feasible set by addressing a subformulation (say, by focusing on one scenario at a time in the above example). The approximation is attained through the use of cutting planes. Such techniques include the L-shaped method (Van

Slyke and Wets, 1969), Benders' decomposition (Benders, 1962), and sampling methods (Linderoth et al., 2006). The committee notes, in passing, that an additional hazard in this formulation is the presence of binary variables used to model the second stage. This casts the formulation as a so-called bilevel program (Bard and Moore, 1990).

CONTROL

Power system controllers are divided into two main categories: protection and control. Protection is associated with event-driven controls to isolate and clear primarily short-circuit faults. Control refers to continuous processes that enable the power system to operate. Controls are further subdivided into primary, secondary, and tertiary control.

Protection

As discussed in "Short-Circuit Analysis" in Chapter 3, when a short circuit occurs, the fault current can be orders of magnitude greater than the full load rating of the equipment. The objective of clearing faults is to limit the damaging thermal heat that will be generated with sustained fault currents. Over the past several decades, significant engineering emphasis has been focused on fast detection and clearing of faults. It is not uncommon for a protective relay to be able to determine if a fault has occurred within a quarter cycle (~5 msec), and fast breakers can clear the fault after about 3 or 4 cycles (~50 msec). Often, particularly for lower-voltage infrastructure, the cost of high-speed fault clearing is not justified, and faults may take much longer to clear (on the order of 100 msec).

Most of the engineering associated with protection is being able to reliably detect a fault and isolate the minimum amount of assets needed to clear it. One key element of protection is to avoid cascading failure in a networked system. Time-overcurrent protection is only used to protect radial circuits, meaning circuits that do not have any electrical path back to the point of protection other than directly back via the radial path out. Another key philosophy is to provide backup protection in case of a failure in the primary scheme—for example, a breaker failing to clear the fault. Therefore, overlapping zones of protection are implemented. Another key philosophy is to not become overly reliant on communications. Therefore, the algorithms that are used depend on communications capability; for example, differential protection is normally used in a substation, and impedance relaying is normally used for transmission lines. Many faults are transient, such as lightning strikes to the power line. Therefore, high-speed reclosing will automatically restore the line after a few cycles, when the fault current plasma path to ground has been extinguished. Given that many of these transient faults are single-line faults, some extra-high-voltage lines can achieve reliability benefits by employing single-pole switching to clear single-line transient faults. There are many more examples. Being selective and secure are the key objectives of a protection engineer, and whole departments at utilities are devoted to this engineering discipline.

Primary Control

Primary control employs fast-acting closed-loop feedback that does not rely on remote communications. It is used to control localized processes. Generation is one element of the power system where primary control is critical to the efficient and reliable operation of the power grid. Two important examples of generator primary controls are the governor and the voltage regulator. The modeling of this primary control is done in transient stability, covered in Chapters 1 and 3.

The governor, which controls the amount of mechanical torque on the generator shaft, controls the speed of the generator when it is not connected to the power system. When the synchronous generator is connected to the power system, its speed is driven by the aggregate response of all generators within the entire interconnection. At that point, the governor regulates the real power output (wattage) of the generator by adjusting the mechanical torque on the shaft.

Droop is an important control concept for governors. When generators were first paralleled over a hundred years ago, it was discovered that there needs to be a mechanism for them to proportionally share variations of load without causing stability issues associated with the generators' response that inadvertently affect the other machines. The solution was to include a relationship between the electrical frequency and power output, whereby when the frequency of the aggregate system decreases there is a programmed response of increased power from each unit. This approach achieves load-balancing stability because incremental changes in the balance between load and generation, including minute variations, will stimulate a reaction by all of the connected governors to control their generators in a way that brings the system to a new stable operating point.

It is important to understand that not all generators may have governor response. Many generators, particularly smaller ones, are designed such that their power output will not vary at all based on system frequency. It is also possible that generators that might normally have governor response will not respond to a frequency variation. For example, a generator with the prime mover control already set at its maximum power setting cannot increase its power any further based on a drop in system frequency.

While traditional governors were mechanical engineering marvels (e.g., a spring-loaded spinning mass connected to a throttle linkage), many have been replaced by digital governor controls. Though these digital systems have advantages with respect to cost, maintenance, and reliability, they have usually been programmed to mimic the operation of the mechanical devices they replaced. Therefore, concepts such as deadband (preventing wear and tear on mechanical linkages by preventing throttle adjustments to small variations in system frequency) and droop could conceivably be replaced in the future by more sophisticated algorithms to provide primary response to frequency variations that might be better optimized from an economic standpoint—that is, perhaps by changing response based on the rate of frequency change or by evaluating the incremental marginal cost of the generation unit compared with the system conditions, or by similar means). A hydro governor block diagram is given in Figure 3.4.

The voltage regulator controls the exciter, which provides the dc field current for the synchronous generator. When the generator is not connected to the system, this directly controls the terminal voltage. When the generator is connected to the system, the dc field current will more directly control the reactive power output of the generator, which is closely coupled to its terminal voltage. During normal operations, the voltage regulator will adjust the reactive power output of the generator to regulate the terminal voltage within the maximum and minimum range of reactive power rating for the unit. The North American Electric Reliability Corporation requires all generators directly connected to a transmission network to be in voltage control.

Again, much like the governor control, not all units are designed and operated to make these adjustments in reactive power output in response to changing system conditions. Particularly for smaller generators, they may operate in constant power factor mode, where the operator sets a desired amount of reactive power output (either leading or lagging) that will not change regardless of the terminal voltage.

There are other examples of primary control that can be installed on generators. The power system stabilizer was developed when solid state exciters became prominent toward the end of the 20th century and the fast-acting nature of exciters required additional control to prevent instability under specific system conditions. Transient excitation boost was installed on some generators to temporarily raise the rotor field current during local faults to boost terminal voltage and help coordinate with the protection scheme.

Other nongenerator examples of primary control include voltage-regulating transformers in distribution substations or static var compensators to continuously regulate voltage at a transmission substation by adjusting the reactive power output of the device. All of these are closed-loop fast controls that achieve their function without remote communications.

Secondary Control

Managing the frequency of an interconnected power system is the main use of secondary controls today. As was previously discussed, the droop characteristics of speed-governor control of the generators will achieve a stable equilibrium operating point associated with any small or large variations between load and generation. However, this equilibrium point will not necessarily be at the desired frequency (60 Hz in North America), particularly if the disturbance is large. Furthermore, the frequency can slowly drift from 60 Hz due to continuous small variations, and while the governors will ensure that there is ample power to meet the load at all times, there nevertheless needs to be a system to maintain the system frequency at 60 Hz.

Each interconnection is divided into balancing areas, and each balancing area is obligated to maintain system frequency by adjusting the generation within that area. This is accomplished by each balancing authority calculating on a continuous basis (usually every 4 sec) a parameter called the area control error (ACE), which was introduced in Chapter 1. The units of ACE are megawatts, and ACE is composed of inadvertent interchange, frequency mismatch multiplied by a bias (an assigned value based on an assessment of that balancing area's ability to affect frequency, whose units are megawatts per hertz), and sometimes includes other parameters.[7] Usually, inadvertent interchange is the dominant term in ACE and is the difference between the actual and scheduled flow of electricity between the balancing area and its neighboring balancing areas.

This calculation of ACE is then disaggregated into desired generation response and dispatched to the individual stations through a process called automatic generation control. Again, this is a continuous process and usually operates at the same timing as the calculation of ACE (every 4 sec). The communications infrastructure to disseminate these new generation set points and to gather the actual interchange data is known as SCADA. Of course, SCADA telemetry is also used for many other things.

Some interconnections, such as the Electric Reliability Corporation of Texas, have only one balancing area within the entire interconnection, and thus concepts such as inadvertent interchange do not apply.

Tertiary Control

Other controls to better optimize power system operations are referred to as tertiary control. These controls are typically slower and take place over several minutes to hours. There is a lot of variation among different system operators depending on the specific needs of their system. Examples include algorithms to adjust voltage set points to minimize reactive power-loop flows and system losses, capacitor bank switching schemes, phase shifting transformer adjustments, trained human operators, and so on.

Relating this back to Chapter 3, a power flow solution needs to have all of these control systems represented in order to determine the new steady-state solution following some system contingency. Usually the primary control is implicitly represented in solving the nonlinear power flow equations. The secondary and tertiary controls are usually modeled as explicit controls, with the power flow algorithm needed to determine the new values for a mixture of discrete and continuous controllers. In transient stability, which is used to determine whether the system can reach a new steady-state solution, the primary control system is explicitly represented along with some of the secondary and tertiary controls.

RISK ANALYSIS, RELIABILITY, MACHINE LEARNING, AND STATISTICS

Power systems are composed of physical equipment that needs to function reliably. Many different pieces of equipment could fail on the power system: Generators, transmission lines, transformers, medium-/low-voltage cables, connectors, and other pieces of equipment could each fail, leaving customers without power, increasing risk on the rest of the power system, and possibly leading to an increased risk of cascading failure. The infrastructure of our power system is aging, and it is currently handling loads that are substantially larger than it was designed for. These reliability issues are expected to persist into the foreseeable future, particularly as the power grid continues to be used beyond its design specifications.

The focus here is on data-driven risk analyses, where data from the power system are used to inform risk assessments related to component failure or other degradations to the system. These data might include sensor measurements from various equipment including PMUs, reports of past equipment inspections, state measurements, failure reports, or other measurements. This section differs from the sections on dynamical systems in that only a partial or even no physical model (no dynamical system, no set of differential equations) might be available, and the predictions are primarily data driven rather than hypothesis driven.

[7] In the Western Interconnection of North America, time error correction is a continuous function included in the ACE calculation. Much like integral control, it is accounting for the sum total of prior frequency mismatch. The Eastern Interconnection still performs time error correction by deliberately operating the system frequency slightly off 60 Hz to account for prior frequency error.

In the subsections below, the committee provides an overview of several core predictive modeling problems and discusses how they are relevant to the power grid. Data-driven methods often grapple with a classic challenge called the "curse of dimensionality," where the complexity of the model trades off with the amount of data available to estimate parameters. For each parameter that needs to be estimated, it is possible that an exponentially larger amount of data is needed in order to estimate it. Thus, there are challenges in handling very large amounts of data in order to fit a model with many parameters (typically requiring large-scale optimization), and challenges in designing more structure to the parameters so that not as many data are required in order to produce a useful model. With the extra structure come more complex optimization problems in order to fit the model to data.

Regression

Regression can be demonstrated by a set of training examples $\{(\mathbf{x}_i, y_i)\}_i$, where the $\mathbf{x}_\mathbf{i} \in \mathbf{R}^p$ are points in a (possibly high-dimensional) real-valued space and y_i is real-valued, $y_i \in \mathbf{R}$. The goal is to construct a function $f : \mathbf{R}^p \rightarrow \mathbf{R}$ such that $f(\mathbf{x})$ predicts y for a new $\mathbf{x}$. Linear models are classical, where $f(\mathrm{x}) = \sum_j \beta_j x_{\cdot j}$, where $x_{\cdot j}$ is the jth component (covariate) of vector $\mathbf{x}$, possibly including a constant intercept term. Regression is a classic problem that is pervasive, and much work in modern statistics and machine learning still revolves around variants of linear regression. The most important regression problems related to the power grid are those of estimating demand:

- *System load forecasting (estimation of demand for a region)*. Reliable energy delivery depends heavily on the estimation of demand, because energy cannot be stored and must be generated to meet the estimated demand. Consider the problem of estimating demand for a city tomorrow, where each $\mathbf{x}$ represents a vector containing information about the weather such as temperature, precipitation, pressure, demand from several past days, day of the week, whether tomorrow is a holiday, and any other information that would be relevant for predicting demand y. This type of multiple regression approach is used by major power companies such as PJM (2014). See also Soliman and Al-Kandari (2010).
- *Individual load forecasting (estimation of demand for a single building)*. Demand also needs to be estimated for each customer if a blackout occurs, for the purpose of resolving lawsuits. In particular, a penalty might need to be paid for power that was not delivered during a blackout, in which case one needs to estimate the customer's demand that would have been realized during the blackout (which is called the counterfactual). This can be difficult, because blackouts often occur when the demand is unusually high (Goldberg, 2015). Predicting individual-level demand is also relevant for making recommendations to customers on whether and how they should reduce their power consumption and for offering demand flexibility programs.

One often desires either the full distribution of y for each $\mathbf{x}$, the mean of y for each $\mathbf{x}$, or a particular quantile of y given $\mathbf{x}$ (quantile estimation).

Classification and Hazard Modeling

Classification and hazard modeling can be demonstrated with a set of training examples $\{(\mathbf{x}_i, y_i)\}_i$, where $\mathbf{x}_i \in \mathbf{R}^p$ and $y_i \in \{0, 1\}$; that is, with the aim to predict a binary outcome. The goal could be either to (i) construct a function f: $\mathbf{R}^p \rightarrow (0, 1)$, where f represents the probability of $y = 1$ given $\mathbf{x}$, or to (ii) construct f: $\mathbf{R}^p \rightarrow \{0, 1\}$, where f represents the decision, either 0 or 1. For (i), logistic regression is a classical approach to modeling probabilities, where the estimated probability is

$$P\left(y = 1 \middle| x\right) = 1/\left\{+ \ \exp[f(x)]\right\}. \tag{5}$$

Linear models are classical, where $f\left(x\right) = \sum_j \beta_j x_{\cdot j}$, where $x_{\cdot j}$ represents the jth component of vector $\mathbf{x}$. For (ii), y can be estimated directly, where the estimator is $\hat{y} = \mathrm{sign}(f(\mathbf{x}))$, and again f is classically a linear model. Many data sources on the power grid are time series, which means that each $\mathbf{x}$ is calculated at a sequence of times, $\mathbf{x}_i = \mathbf{x}(t_i)$, and the value y_i is 1 when an event happens at time t_i. In that case, survival analysis could be used. The hazard function $\lambda[t|\mathbf{x}(t)]$ is defined to be the failure rate at time t conditioned on survival until at least time t, where $\mathbf{x}(t)$ are a set of time-dependent covariates. If modeling equipment failure, covariates would encode (for instance) the manufacturer of the equipment, type of equipment, settings, and other factors that are potentially predictive. The time-dependent Cox hazard model assumes the hazard rate to be of the form

$$\lambda\left[t \middle| \mathbf{x}(t)\right] = \exp\left[\sum_j \beta_j x_j(t)\right] \tag{6}$$

where j represents the jth covariate (Martinussen and Scheike, 2006). Important classification and reliability problems related to the power grid include the following:

- *Asset failure prediction and condition-based maintenance*. Many pieces of equipment provide data throughout their lifetimes, including sensor measurements on the equipment such as from SCADA or PMUs, past failures and warnings, and past inspection reports and repairs. These data can be used to predict failures before they occur, which can be used to inform maintenance policies. Covariates can include the number of failures within the past year, the average of the signal from each sensor within the past 3 days, the maximum of the signal from each sensor within the past week, whether the equipment was made by a specific manufacturer, and the age of the equipment. Failure prediction methods can be useful for almost any type of equipment, from electrical cables to generators, transformers, manholes, wind turbine components, solar panels, and so on. (One source of material on learning in power systems is Wehenkel, 1998.)

  The application of "waveform analytics" to predict distribution system failures, along with several case studies, is given in Wischkaemper et al. (2014, 2015); specific case study examples include the detection of prefailure of a capacitor vacuum switch, detection of a failing service transformer, and the detection of fault-induced conductor slap in which fault current in a feeder induces magnetic forces in another location of the feeder causing the conductors to slap together.
- *Customer rebate adoption*. Many power companies offer rebates to customers in order to conserve power during periods of peak demand. These rebates are (and have already proved to be) critical in ensuring the reliability of the grid. The companies thus need to predict which customers will adopt a power conservation program when it's offered. These programs may provide some form of rebates for customers who allow the power company to curtail their consumption on certain days (demand-side flexibility). The question is whether it is possible to predict which customers will be receptive to which type of rebate. Consider a database of customers who were offered rebates. Each customer will be represented by a vector $\mathbf{x}$ that represents the household (number of adults, number of children), types of appliances, historical consumption patterns, and the like. The label y represents whether the customer responded to the rebate.
- *Energy theft*. One of the most important goals set by governments in the developing world is universal access to reliable energy (World Bank, 2010). While energy theft is not a significant problem in the United States, some utilities cannot provide reliable energy because of rampant theft, which severely depletes their available funding to supply power (World Bank, 2009). Customers steal power by threading cables from powered buildings to unpowered buildings. They also thread cables to bypass meters or tamper with the meters directly, for instance, by pouring honey into them to slow them down. Power companies need to predict which customers are likely to be stealing power and determine who should be examined by inspectors for lack of compliance. Again, each customer can be represented by a vector $\mathbf{x}$ that represents the household, and the label y is the result of an inspector's visit (the customer is either in compliance or not in compliance).

Causal Inference

The issues of customer rebate adoption and energy theft, described above, are related to questions of causal inference. Here, the goal is to determine a cause-and-effect relationship such as "this rebate causes these customers to alter their energy consumption behavior." There are many approaches to causal inference. For instance, to answer the question about rebates, one could design an experiment in which some customers receive the rebate and some do not. The committee could use the consumption patterns of these two sets of customers to determine whether there is a significant change in consumption as a result of the rebate. If no such experiment can be conducted to collect these data, the committee could look retrospectively at observational data and match each customer who received the rebate with a similar customer who did not receive one. While some causal inference problems might be solved with a simple hypothesis test, other problems could be very complex; for instance, if the causal effect depends on covariates $\mathbf{x}$, one might need separate regressions to estimate the outcome when the rebate is given and the outcome when it is not given.

Clustering

In clustering, a set of entities is given, possibly vectors $\{\mathbf{x}_i\}_i$, where $\mathbf{x}_i \in \mathbf{R}^p$, and the goal is to assign each of these points to an integer representing its cluster label. Since there are no ground truth labels for clustering, there are many different viable methods for measuring the quality of a clustering. One application of clustering for the power grid is energy disaggregation and nonintrusive load monitoring. The goal of energy disaggregation is to take a single power consumption signal (for instance from a house) and break it into usage by the various appliances (Sultanem, 1991; Hart, 1992). Figure 4.1 illustrates the energy consumption over time of a single household, where there are several appliances contributing to the overall consumption. One can see that each appliance has a unique signature—for instance, the stove burners have short repetitive cycles. Knowledge of detailed energy usage allows (and encourages) consumers to better conserve energy (Darby, 2006; Neenan and Robinson, 2009; Armel et al., 2012).

There are many variations of the energy disaggregation problem (Ziefman and Roth, 2011), but clustering is a key step in several established algorithms (Hart, 1992). The data are whole-house power consumption signals, which can be written as two time series: one for real power, real(t), and one for reactive power, reactive(t), both of which can be calculated from current and voltage signals. Edge detection filters of the form f_j[real($t-\Delta$), real($t+\Delta$)] are used to sense different types of changes in the real and reactive power levels, where each j is different. Edge detectors can be useful in finding, for instance, on and off cycles of a dishwasher or a clothes washing machine. Each signal is then represented as a vector of its f_j values. Cluster analysis is then useful for determining which edge signals belong to the same appliances. There are many cases in which data alone do not suffice to build realistic models; domain knowledge can often be used to build realistic structure into the model. The field of reliability engineering exemplifies this (see, for instance, Trivedi, 2001).

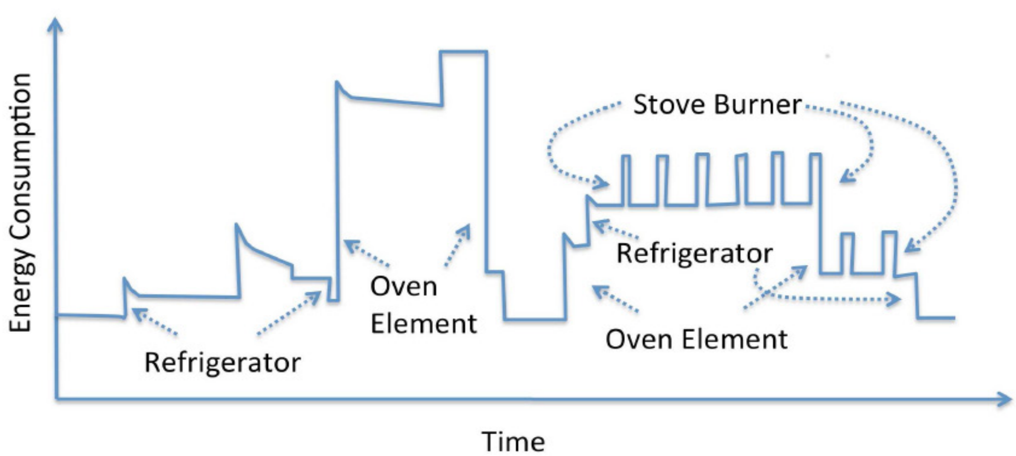

FIGURE 4.1 Energy consumption of a household over time. SOURCE: Courtesy of Cynthia Rudin.

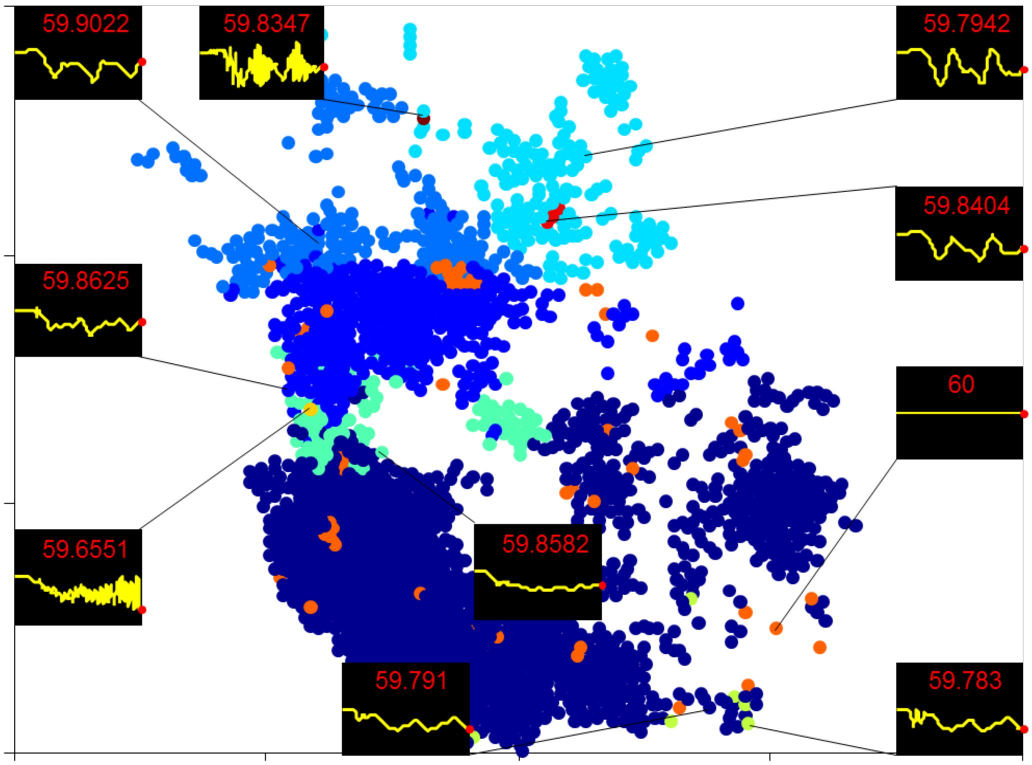

FIGURE 4.2 Example of cluster usage to group transient stability generator responses. SOURCE: Dutta and Overbye (2014). Courtesy of Thomas J. Overbye, University of Illinois, Urbana-Champaign.

Another use of clustering in the power grid is for visualization and feature extraction from large data sets (Dutta and Overbye, 2014). As an example, Figure 4.2 takes the time-domain bus frequency responses for 16,000 buses from a transient stability study, much as Figure 1.21 showed a time-domain frequency response for generators and then used clustering to group the frequency responses of the buses, showing the results in their geographic context. In the figure the circles correspond to the geographic location of the buses, with the color used to show the buses with a similar response. Sparklines show the average frequency response for each cluster. With this approach the aggregate behavior is apparent, along with identification of the small number of outlier buses.

Reliability Modeling with Physical Models

If a model of how the physical parts of mechanical equipment interact can be created, it can be used to estimate failure probabilities or time until failure. Consider an interconnected network of equipment. Each piece of equipment is represented by a random variable characterizing its hazard rate. This rate depends on neighboring equipment in the network and whether any of them has failed; in that sense, the model is Markovian. One can fit the parameters of this model to data if they are available. In the case of the power network, each piece of equipment on the power grid (substations, transformers, wind turbines, consumers) and their influence on one another would be modeled.

If the equipment has a combinatorial structure (there are components of components, or a logical structure for how one component causes another specific component to fail), this structure can be formalized into reliability block diagrams, reliability graphs, or fault trees, which are specific types of models that govern how components interact and how failures occur in the system. If one does not want to make many assumptions about the components of the equipment, Markov models of a more general form can also be used for a more data-driven approach, but possibly at the cost of a much larger set of parameters to estimate. Because these physical models completely characterize the physical system, they would allow estimates to be made about failures even when none have ever occurred; updated with data, these models become more powerful.

Cascading Failures

Analysis of the risk posed by cascading failures is a challenging problem that spans several scientific disciplines. At one end the determination of which contingencies (or, perhaps, multiple contingencies) could cause a dangerous cascade must be faced. A large-scale power system example of a cascading failure is the August 14, 2003, Northeast blackout that affected an estimate 50 million people with total costs estimated between $4 billion and $10 billion (U.S.-Canada Power System Outage Task Force, 2004).

The traditional methodology is to enumerate all possibilities of multiple element contingencies, where an individual element could be a transmission line, transformer, or generator. Let k denote the number of simultaneous element outages being considered. As introduced in Chapter 3, N – k contingency analysis seeks to determine whether there are sets with k or fewer elements that when simultaneously taken out of service could cause a cascade. This approach is computationally very expensive even for small values of k. Moreover, it may be too conservative to insist that a system endure any such k contingencies. On the other hand, this criterion may not capture complex network conditions (such as extreme weather), which can impinge on grid operation by triggering multiple contingencies at the same time.

Estimating the consequences of possible multiple contingencies and cascades could prevent those that would be most dangerous. Resources can then be concentrated on measures for preventing those judged to have the highest risk. Comprehensive analyses attempt to predict the behavior of a cascade. This is an especially complex undertaking because it takes place on multiple time scales. Phenomena such as generator tripping are subsecond events, whereas unintentional line tripping may require several minutes. Moreover, both phenomena are subject to noise and unpredictable exogenous elements such as human error and environmental influence. Race conditions between discrete events (when two events are almost simultaneous but it is necessary to determine which actually occurred first) may also need to be modeled to determine the order in which discrete events will occur. These events should be modeled as stochastic events with probability distributions associated with different levels of severity.

Data Assimilation

Dynamical models have states and parameters that must be estimated to perform simulations. Methods for performing these estimations have been extensively developed for linear systems and have become part of standard engineering practice. Data assimilation (DA) extends these methods to incorporate data from observational measurements into model predictions or estimations of states of an ongoing predictive simulation. DA has been used most extensively in weather prediction, where current and recent observations are used to recalibrate the state of large operational forecast models. The key point is that two sources of information are balanced to produce an optimal estimate. There is assumed to be an underlying model that is derived from physical principles and, from the model's viewpoint, the incorporation of data is accounting for the physics that it is missing. From the viewpoint of the data, the model is providing an extrapolation of the state to times and locations where no measurements are available. This balance between the two types of information, that derived from measurements and that from physical laws, distinguishes data assimilation from other statistical methods such as state estimation, machine learning, and data-driven modeling.

DA is applied in one of two ways: (1) sequentially, in which the state of the system is governed by an evolution equation (the physical model), which is then reinitialized at regular observation times with a new system state that is formed from a systematic interpolation between the state as forecast by the model and the observational data available at that time; and (2) retroactively, where a new state of the system over an entire period of time is estimated as an interpolation between the model state and all available observational data during that period. The former approach is known as sequential DA and the latter as *reanalysis*. Both approaches rely heavily on ideas from dynamical systems, and the research at the interface of DA and dynamic stability is particularly active and promising.

There is extraordinary promise in applying DA to issues in the electric grid, because it shares with other areas—namely, weather and climate—where DA has developed a strong underpinning in physical laws combined with considerable observational data. Despite this fact, it seems that DA has had little application, arguably because it is hard to acquire good data that are not proprietary. Nevertheless, its promise as a powerful tool indicates that efforts to develop it for grid models will be bear significant fruit. This can be achieved using partial, available, and/or synthetic data. The fundamental questions can be addressed using data that mimic the actual grid data.

In DA generally, the issues that are receiving most attention are related to the natural tension that occurs between nonlinearity and dimension. In weather and climate modeling, dimensions of the state space are particularly large and growing larger with increasing computational capacity because they are primarily determined by the resolution of the computational model being used. Methods of DA that are effective in these dimensions typically have to make compromises that involve some level of linearization. Nonlinear features of the model can then either be missed or, worse still, cause breakdown of the assimilation process.

The issue of nonlinearity will arise naturally in applying DA in the power system because the underlying equations are highly nonlinear and discontinuous. Indeed, this will be a much more serious issue than in forecasting weather, and the compromises made there may be completely inappropriate for grid predictions and estimates. Moreover, an optimal estimate, as is often sought in DA, will not suffice, because the uncertainties will be as important. A full Bayesian approach would then be necessary, and the issue shifts to whether Gaussian approximations will work.

The ideas of dynamical systems are coming into DA as a way to reduce the dimensions of the necessary calculations. An example is the use of Lyapunov vectors, which capture the most significant unstable directions. With all that is understood about the underlying grid equations as systems of coupled oscillators, this is a particularly promising direction to pursue.

The assimilation of different types of data that are manifest in forms other than state variables is also a challenge to current DA methods. This observational space may play a similar role in determining the dimension reduction. There is considerable structure in the physical model that will both aid and challenge the systematic development of DA strategies. For instance, DA has been mainly developed for large computational models based on fluid dynamics and not for networks. Adapting to this context is, in itself, a challenge and will likely indicate approaches to effective dimension reduction.

COMPLEXITY AND MODEL REDUCTION IN THE TIME OF BIG DATA

Realistic modeling involves large numbers of equations and degrees of freedom improving spatial and temporal resolution of the model, incorporating more physical effects, and describing uncertainty and noise. What is considered "large" in a model has progressed from thousands to tens and even hundreds of millions of variables as computational power has been systematically expanding. Yet there is an unwritten hypothesis of "effective simplicity"—the assumption that an expert is capable of selecting a drastically reduced number of key observables that are sufficient to make useful predictions and to take crucial actions. Model reduction is ubiquitous in applied mathematics: Making a problem as simple as possible, but not simpler, permeates our entire education as modelers. The mathematical underpinnings of model reduction (such as separation of time scales or averaging) have given rise to tools and techniques (quasi-steady-state approximation for reducing chemical kinetics, homogenization for the description of heterogeneous materials) that dramatically affect the ability to simulate, design, control, and rationalize the behavior of complex dynamical systems. Aggregation of degrees of freedom (e.g., of the many loads in a household within a distribution system, or reducing the effects of a distribution system on a transmission network to a single effective degree of freedom), so that coarse-grained, effective grid models can be usefully analyzed, is already common practice.

The ability to systematically derive useful reduced models has traditionally, however, been very restricted: The symbolic computation of, say, center manifolds for dynamical systems, while conceptually powerful, becomes quickly intractable, and an "expert" is almost by definition someone who can extract meaningful reduced models and the understanding that comes with them from extensive personal experience/wisdom.

High-performance computing, on the one hand, and data mining/machine learning on the other, are gradually (but at an accelerating pace) changing the way this expert knowledge is obtained. Useful reduced models are now increasingly obtained on the fly as a by-product of detailed scientific computation and/or by the mining of resolved computational/observational data. And while the resulting models and computations may not reflect the intuition or experience of human experts, they can be reliably tested and used. It is remarkable that many large-scale scientific computation techniques, from multigrid to preconditioning to Krylov integrator methods, implicitly rely on this online reduction process to successfully accelerate computation.

Instead of searching for reduction in the structure of dynamic models (e.g., looking for fast equilibrating processes so that they can be modeled at quasi-equilibrium and thus reduce the problem size), it is now possible to search for low dimensionality and reduction directly in both observational and simulation data. Principal component analysis held pride of place in data reduction for a hundred years; then, in the 1980s, neural networks briefly held unfulfilled promise before coming back as deep learning today. The manifold learning tools that have arisen in the last two decades (from ISOMAP to locally linear embedding to kernel principal component analysis to diffusion maps and their extensions) have used semidefinite programming relaxations to open a new window in the way effective simplicity can be discovered and algorithmically perform model reduction as a wrapper around detailed simulation codes. Smart precomputation and tabulation techniques, such as in situ adaptive tabulation and intrinsic low-dimensional manifolds, along with the associated fast database search/access, further exploit these new tools to accelerate computation. Instead of trying to derive smart reduced models (that is, deduce/discover the truly important observable quantities and then express the dynamics in terms of these quantities), detailed evolution models are written with the best physics available. These models are used in a sense as computational experiments: Instead of running them for all times, all feasible parameter values, and all reasonable initial conditions, one can run them in brief simulation bursts, process the resulting data to discover online their local useful reduction, and use this information to design the next informative short simulation burst. In effect, one is implementing singular perturbation or averaging techniques computationally rather than symbolically. The mathematics of reduction are the same as when these models can be explicitly performed on formulas, but now the reductions are based on computational observations.

In this spirit, detailed grid simulations can be effectively reduced online so that computations to identify stationary states, stability results, or control designs become tractable. This online reduction in complexity relies first on mining of computational or observational data to discover the right local variable aggregation and then on brief bursts of simulation by the disaggregated simulator to extract new, useful aggregated information. In effect, one solves the reduced model without ever writing it down in closed form and finds the relevant coarse-grained variables without necessarily being able to articulate their description.

This type of effective modeling, which crucially relies on machine learning for the detection of the important observables parameterizing the high-dimensional state data, holds promise for many complex but effectively simple systems models, especially for the power grid. A nontrivial twist arises from the fact that some of the data that one wishes to effectively compress are not just real numbers (components of a vector in $\mathbf{R}^n$ for very large n) but are also integers (such as the adjacency matrices defining the connectivities of networks of units). In this case, mining data in the form of networks/graphs rather than just as points in high-dimensional spaces becomes a challenge. What is also a challenge is the issue of comprehensibility to human experts: discovering descriptions that are isomorphic to the mathematically obtained descriptions yet are easily interpretable by humans.

The promise of this approach lies not in better modeling but in smarter, faster, possibly more easily understandable extraction of relevant information from the best written models. One of its main strengths is its "wrapper" philosophy: Since we are wrapping postprocessing algorithms (as accelerators/enablers) around our best simulators, we can still make improvements to those simulators—reflecting our best understanding of physical processes and details of simulator dynamics—and it won't require changes to the wrapper algorithms. As a matter of fact, part of the original motivation in devising this computational enabling technology was the desire to force legacy codes to perform tasks they were not intrinsically designed to do. As such, these algorithms can be wrapped around existing validated simulators, and they enhance the speed and reliability of extracting information from the simulations without needing to validate them anew.

UNCERTAINTY QUANTIFICATION

Viewing the power system as a complex, many-degree-of-freedom, nonlinear dynamical system gives rise to a hierarchy of model formulations, each with its own uses in both the design and operation phases. Deterministic steady-state formulations (giving rise to large sets of algebraic equations) and deterministic time-dependent formulations (giving rise to large sets of nonlinear differential, or differential algebraic equations) have been, and continue to be, useful. Still, taking into account and successfully modeling the inherently stochastic nature of the grid operation is increasingly the driver in pushing the boundary of the state of the art both in mathematical modeling and in algorithm development for grid design and operation.

There are aspects of uncertainty that have always been present in grid modeling and cannot be prescribed or measured accurately in advance. The uncertainty of grid loads is ever present, but the uncertainty inherent in renewable energy production, especially in solar panels and wind turbine farms owing to weather variability, poses a new set of challenges. Furthermore, the anticipated real-time interaction of renewable energy generation with energy pricing (especially as storage technology options develop) brings an additional level of complexity.

Fortunately, there is an explosion in the mathematics and computation of uncertainty quantification driven by many different natural, technological, and financial problems (weather prediction being a crucial driver). Once deterministic dynamical systems tools became more mature, uncertainty and stochasticity became the natural research frontier. It might be useful to make a distinction between uncertain parameters (unknown fixed values, for which a distribution is known, estimated, or postulated) and uncertain, time-dependent processes. For uncertain parameters, techniques based on Wiener's polynomial chaos (PC) expansion—expanding the uncertain solution of a given problem as a function of the uncertain parameter(s), using as basis functions orthogonal polynomials constructed from the parameter uncertainty distribution—have experienced a broad resurgence over the last two decades. Used for linear uncertain problems for many years (especially in civil engineering), the approach is now being used in nonlinear problems across many sciences. It has been implemented in publicly available software and has driven important numerical developments in the use of sparse (Smolyak) grids in what are called non-intrusive computational methods, circumventing the stochastic Galerkin approximation through high-dimensional collocation. Indeed, if the distributions of uncertainty of different parameters are independent, then a relatively large number (O(100)) of uncertain parameters can be practically modeled. Such Wiener PC tools can also be used in the case of uncertain processes (e.g., in stochastic partial differential equation with Brownian motion forcing). Alternative approaches, such as analysis of variance techniques (including adaptive ones) and the optimal uncertainty quantification framework, are being developed and find applications in the context of the power system (e.g., in quantifying uncertainty for large-scale dynamic simulations of power systems).

It is crucial to recognize, however, that the number of uncertain parameters/stochastic variables that enter modern power system models is, and will remain, beyond what one can usefully compute, and that therefore reduction techniques (and, importantly, data-driven reduction techniques as well as state estimation and particle filtering techniques) will play an important role in making computational uncertainty quantification sufficiently manageable as to be practically useful. This will involve building smart reduced parameterizations of uncertainty in dynamic models, and methods for estimating probability density functions or covariance matrices for the forcing terms in the models based on finite measurements—a context where, for example, multidimensional Gaussian processes and Bayesian approaches become relevant.

But while parameterizing uncertainty in a model (both for the inputs, and the loads, and the outputs) is crucial, it is only the first step necessary for formulating the problem. The most important component is the development of simulation/optimization algorithms that will actually solve the problem (plan energy production and decide dynamically on net interchange schedules and allocations). These will operate at the hardware design level (constituting the network, building in the reconfigurability/safeguard capabilities that will guarantee safe/acceptable/optimal operation) and also, more challengingly, at the day-to-day and hour-to-hour operation planning level.

The computational scalability of stochastically optimizing the design and operation of complex energy systems is an important frontier; usefully decomposing stochastic optimization problems on a scenario basis and successfully distributing scenarios/sampling strategies over processors hold promise for solving, for example, stochastic economic dispatch problems that use hourly wind forecasts to integrate wind power with the real-time energy market.

Ultimately, the biggest challenge will be at the confluence of power system hardware, renewable energy production uncertainty, and interactions with real-time prices as renewable energy production becomes more and more integrated in the grid. The relation between, say, day-ahead and expected real-time prices, and the ways this may bias economic incentives, induces an additional uncertainty quantification/uncertainty propagation layer that makes large-scale problems simply unsolvable today.

REFERENCES

Andersson, G. 2004. *Modelling and Analysis of Electric Power Systems*. Power Systems Laboratory, ETH, Zurich. March.

Armel, K.C., A. Gupta, G. Shrimali, and A. Albert. 2012. "Is Disaggregation the Holy Grail of Energy Efficiency? The Case of Electricity." Technical Paper Series: PTP-2012-05-1. Precourt Energy Efficiency Center, Stanford University, Calif. Published in 2013 in *Energy Policy* 52(C):213-234.

Bard, J.F., and J.T. Moore. 1990. A branch and bound algorithm for the bilevel programming problem. *SIAM Journal on Scientific and Statistical Computing* 11(2):281-292.

Belotti, P., C. Kirches, S. Leyffer, J. Linderoth, J. Luedtke, and A. Mahajan. 2013. Mixed-integer nonlinear optimization. *Acta Numerica* 22:1-131.

Benders, J.F. 1962. Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik*, 4:238–252.

Ben-Tal, R., and A. Nemirovski. 2000. Robust solutions of linear programming problems contaminated with uncertain data. *Mathematical Programming* 88:411-424.

Bertsimas, D., D.B. Brown, and C. Caramanis. 2011. Theory and applications of robust optimization. *SIAM Review* 53(3):464-501.

Bienstock, D. 2013. Progress on solving power flow problems. *Optima Mathematical Optimization Society Newsletter* 93:1-7.

Bienstock, D., M. Chertkov, and S. Harnett. 2014. Chance-constrained optimal power flow: Risk aware network control under uncertainty. *SIAM Review* 56(3):461-495.

Birge, J.R., and F. Louveaux. 2011. *Introduction to Stochastic Programming*. Springer Series in Operations Research and Financial Engineering. Springer-Verlag, New York, N.Y.

Bixby, R.E. 2015. "Mixed-Integer Programming: It Works Better Than You May Think." Presentation to the National Research Council workshop Analytical Research Foundations for the Next-Generation Electric Grid on February 11-12, 2015, in Irvine, Calif.

Cain, M., R.P. O'Neill, and A. Castillo. 2013. History of optimal power flow and formulations, Optimal power flow paper 1. Federal Regulatory Commission.

Castillo, A. and R.P. O'Neill. 2013a. Survey of approaches to solving the ACOPF, Optimal power flow paper 4. Federal Regulatory Commission.

Castillo, A. and R.P. O'Neill. 2013b. Computational performance of solution techniques applied to the ACOPF, Optimal power flow paper 5. Federal Regulatory Commission.

Darby, S. 2006. *The Effectiveness of Feedback on Energy Consumption. A Review for DEFRA of the Literature on Metering, Billing, and Direct Displays*. Technical report. Environmental Change Institute, University of Oxford. April.

Dobson, I. 1992. Observations on the geometry of saddle node bifurcation and voltage collapse in electrical power systems. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications* 39(3).

Dutta, S., and T.J. Overbye. 2014. Feature extraction and visualization of power system transient stability results. *IEEE Transactions on Power Systems* 29(2):966-973.

Ferris, M.C., Z. Friedman, D. Molzahn, C. DeMarco, and B. Lesieutre. ACOPF models: Extending data, formulations, and solution methodology, presentation given June 2013. See http://pages.cs.wisc.edu/~ferris/talks/ferc-jun.pdf. Accessed December 1, 2015.

Goldberg, M. 2015. Presentation to the National Research Council Workshop on Analytical Research Foundations for the Next-Generation Electric Grid in February in Irvine, Calif.

Glover, J.D., M.S. Sarma, and T.J. Overbye. 2012. *Power System Analysis and Design*, Fifth Edition. Cengage Learning.

Guckenheimer, J., and P. Holmes. 1983. *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields (Applied Mathematical Sciences)*, 1st ed. Springer-Verlag, New York, N. Y.

Hairer, E., and G. Wanner. 2004. *Solving Ordinary Differential Equations (I)*. Springer-Verlag, New York, N. Y.

Hairer, E., and G. Wanner. 2009. *Solving Ordinary Differential Equations (II)*. Springer-Verlag, New York, N. Y.

Hart, G. 1992. Nonintrusive appliance load monitoring. *Proceedings of the IEEE* 80(12):1870-18912.

Hedman, K.W., S.S. Oren, and R.P. O'Neill. 2011. "A Review of Transmission Switching and Network Topology Optimization." IEEE PES General Meeting, Detroit.

Heidel, T. 2014. "Advanced Research Projects Agency-Energy," Grid Optimization Workshop Introduction. Presentation to the Advanced Bulk Power System Optimization Technologies Workshop on November 13-14, 2014, U.S. Department of Energy, http://arpa-e.energy.gov/?q=workshop/advanced-bulk-power-system-optimization-technologies-workshop.

Higle, J.L., and S. Sen. 1991. Stochastic decomposition: An algorithm for two-stage linear programs with recourse. *Mathematics of Operations Research* 16:650-669.

Kuznetsov, Y.A. 2004. *Numerical Analysis of Bifurcations*. Springer, New York.

Linderoth, J.T., A. Shapiro, and S.J. Wright. 2006. The empirical behavior of sampling methods for stochastic programming. *Annals of Operations Research* 142:219-245.

Low, S. 2014. Convex relaxation of optimal power flow, Part I: Formulations and equivalence. *IEEE Transactions on Control of Network Systems* 1(1):15-27.

Martinussen, T., and T.H. Scheike. 2006. *Dynamic Regression Models for Survival Data*. Springer, New York, N.Y.

NASPI. 2015. *NASPI 2014 Survey of Synchrophasor System Networks—Results and Findings*. Technical report. Data and Network Management Task Team, Network Systems Group. July. https://www.naspi.org/documents.

Neenan, B., and J. Robinson. 2009. *Residential Electricity Use Feedback: A Research Synthesis and Economic Framework*. Technical report. Electric Power Research Institute.

Nemhauser, G.L., and L.A. Wolsey. 1998. *Integer and Combinatorial Optimization*. Wiley, New York, N.Y.

Nocedal, J., and S. Wright. 2006. *Numerical Optimization*, 2nd ed. Springer.

Oliveira, W., C. Sagastizbal, and S. Scheimberg. 2011. Inexact bundle methods for two-stage stochastic programming. *SIAM Journal on Optimization* 21:517-544.

Papavasiliou, A. and S.S. Oren. 2013. A comparative study of stochastic unit commitment and security-constrained unit commitment using high performance computing. *Proceedings of the 2013 European Control Conference (ECC)*, pp. 2507-2512.

PJM. 2014. *Load Forecasting and Analysis*, Manual 19, Revision 26. Effective date November 1. Prepared by Resource Adequacy Planning.

Prékopa, A. 1995. *Stochastic Programming*. Kluwer.

Rockafellar, R.T. 1993. Lagrange multipliers and optimality. *SIAM Review* 35(2):183-238.

Schrijver, A. 1998. *Theory of Linear and Integer Programming*. Wiley, London, U.K.

Shapiro, A., D. Dentcheva, and A. Ruszczyński. 2009. *Lectures on Stochastic Programming: Modeling and Theory*. MOS-SIAM Series on Optimization. Society for Industrial and Applied Mathematics. http://dx.doi.org/10.1137/1.9780898718751.

Sheble, G.B., and G.N. Fahd. 1994. Unit commitment literature synopsis. *IEEE Transactions on Power Systems* 9(1):128-135.

Shiina, T., and J.R. Birge. 2004. Stochastic unit commitment problem. *International Transactions in Operational Research* 11:19-32.

Soliman, S., and A. Kandari. 2010. Electrical Load Forecasting: Modeling and Model Construction. Elsevier. ISBN 978-0-12-381543-9.

Sultanem, F. 1991. Using appliance signatures for monitoring residential loads at meter panel level. *IEEE Transaction on Power Delivery* 6(4):1380-1385.

Trivedi, K. 2001. *Probability and Statistics with Reliability, Queueing, and Computer Science Applications*, 2nd ed. Wiley, New York, N.Y.

U.S.-Canada Power System Outage Task Force. 2004. *Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations*. April. http://www.nerc.com/docs/docs/blackout/ch1-3.pdf.

Van Slyke, R., and R. J-B. Wets. 1969. L-shaped linear programs with application to optimal control and stochastic programming. *SIAM Journal of Applied Mathematics* 17:638–663.

Wang, Q., Y. Guan, and J. Wang. 2012. A chance-constrained two-stage stochastic program for unit commitment with uncertain wind power output. *IEEE Transactions on Power Systems* 27:206-215.

Wehenkel, L.A. 1998. *Automatic Learning Techniques in Power Systems*. Kluwer International Series in Engineering and Computer Science. Kluwer Academic Publishers, Norwell, Mass.

Wischkaemper, J.A., C.L. Benner, B.D. Russell, and K.M. Manivannan. 2014. "Waveform Analytics-based Improvements in Situational Awareness, Feeder, Visibility, and Operational Efficiency," *Proceedings of the 2014 IEEE PES T&D Conference and Exposition*. doi:10.1109/TDC.2014.6863349.

Wischkaemper, J.A., C.L. Benner, B.D. Russell, and K. Manivannan. 2015. Application of waveform analytics for improved situational awareness of electric distribution feeders. *IEEE Transactions on Smart Grid* 6(4):2041-2049.

Wolsey, L.A. 1998. *Integer Programming*. Wiley, New York, N.Y.

Zeifman, M., and K. Roth. 2011. Nonintrusive appliance load monitoring: Review and outlook. *IEEE Transactions on Consumer Electronics* 57(1):76-84.

5

Preparing for the Future

INTRODUCTION

Planning the foundations of analytic research for the next-generation electric grid requires consideration of how the grid is likely to change in the future. That is the focus of this chapter. However, in preparing for the future it is prudent to contemplate a quote from Winston Churchill: "It is always wise to look ahead, but difficult to look further than you can see."[1] The future is always uncertain, and in planning foundational research it would be a mistake to think that the future is just some extrapolated view of the present. There will always be disruptive technologies, as central station electricity itself was disruptive to the economy of the 1880s. Hence the focus of this chapter is not to try to predict a single most likely scenario, but rather to explore the range of uncertainties that could unfold. The ultimate goal is to present research foundations that can future-proof the grid so that regardless of how the grid evolves, the United States is prepared.

UNCERTAINTY IN WHAT LIES AHEAD

The grid of today is changing with the rapid integration of renewable energy resources such as wind and solar photovoltaic (PV) and the retirement of substantial amounts of coal generation. For example, in early 2015 in the United States, there was installed capacity of about 65 GW of wind and 9 GW of solar PV (out of a total of 1,070 GW), from less than 3 GW of wind and 0.4 GW of solar just 15 years back (EIA, 2009). However, this needs to be placed in context by noting that during the natural gas boom in the early 2000s, almost 100 GW of natural gas capacity was added in just 2 years! And solar thermal, which seemed so promising in 2009, has now been mostly displaced by solar PV because of dropping prices for the PV cells. Further uncertainty arises because of the greater coupling of the electric grid to other infrastructures such as natural gas, water, and transportation. Finally, specific events can upset the best predictions. An example is the Japanese tsunami in 2011, which (among other factors) dimmed the prospects for a nuclear renaissance in the United States and elsewhere.

Some of the uncertainty currently facing the industry is illustrated in Figure 5.1. The drivers of this uncertainty are manifold: (1) cyber technologies are maturing and are becoming available at reasonable cost—these include sensing, such as phasor measurement units (PMUs), communications, control, and computing; (2) emergence of

[1] National Churchill Museum, "Winston Churchill and the Cold War," https://www.nationalchurchillmuseum.org/winston-churchill-and-the-cold-war.html. Accessed September 15, 2015.

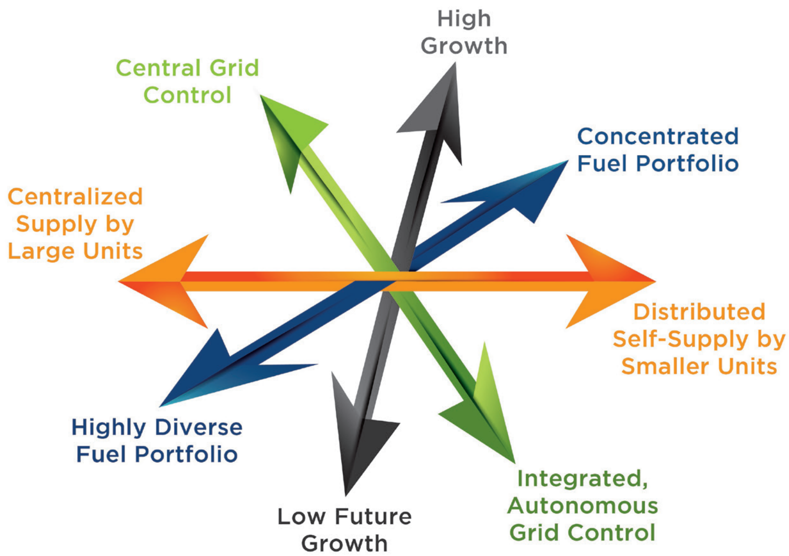

FIGURE 5.1 Electric grid uncertainty. SOURCE: Courtesy of PJM Interconnection.

qualitatively new resources, such as renewable distributed energy resources (DERs)—PVs, wind generation, geothermal, small hydro, biomass, and the like; (3) new quest for large-scale storage—stationary batteries, as well as low-cost storage batteries such as those for use in electric vehicles; (4) changing transmission technologies such as increased use of flexible ac transmission system (FACTS) technologies and/or increased use of high-voltage direct current (HVDC) lines and the integration of other dc technologies; (5) environmental objectives for reducing pollutants; (6) industry reorganization, from fully regulated to service-oriented markets; and (7) the need for basic electrification in developing countries, which affects the priorities of equipment suppliers. Given these drivers, it is hard to predict exactly long-term power grid scenarios. However, to help future-proof the grid, the committee offers the advice that follows in the remaining sections of this chapter.

TECHNOLOGIES THAT WILL ENHANCE THE OBSERVABILITY OF THE GRID

Since the advent of the electric power grid, measurement technologies have been a necessary component of the system for both its protection and its control. For example, measuring the currents flowing in the power system wires and the bus voltages are two key quantities of importance. The currents are measured using current transformers, which convert the magnetic field of the primary circuit to a proportionally smaller current suitable for input to instrumentation. The voltages are measured using potential transformers (PTs), which utilize traditional transformer technology of two windings coiled on a common magnetic core to similarly proportionally reduce the line voltage to a voltage suitable for instrumentation. Through the middle of the 20th century higher voltages and coupled capacitive voltage transformers used capacitors as a voltage divider as a more practical alternative to a PT for extra-high-voltage transmission. Other instruments exploiting either the electric or the magnetic fields have been developed. More recently, optical sensors can convert the voltages and currents as a directly measured quantity (Niewczas and McDonald, 2007).

Bringing these measurements to a central location has been possible for many decades. Technologies such as Supervisory Control and Data Acquisition (SCADA) use specialized protocols to transmit the information gathered in substations through analog-to-digital conversion in various sensors that are directly connected to remote terminal units (RTUs). A typical SCADA architecture exchanges both measurement and control information between the front end processor in the control center and the RTUs in the substations. Modern SCADA protocols support reporting of exceptions in addition to more traditional polling approaches. These systems are critical to providing control centers with the information necessary to operate the grid and to providing control signals to the various devices in the grid to support centralized control and optimization of the system.

SCADA systems in use today have two primary limitations. First, they are relatively slow. Most systems poll once every 4 sec, with some of the faster implementations gathering data at a 2-sec scan rate. Second, they are not time synchronized. Often, the data gathered in the substation and passed to the central computer are not time-stamped until they are registered into the real-time database at the substation. And as the information is gathered through the polling cycle, sometimes there can be a difference between the pre- and postevent measurements if something happens during the polling cycle itself.

First described in the 1980s (Phadke et al., 1983), the PMUs mentioned in earlier chapters utilize the precise time available from systems such as the Global Positioning System. The microsecond accuracy available is reasonable for the accurate calculation of phase angles of various power system quantities. More broadly, high-speed time-synchronized measurements are broadly referred to as wide area measurement systems. These underwent significant development beginning in the 1990s and can now provide better measurements of system dynamics with typical data collection rates of 30 or more samples per second. Significant advances in networking technology within the past couple of decades have enabled wide area networks by which utilities can share their high-speed telemetry with each other, enabling organizations to have better wide area situational awareness of the power system. This is addressing one of the key challenges that was identified and formed into a recommendation following the August 14, 2003, blackout (U.S.-Canada Power System Outage Task Force, 2004).

There are several benefits of wide area measurement systems. First, because of the high-speed measurements, dynamic phenomena can be measured. The 0.1- to 5-Hz oscillations that occur on the power system can be compared to simulations of the same events, leading to calibration that can improve the power system models. It is important to have access to accurate measurements corresponding to the time scales of the system. Second, by providing a direct measure of the angle, there can be a real-time correlation between observed angles and potential system stress.

The measurements from PMUs, known as synchrophasors, can be used to manage off-normal conditions such as when an interconnected system breaks into two or more isolated systems, a process known as "islanding." For example, during Hurricane Gustav, in September 2008, system operators from Entergy (the electric utility company serving the impacted area in Louisiana) were able to keep a portion of the grid that islanded from the rest of the Eastern Interconnection operating after the storm damage took all of the transmission lines out of service, isolating a pocket of generation and load. The isolated area continued to operate by balancing generation and load. The system operators credited synchrophasor technology with allowing them to keep this island operational during the restoration process (NERC, 2010).

Researchers are looking at PMU data to expedite resolution of operating events such as voltage stability and fault location and to quickly diagnose equipment problems such as failing instrument transformers and negative current imbalances. More advanced applications use PMU data as inputs to the special protection systems or remedial action schemes, mentioned in Chapter 3 for triggering preprogrammed automated response to rapidly evolving system conditions.

All telemetry is subject to multiple sources of error. These include but are not limited to measurement calibration, instrumentation problems, loss of communications, and data drop-outs. To overcome these challenges, state estimation, introduced in Chapter 3, is used to compute the real-time state of the system. This is a model-fitting exercise, whereby the available data are used to determine the coefficients of a power system model. A traditional state estimator requires iteration to fit the nonlinear with the available measurements. With an overdetermined set of measurements, the state estimation process helps to identify measurements that are suspected of being inaccurate. Because synchrophasors are time aligned, a new type of linear state estimator has been developed and is now undergoing widespread implementation (Yang and Bose, 2011). The advantage of "cleaning" the measurements through a linear state estimator is that the application is not subject to the data quality errors that can occur with the measurement and communications infrastructure. Additional advances are under way, including distributed state estimation and dynamic state estimation.

One of the more recent challenges has been converting the deluge of new measurements available to a utility, from synchrophasors and other sources, into actionable information. Owing to the many more points of measurement available to a utility from smart meters and various distribution automation technologies, all organizations involved in the operation of the electric power grid are faced with an explosion of data and are grappling with techniques to utilize this information for making better planning and/or operational decisions. Big data analytics is being called on to extract information for enhancing various planning and operational applications.

One such challenge includes the improved management of uncertainty. Whether it be the uncertainty associated with estimating future load or generation availability or the uncertainty associated with risks such as extreme weather or other natural or manmade disaster scenarios that could overtake the system, more sophisticated tools for characterizing and managing this uncertainty are needed.

Better tools to provide more accurate forecasting are also needed. One promising approach is through ensemble forecasting methods, in which various forecasting methods are compared with one another and their relative merits used to determine the most likely outcome (with appropriate confidence bounds). One such example is an ensemble-based Bayesian model averaging technique (Vlachopoulou et al., 2013).

Finally, better decision support tools, including intelligent alarm processors and visualization, are needed to enhance the reliability and effectiveness of the power system operational environment. Better control room automation over the years has provided an unprecedented increase in the effectiveness with which human operators handle complex and rapidly evolving events. During normal and routine situations, the role of the automation is to bring to the operator's attention events that need to be addressed. However, during emergency situations, the role of the automation is to prioritize actions that need to be taken. Nevertheless, there is still room for improving an operator's ability to make informed decisions during off-normal and emergency situations. More effective utilization of visualization and decision-support automation is still evolving, and much can be learned by making better use of the social sciences and applying cognitive systems engineering approaches.

TECHNOLOGIES THAT WILL ENHANCE THE CONTROLLABILITY OF THE GRID

The value of advanced analytics is only as good as our ability to effect change in the system based on the result of those analytics. Whether it is manual control with a human in the loop or automated control that can act quickly to resolve an issue, effective controls are essential. The power system today relies on the primary, secondary, and tertiary hierarchical control strategies that were introduced in Chapter 4 to provide various levels of coordinated control. This coordination is normally achieved through temporal and spatial separation of the various controls that are simultaneously operating. For example, high-speed feedback in the form of proportional-integral-derivative controls operates at power plants to regulate the desired voltage and power output of the generators. Supervisory control in the form of set points (e.g., maintain this voltage and that power output) is received by the power plant from a centralized dispatcher. Systemwide frequency of the interconnected power system is accomplished through automatic generation control, which calculates the desired power output of the generating plants every 4 sec.

Protection schemes that are used to isolate faults rely on local measurements to make fast decisions, supplemented by remote information through communications to improve the accuracy of those decisions. Various teleprotection schemes and technologies have been developed over the past several decades to achieve improved reliability by leveraging available communications technologies. In addition, microprocessor-based protective relays have been able to improve the selectivity and reliability of fault isolation, including advanced features such as fault location. One example is the ability to leverage traveling wave phenomena that provide better accuracy than traditional impedance-based fault location methods (IEEE, 2015).

All of these methods described above have one thing in common: judicious use of communications. For historical reasons, when communications were relatively expensive and unreliable, more emphasis was placed on local measurements for protection and control. Communications were used to augment this local decision making. With the advent of more inexpensive (and reliable) communication technologies, such as fiber-optic links installed on transmission towers, new distributed control strategies are beginning to emerge. Additionally, classical control approaches are being challenged by the increased complexity of distribution networks, with more distributed generation, storage, demand response, automatic feeder switching, and other technologies that are dramatically changing the distribution control landscape. It will soon no longer be possible to control the power system with the control approaches that are in use today.[2]

One of the approaches being considered is more comprehensive application of market-based, real-time control signals that traverse the entire electricity delivery infrastructure, from transmission to distribution, including supply and end use. One such approach that has been proposed is "transactive energy" or, alternatively, "transactive control."[3] This approach will enable the supplier to communicate the availability of power given the various real-time constraints at the point of delivery and provide an opportunity for consumers to communicate their willingness to curtail the delivery of their power as conditions dictate.

Perhaps the biggest challenge underlying the mathematical and computational requirements for this research is the fact that any evolution from today's operating and control practices will require that newly proposed methods cannot be best-effort methods; instead, a guaranteed performance (theoretical and tested) will be required if any new methods unfamiliar to the system operators are to be deployed. Today there is very little theoretical foundation for mathematical and computational methods capable of meeting provable performance goals over a wide range of operating conditions. More specifically, to arrive at the new mathematical and computational methods needed for the power system, one must recognize that the power system represents a very large-scale, complex, and nonlinear dynamic system with multiple time-varying interdependencies. A systematic framework for modeling, defining performance objectives, ensuring control performance, and providing multidimensional optimization will be helpful.

EFFECTS OF CLIMATE CHANGE

Many of the assumptions associated with the long-term operation of the electricity infrastructure are based on climatic conditions that prevailed in the past century. Climate changes appear likely to change some of those basic planning assumptions.

If policy changes are made to mitigate carbon emissions, parallel changes to the entire power generation infrastructure and the transmission infrastructure connecting our sources of electricity supply will be necessary. This gets into institutional issues such as the availability of capital investment to accommodate these changes, and policies associated with how to recover the costs of the investments. The traditional utility business model would need to be changed to accommodate these developments (Finnigan, 2014).

If the average intensity of storms increases, or if weather events become more severe (hotter summers and/or colder winders), basic assumptions about the cost effectiveness of design trade-offs underlying the electric power infrastructure would need to be revisited. Examples of this are the elements for hardening the system against wind or water damage, the degree of redundancy that is included to accommodate extreme events, and the extent to which dual-fueled power plants are required to minimize their dependency on natural gas.

MATHEMATICAL AND COMPUTATIONAL CHALLENGES IN GRID ARCHITECTURES

At present, the system is operated according to practices whose theoretical foundations require reexamination. In one such practice, industry often uses linearized modes in order to overcome nonlinear temporal dynamics. For example, local decentralized control relies on linear controls with constant gain. While these designs are simple and straightforward, they lack the ability to adapt to changing conditions and are only valid over the range of operating conditions that their designers could envision. If the grid is to operate in a stable way over large ranges of disturbances or operating conditions, it will be necessary to introduce a systematic framework for deploying more sensing and control to provide a more adaptive and nonlinear dynamics-based control strategy. Similarly, to overcome nonlinear spatial complexity, the system is often modeled assuming weak interconnections of subsystems with stable and predictable boundary conditions between each, while assuming that only fast controls are

[2] Hawaiian Electric Company, Inc., "Issues and Challenges," http://www.hawaiianelectric.com/heco/Clean-Energy/Issues-and-Challenges. Accessed December 1, 2015.

[3] Pacific Northwest SMART GRID Demonstration Project, "Our Electricity System Is Changing," last modified April 2015, http://www.pnwsmartgrid.org/transactive.asp. Accessed December 1, 2015.

localized. Thus, system-level models used in computer applications to support various optimization and decision-support functions generally assume steady-state conditions subject to linear constraints. As power engineers know, sometimes this simplifying assumption is not valid.

Other open mathematical and computational challenges include integrating more nondispatchable generation in the system or other optimized adjustment of devices or control systems. These opportunities for advancing the state of the art for computing technologies could be thoughts of as “deconstraining technologies”: The nonlinear ac optimal power flow can be used to help reduce the risk of voltage collapse and enable lines to be used within the broader limits; FACTS, HVDC lines, and storage technology can be used for eliminating stability-related line limits; and so on.

The problem of unit commitment and economic dispatch subject to plant ramping rate limits needs to be revisited in light of emerging technologies. It is important to recognize that ramping rate limits result from constraints in the energy conversion process in the power plant. But these are often modeled as static predefined limits that do not take into account the real-time conditions in the actual power generating facility. This is similar to the process that establishes thermal line limits and modifies them to account for voltage and transient stability problems. As the dynamic modeling, control, and optimization of nonlinear systems mature, it is important to model the actual dynamic process of energy conversion and to design nonlinear primary control of energy conversion for predictable input-output characteristics of the power plants.

In closing, instead of considering stand-alone computational methods for enhancing the performance of the power system, it is necessary to understand end-to-end models and the mathematical assumptions made for modeling different parts of the system and their interactions. The interactions are multitemporal (dynamics of power plants versus dynamics of the interconnected system, and the role of control); multispatial (spanning local to interconnection-wide); and contextual (i.e., performance objectives). It will be necessary to develop a systematic framework for modeling and to define performance objectives and control/optimization of different system elements and their interactions.

MATHEMATICAL AND COMPUTATIONAL CHALLENGES IN LOCAL DISTRIBUTION GRID ARCHITECTURES

Today transmission and distribution are often planned and operated as separate systems. The fundamental assumption is that the transmission system will provide a prescribed voltage at the substation, and the distribution system will deliver the power to the individual residential and commercial customers. Historically, there is very little feedback between these separate systems beyond the transmission system operator needing to know the amount of power that needs to be delivered and the distribution system operator knowing what voltage to expect. It has been increasingly recognized, however, that as different types of distributed energy resources, including generation, storage, and responsive demand, are embedded within the distribution network, different dynamic interactions between the transmission and distribution infrastructure may occur. One example is the transient and small-signal stability issues of distributed generation that changes the dynamic nature of the overall power system (Donnelly et al., 1996; Cardell and Ilic, 2004; Nazari and Ilic, 2010). It will be important in the future to establish more complete models that include the dynamic interactions between the transmission and distribution systems.

In addition, there is a need for better planning models for designing the sustainable deployment and utilization of distributed energy resources. It is critical to establish such models to support the deployment of nondispatchable generation, such as solar, with other types of distributed energy resources and responsive demand strategies. To illustrate the fundamental lack of modeling and design tools for these highly advanced distribution grids, consider a small, real-world, self-contained electric grid of an island (Ilic et al., 2013). Today’s sensing and control are primarily placed on controllable conventional power plants since they are considered to be the only controllable components. Shown in Figure 5.2a is the actual grid, comprising a large diesel power plant, small controllable hydro, and wind power plant. Following today’s modeling approaches, this grid gets reduced to a power grid, shown in Figure 5.2b, in which the distributed energy resources are balanced with the load. Moreover, if renewable plants (hydro and wind) are represented as a negative predictable load with superposed disturbances, the entire island is represented as a single dynamic power plant connected to the net island load (Figure 5.2c).

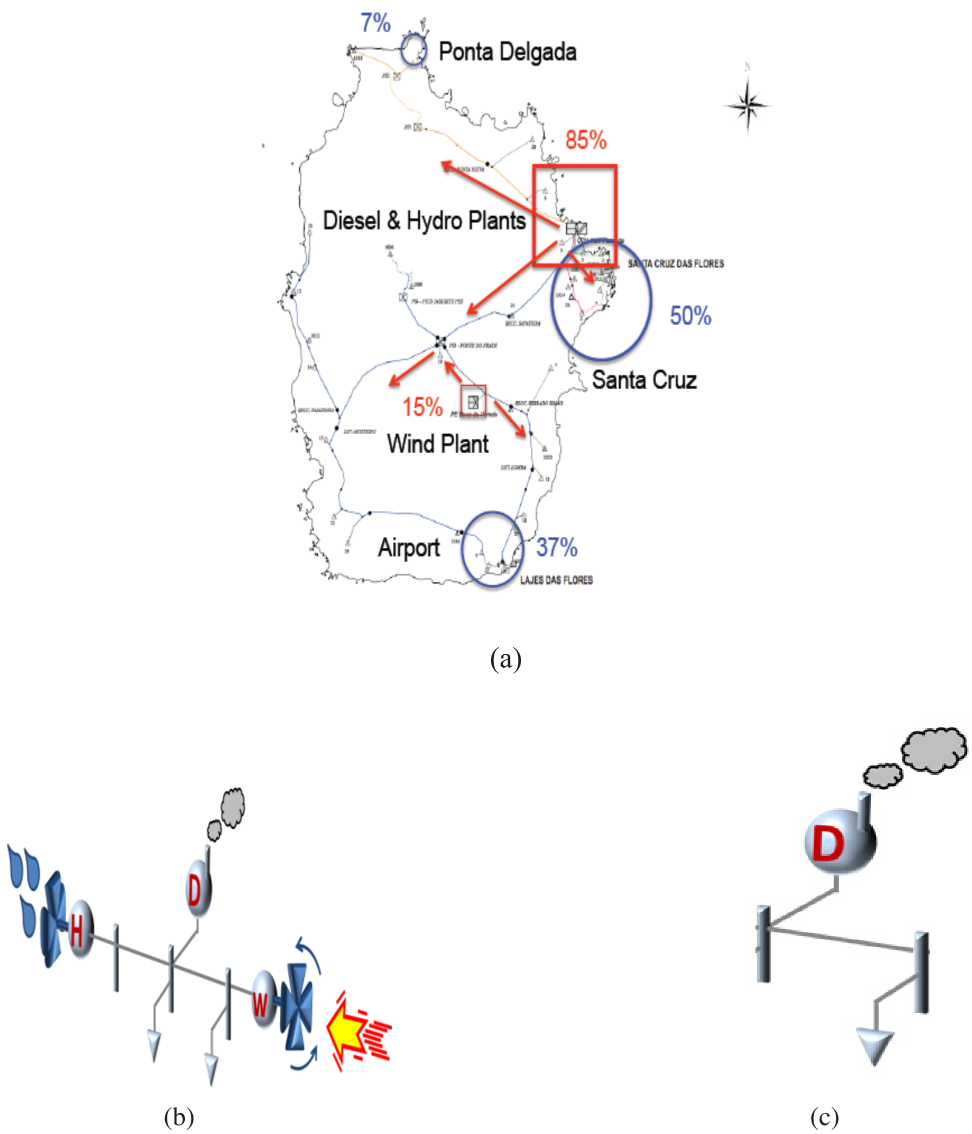

FIGURE 5.2 A representative stand-alone electric power grid (a); its dynamical model (b); and a dynamical model that represents renewable resources as negative constant power load (c). SOURCE: Ilic et al. (2013); Courtesy of Marija Ilic, Carnegie Mellon University.

In contrast with today's local grid modeling, consider the same island grid in which all components are kept and modeled (see Figure 5.3). The use of what is known as advanced metering infrastructure (AMI) allows information about the end user electricity usage to be collected on an hourly (or more frequent) basis. Different models are needed to exploit this AMI-enabled information to benefit the operating procedures used by the distribution system operator (DSO) in charge of providing reliable uninterrupted electricity service to the island. Notably, the same grid becomes much more observable and controllable. Designing adequate SCADA architecture for integrating more PVs and wind power generation and ultimately retiring the main fossil power plants requires such new models. Similarly, communication platforms and computing for decision making and automation on the island require models that are capable of supporting provable quality of service and reliability metrics. This is particularly

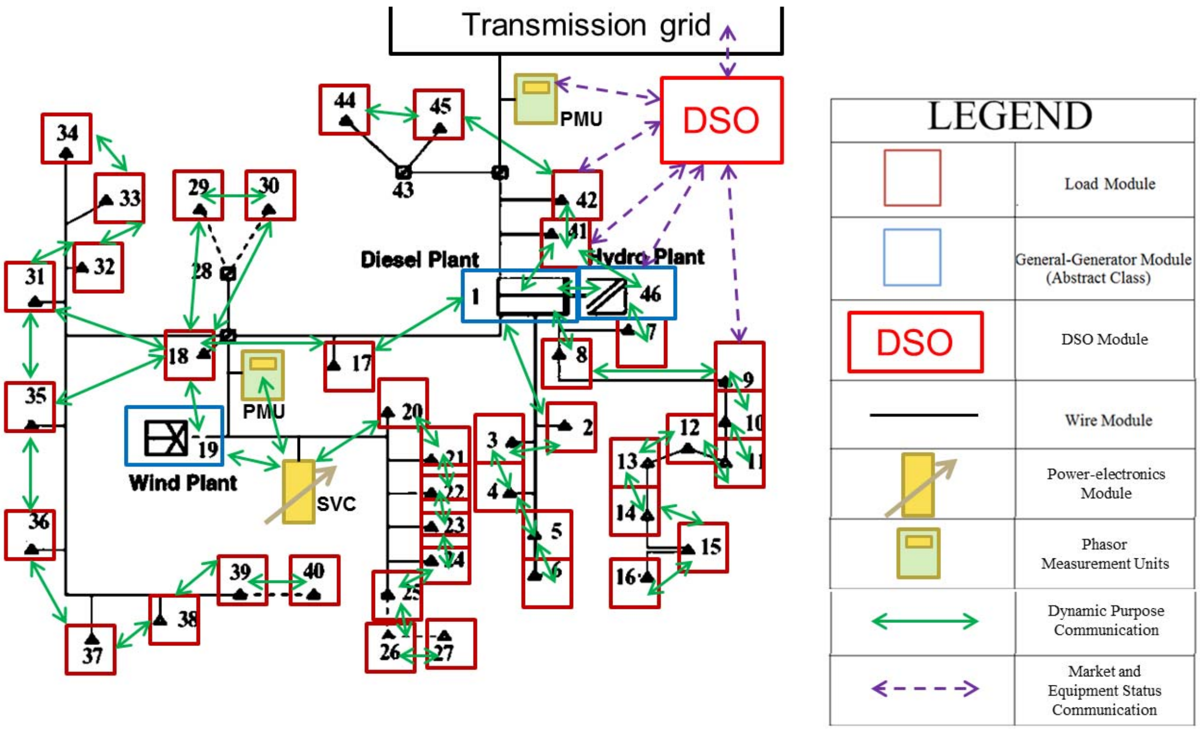

FIGURE 5.3 An island distribution grid representation for modeling and controlling DERs. SOURCE: Ilic et al. (2014). Reprinted, with permission, copyright 2014, IEEE.

important for operating the island during equipment failures and/or unexpected variations in power produced by the distributed energy resources. The isolated grid must remain resilient and have enough storage or responsive demand to ride through interruptions in available power generation without major disruptions. Full distribution automation also includes reconfiguration and remote switching.

MATHEMATICAL AND COMPUTATIONAL CHALLENGES IN MANAGING INTERDEPENDENCIES BETWEEN THE TRANSMISSION AND LOCAL DISTRIBUTION GRIDS/MICROGRIDS

Based on the preceding description of representative power grid architectures, it is fairly straightforward to recognize that different grid architectures present different mathematical and computational challenges for the existing methods and practices. These new architectures include multiscale systems that range temporally between the relatively fast transient stability–level dynamics and slower optimization objectives. They consist, as well, of nonlinear dynamical systems, where today's practice is to utilize linear approximations, and large-scale complexity, where it is difficult to completely model or fully understand all of the nuances that could occur, if only infrequently, during off-normal system conditions but that must be robustly resisted in order to maintain reliable operations at all times. In all these new architectures the tendency has become to embed sensing/computing/control at a component level. As a result, models of interconnected systems become critical to support communications and information exchange between different industry layers. These major challenges then become a combination of (1) sufficiently accurate models relevant for computing and decision making at different layers of such complex, interconnected grids, (2) sufficiently accurate models for capturing the interdependencies/dynamic interactions, and (3) control theories that can accommodate adaptive and robust distributed, coordinated control. Ultimately, advanced mathematics will be needed to design the computational methods to support various time scales of decision making, whether it be fast automated controls or planning design tools.

REFERENCES

Cardell, J., and M. Ilic. 2004. Maintaining stability with distributed generation in a restructured industry. Pp. 2142-2149 in *IEEE Power Engineering Society General Meeting*. doi:10.1109/PES.2004.1373260.

Donnelly, M.K., J.E. Dagle, D.J. Turdnowski, and G.J. Rogers. 1996. Impacts of the distributed utility on transmission system stability. *IEEE Transactions on Power Systems* 11(2):741-746.

EIA (Energy Information Administration). 2009. *Electric Power Annual*. http://www.eia.gov/electricity/annual/. Accessed December 1, 2015.

Finnigan, J. 2014. Changing times for electric utilities. *Forbes Magazine*. March 7.

IEEE. 2015. *IEEE Guide for Determining Fault Location on AC Transmission and Distribution Lines*. IEEE Standard C37.114-2014. doi:10.1109/IEEESTD.2015.7024095.

Ilic, M., L. Xie, and Q. Liu (eds.). 2013. *Engineering IT-Enabled Sustainable Electricity Services: The Tale of Two Low-Cost Green Azores Islands*. Springer, New York, N.Y.

Ilic, M., K. Bachovchin, M. Cvetkovic, X. Miao, and A. Hsu. 2014. Physics-based foundations for cyber and market design in complex electric energy systems. Pp. 4635-4654 in *IEEE Proceedings of the 53rd Annual Conference on Decision and Control (CDC)*. doi:10.1109/CDC.2014.7040113.

Nazari, M.H., and M. Ilic. 2010. Potential for efficiency improvement of future electric energy systems with distributed generation units. Pp. 1-9 in *IEEE 2010 Power and Energy Society General Meeting*. doi:10.1109/PES.2010.5589994.

NERC (North American Electric Reliability Corporation). 2010. *Real-Time Application of Synchrophasors for Improving Reliability*. Princeton, N.J. October 18. http://www.nerc.com/docs/oc/rapirtf/RAPIR%20final%20101710.pdf.

Niewczas, P., and J.R. Mcdonald. 2007. Advanced optical sensors for power and energy systems applications. *IEEE Instrument and Measurement Magazine* 10(1):18-28.

Phadke, A.G., J.S. Thorp, and M.G. Adamiak. 1983. A new measurement technique for tracking voltage phasors, local system frequency, and rate of change of frequency. *IEEE Transactions on Power Apparatus and Systems* PAS-102(5):1025-1038.

U.S.-Canada Power System Outage Task Force. 2004. *Final Report on the August 14, 2003 Blackout in the Unites States and Canada: Causes and Recommendations*. April. http://energy.gov/sites/prod/files/oeprod/DocumentsandMedia/BlackoutFinal-Web.pdf.

Vlachopoulou, M., L. Gosink, T.C. Pulsipher, T. Ferryman, N. Zhou, and J. Tong. 2013. An ensemble approach for forecasting net interchange schedule. Pp. 1-5 *2013 IEEE Power and Energy Society (PES) General Meeting*. doi:10.1109/PESMG.2013.6672760.

Yang, T., and A. Bose. 2011. Fast estimation of the state of the power grid using synchronized phasor measurements. Pp. 5960-5963 in *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. doi:10.1109/ICASSP.2011.5947719.

6

Mathematical Research Priorities Arising From the Electric Grid

INTRODUCTION

This chapter describes mathematical and computational problems where advances are expected to have significant positive impact on the electric grid. Many of the report's recommendations appear here. Previous chapters have described rapid changes to the grid that call for reexamining its architecture and operating procedures. Those changes included dramatic increases in the numbers of physical variables that must be measured and monitored; intermittent generation from renewable, distributed energy sources; and smart grids that extend control to loads as well as generation. The charge of this committee has been to identify the mathematical research that is necessary in order to enable effective management of these changes. How, then, do mathematics and computation enable what is feasible, and what limitations will remain? What are some research challenges for mathematics where timely progress will enable new technologies? Is the power industry poised for a sweeping transformation comparable to that which occurred in the telecommunications industry? Anticipating where bottlenecks occur in mathematical and computational tools will help answer these questions.

The role of optimization in the power grid industry features prominently in this report because optimization algorithms have become a central aspect of wholesale electricity markets, and the limitations of these algorithms slow progress in accomplishing more. Procedures for these markets would be significantly improved if there were a robust, efficient solution to the alternating current optimal power flow (ACOPF) problem discussed throughout this report. But finding such a solution requires either fundamental advances in general algorithms for nonlinear, nonconvex optimization problems or insights that rely on the network structure of the power grid.

The electricity markets pose challenges that go beyond improvements in optimization algorithms. The committee's formulation of the optimization algorithms is based on imposed security constraints ensuring that equipment operates within design limits and that the network does not fail. In current practice, these constraints are set conservatively and there is little testing to learn where the constraints can be relaxed to achieve significantly better performance. The availability of finer resolution data that measure the state of the network on time scales comparable to the ubiquitous 60-Hz ac frequency creates opportunities to do better. Learning how to exploit these opportunities has only begun, in part because they demand a conceptual framework for the control and operation of the grid that incorporates more of the grid's dynamics than is currently the case.

Innovation in the electricity industry requires more than the solution of technical problems. The industry has many commercial stakeholders with vested interests overseen by a complex set of regulatory authorities. The section below, "Synthetic Data for Facilitating the Creation, Development, and Validation of New Power System

Tools for Planning and Operations," addresses the balance between security and financial incentives to keep data confidential on the one hand and open on the other to satisfy researchers' needs for access to data. The path proposed here is to create synthetic data sets that retain the salient characteristics of confidential data without revealing sensitive information. Because developing ways to do this is in itself a research challenge, the committee gives one example of recent work to produce synthetic networks with statistical properties that match those of the electric grid.

Ideally, one would like to have real-time, high-fidelity simulations for the entire grid that could be compared to current observations. However, that hardly seems feasible any time soon. Computer and communications resources are too limited, loads and intermittent generators are unpredictable, and accurate models are lacking for many devices that are part of the grid. The section "Data-Driven Models of the Electric Grid" discusses ways to use the extensive data streams that are increasingly available to construct data-driven simulations that extrapolate recent observations into the future without a complete physical model. Not much work of this sort has yet been done: Most attempts to build data-driven models of the grid have assumed that it is a linear system. However, there are exceptions that look for warning signs of voltage collapse by the monitoring of generator reactive power reserves. The potential payoff for work in this direction depends on the results and should be regarded as high risk, high reward at this point. If it does turn out that large parts of the grid can be reduced to low-dimensional models, this information would be very useful for making short-term forecasts of grid behavior.

SYNTHETIC DATA FOR FACILITATING THE CREATION, DEVELOPMENT, AND VALIDATION OF NEW POWER SYSTEM TOOLS FOR PLANNING AND OPERATIONS

Data of the right type and fidelity are the bedrock of any operational assessment or long-range planning for today's electric power system. In operations, *assessment* through simulation and *avoidance* of potentially catastrophic events by positioning a system's steady-state operating point based on that assessment is the mantra that has always led to reliability-constrained economical operation. In the planning regime, simulation again is key to determining the amount and placement of new generation, transmission, and distribution. The data used to achieve the power industry's remarkable record of universal availability of electricity has been relatively simple compared to future data needs, which will be characterized by a marked increase in uncertainty, the need to represent new disruptive technologies such as wind, storage, and demand-side management, and an unprecedented diversity in policy directions and decisions marked by a tension between the rights of states and power companies versus federal authority. The future grid is likely to be characterized by a philosophy of command and control rather than assessment and avoidance, which will mean an even greater dependence on getting the data right.

The U.S. electric power system is a critical infrastructure, a term used by the U.S. government to describe assets critical to the functioning of our society and economy. The Patriot Act of 2001 defined critical infrastructure as "systems and assets, whether physical or virtual, so vital to the United States that the incapacity or destruction of such systems and assets would have a debilitating impact on security, national economic security, national public health or safety, or any combination of those matters." Although the electric grid is perhaps the most critical of all the critical infrastructures, much of the data needed by researchers to test and validate new tools, techniques, and hypotheses is not readily available to them because of concerns about revealing too much data about critical infrastructures.

This lack of easily available high-quality, realistic power grid data has become a thorny issue in a number of research communities that seek to address mathematical problems arising in power engineering, advance the state of the art in computation and simulation, and test new market designs that may result in better economic efficiencies. Some researchers with good access to power utilities enjoy the ability to test ideas on real data, but even here they are usually unable to publish full results from such tests in the open literature. Many other researchers do not even have such access, and are instead forced to rely on public repositories. Examples of such repositories include the family of steady-state power flow case files that ship with the free (and very useful) Cornell University MATPOWER suite (Zimmerman et al., 2011); the free power flow and transient stability cases used with some textbooks (PowerWorld Corporation, "Glover Overbye Sarma," http://www.powerworld.com/gloveroverbyesarma. accessed March 23, 2016); power flow and transient stability cases from the University of Illinois (ICSEG, 2015); and cases available at the University of Washington Power Systems Test Case Archive or at the University of

Edinburgh's School of Mathematics' Power Systems Test Case Archive (NetworkData). Market data are even harder, if not impossible, to obtain.

The alternative to using public data is to enter into a nondisclosure agreement (NDA) with a consenting utility to use its data for testing ideas. The use of NDAs to obtain data, while useful in many instances, limits the important ability of the broader scientific community[1] to check the reproducibility of results and to use important findings to advance their own work. This reproducibility principle has been a standard for scientific advancement since the time of Aristotle. In fact, the International Council for Science has made the following recommendations (among others):

> Science publishers and chief editors of scientific publications should require authors to provide explicit references to the datasets underlying published papers, using unique persistent identifiers. They also should require clear assurances that these datasets are deposited and available in trusted and sustainable digital repositories. Citing datasets in reference lists using an accepted standard format should be considered the norm (ICSU, 2014, p. 2).

However, while openness and transparency are desirable, the mechanisms for achieving open access will necessarily vary by discipline. There are legitimate constraints on open access to some research data and, in some cases, to the research findings themselves. Whatever the case, transparency should be the norm, to be deviated from only with good reason. In those latter cases the use of synthetic data—that is, data that are not derived from real measurements but rather are synthesized from real measurements—could be an option provided the data sets are sufficiently rich enough to support new findings that would have been discovered by using real data while masking certain information such as private information (health, census, and the like), sensitive economic information, or specific parameters and topologies associated with a critical infrastructure. Overcoming the challenges in generating and using synthetic data is one way to resolve some of the restrictions associated with critical infrastructure data.

Other fields have had to deal with this data problem.[2] The U.S. Bureau of the Census, for example, has set up a number of secure national data centers (the Federal Statistical Research Data Centers, or FSRDCs) at various universities, nonprofit organizations, and government entities across the country. Researchers wishing to use highly granular Census data can apply for permission to access the data on-site at one of the FSRDCs. This provides scientific access to highly valuable data while respecting confidentiality concerns. Fienberg (2006) summarized the technical goals of disclosure limitation techniques as follows: (1) inferences should be the same as if the researcher had complete original data; (2) researchers should have the ability to reverse disclosure protection mechanisms, not for individual identification but for inferences about parameters in statistical models; (3) there should be sufficient variables to allow for proper multivariate analyses; and (4) researchers should not only have the ability to assess the fit of the models but also be provided with most summary information, such as residuals (to identify outliers). The core guiding principle should be to generate released data that are as close as possible to illuminating the research frontier. These principles hold just as much for microdata as for synthetic data. From a scientific perspective, then, any field that may advance from analysis of, and experiments using, high-quality data is significantly limited in the absence of such data.

The electric industry perspective is that actual electric grid data are too sensitive to freely disseminate, a claim that is clearly understandable and justifiable. Network data are especially sensitive when they reveal not only the topology (specific electrical connections and their locations) but also the electrical apparatuses present in the network along with their associated parameters. Revealing these data to knowledgeable persons reveals information an operator would need to know to ensure a network is reliable as well as the vulnerabilities an intruder would like to know in order to disrupt the network for nefarious purposes.

There is also some justifiable skepticism that synthesized data might hide important relations that a direct use of the confidential data would reveal. This makes the development of a feedback loop from the synthetic data to the confidential data essential to develop confidence in the products resulting from synthetic data and to ensure their continuous improvement. A natural question is therefore what, if anything, can be done to alter realistic data

[1] Including the disciplines of engineering and science engaged in invention and discovery of new methods and techniques of design.

[2] See, for example, the work by Agrawal and Srikant (1994), which exemplifies the work on privacy for preserving data mining.

so as to obtain synthetic data that, while realistic, do not reveal sensitive details. An interesting methodology used to mask grid data is described in Borden et al. (2012). ARPA-E recognized this problem in a 2015 Funding Opportunity Announcement (DE-FOA-0001357), Generating Realistic Information for Development of Distribution and Transmission Algorithms (GRID DATA).

Industry's hesitation to reveal too much data might also indicate a view of what problems need to be solved that differs from the committee's view. For example, one may consider the ACOPF problem as a multilayered engineering problem, where producing a workable solution is the challenge, and that requires data of a certain fidelity. From that perspective, one could conclude that a mathematical representation, absent simplifying assumptions, cannot readily be solved. But the committee believes that different mathematical approximations to the general ACOPF problem can provide value, such as for carrying out what-if analyses that help guide engineering decisions. With that view, data sets of less-than-perfect fidelity can provide important value as long as they reflect characteristics of real-world data.

It is clear that the availability of realistic data is pressing, critical, and central to enabling the power engineering community to rely on increasingly verifiable scientific assessments. In an age of Big Data such assessments may become ever more pressing, perhaps even mandatory, for effective decision making.

Recommendation 4: Given the critical infrastructure nature of the electric grid and the critical need for developing advanced mathematical and computational tools and techniques that rely on realistic data for testing and validating those tools and techniques, the power research community, with government and industry support, should vigorously address ways to create, validate, and adopt synthetic data and make them freely available to the broader research community.

Random Topology Networks

This subsection describes a specific problem as an example of the type of mathematical research required to generate synthetic data. There has been extensive research recently on characterizing statistical properties of networks in varied domains. The idea here is to apply such analysis to the electric grid and then to generate fictional networks that share all the properties of the real ones. If the analysis is valid, ensembles of real and synthetic networks could be generated where real and synthetic networks are indistinguishable.

It has been shown that electric power grids have distinct topological characteristics (Hines et al., 2010). Wang et al. (2010a) systematically investigated both the topological and electrical characteristics of power grid networks based on available, real-world power grid system data. First, power grids have salient "small-world" properties, since they feature much shorter average path length (in hops) and much higher clustering coefficients than that of Erdos-Renyi random graphs with the same network size and sparsity. Second, their average node degree does not scale as the network size increases, which indicates that power grids are more complex than small-world graphs; in particular, the node degree distribution is well fitted by a mixture distribution coming from the sum of a truncated geometric random variable and an irregular discrete random variable. In Wang et al. (2010b) the deviation of the node degree distribution of power grids from a pure geometric distribution is investigated, with the result that the deviation substantially affects the topological vulnerability of a network under intentional attacks when nodes with large degrees become first targets of an attack. Another important characteristic of a power grid network is its distribution of line impedances, whose magnitude exhibits a heavy-tailed distribution that is well fitted by a clipped double-Pareto-lognormal distribution (Wang et al., 2010a).

Using recent advances in network analysis and graph theory, many researchers have applied centrality measures to complex networks in order to study network properties and to identify the most important elements of a network (see, for example, Wang et al., 2012). Real-world power grids experience changes continuously. The most dramatic evolution of the electric grid in the coming 10 to 20 years will possibly be seen from both the generation side and the smart grid demand side. Evolving random topology grid models would be significantly enhanced and improved and made even more useful if, among other things, realistic generation and load settings with dynamic evolution features, which can truly reflect the generation and ongoing load changes, could be added.

DATA-DRIVEN MODELS OF THE ELECTRIC GRID

Operating the electric grid is a collective enterprise of a large industry with well over a century of experience. Models, data, and controls all play their part, but their interactions are hardly seamless, because the mathematics on which they draw reflects theoretical communities that have each gone their own way, in some cases duplicating work as a result of different and even conflicting terminology. This section discusses the opportunities in a coherent amalgamation of these multiple mathematical perspectives. To set the stage, consider three problems from other areas that bear upon the interactions of models and data: (1) planning spacecraft trajectories, (2) discovering the causes of cancer, and (3) predicting weather. The first of these is (almost) solved entirely with *the* classical dynamical model, namely, Newton's laws. There may be small stochastic influences from the solar wind and so forth and bits of control theory used to compute midcourse corrections, but these are minor compared to the dynamics. For cancer and other diseases, statistical analysis of data is primarily used to sort out what appears to be the relative contributions of different causative factors. It is therefore surprising when dynamic models contribute to developing treatment protocols, as in the synergistic effects of multiple drugs in the HAART treatment of HIV infections. Weather prediction lies in between, and data assimilation methods that intertwine data and models have led to marked improvements in forecasts. Note here that theory and simulation have been used to argue that forecasts beyond short and medium forecasts will never be feasible owing to the chaotic nature of atmospheric dynamics. Moreover, ensemble forecasting methods are used to estimate the uncertainty of forecasts by analyzing the effects of small changes in initial data or parameters of simulations.

Where does the power grid fit on this graph one of whose axes is the accuracy of models and the other is the statistical analysis of (large) data sets? Grid models are complex and make approximations that are not always completely understood. Loads are the biggest modeling uncertainty: They are constantly changing on multiple time scales and are hardly controlled in today's grid. Still, dynamical models are essential to planning and operation of the grid, as evidenced by the transient stability load model mentioned in Chapter 3 that has 100 or so parameters. Improvements in model fidelity and confidence in their predictive capabilities will lead to more efficient and reliable power systems. Although the mathematical challenges are enormous, there are good reasons to be optimistic that substantial progress can be made with integrative methods that draw on dynamical systems theory, control theory, and machine learning. This subsection sketches some of the things that might be done.

Normal operation of the grid maintains system frequency and voltage through an elaborate set of controls designed to maintain stability. Even though system models have a vast number of degrees of freedom as a dynamical system, most of these are expected to be highly damped so that highly reduced and coarse-grained models can faithfully capture many of the dynamics related to more stressed operation, some large-scale oscillations, and other instabilities. Of course, sometimes the full model details are needed to represent quite stressed situations, such as motor stalling and recovery behavior during and after faults. Excepting these more unusual situations, the system can be engineered so that the more reduced models are adequate and can seek to do so even more in the future. Important questions are (1) How much of a reduction is possible? and (2) How can a reduced model be produced effectively? Analogous questions have been posed in the context of attractors for autonomous dynamical systems. There, nonlinear time-series analysis (Kantz and Schreiber, 2004) methods have been developed that use trajectory data to measure characteristics such as dimension, entropy, and Lyapunov exponents of the attractors. Heuristically, one thinks of trajectories in a high-dimensional phase space evolving onto attractors that may be fractal sets of much lower dimension.[3] Tubes of nearby trajectories evolve to become squashed in some directions and elongated in others. Finite-time Lyapunov exponents quantify this geometry, which results in a cloud of initial conditions flowing to become almost a multidimensional ribbon. In the language of machine learning, the flow map F advances trajectories by a time t as an approximation of a low-rank map. Singular value decomposition of the Jacobian of F gives a linearized view of the squashing that occurs near this trajectory. One of the objectives for nonlinear time-series analysis is to reconstruct this structure from trajectories.

For observations of an ongoing system, the only way to obtain information about the phase space geometry near a given location on an attractor is to pick out trajectory segments in the region of interest. The time between

[3] There are several different definitions of dimension, but the differences are unimportant here.

recurrences to typical regions increases at an exponential rate as the dimension of the system attractor increases. Consequently, attractor reconstruction for a system is feasible only when the dimension of the attractor is small. The recurrence time will be astronomical for attractors whose dimension is larger than, say, five or six. Because the power grid has changing inputs and is so high dimensional with loosely coupled parts, it seems to be an unlikely candidate for this type of analysis. On the other hand, it is controlled to operate at steady state, and stable states typically bifurcate when a single mode becomes unstable. This observation has been used to develop methods that anticipate voltage collapse and take control actions to prevent the collapse from occurring (Dobson and Chiang, 1989). Further development of these methods that take account of stochasticity in the system is a hot research area at this time, especially in the context of ecological regime shifts and tipping points for climate change. Additionally, "sandpile" models rooted in statistical physics are a useful tool in studying cascading failures in networks. In both of these instances, theory points to models of sufficient generality that they serve as a guide for what to look for in data, even in the absence of quantitative circuit models of the grid.

Apparently, little research has been done to monitor the transmission grid and analyze its dynamics from the perspective of nonlinear time-series analysis. This is hardly surprising because the grid is always changing: Its asymptotic behavior as an autonomous dynamical system on time scales of days and longer is of less interest than its predictability on time scales of minutes and hours. Indeed, the unit commitment auctions of the system operators "reset" the grid as a dynamical system daily. Moreover, the data for fine-scale monitoring of nonlinear dynamics of the grid is only now becoming available with the widespread use of synchrophasors. Monitoring methods have primarily relied on a linear perspective in which transients typical of linear systems are fit to data. Further work to identify modes that are close to marginal stability could provide the basis for establishing monitors that trigger alarms signaling the need for intervention to prevent the instability and provide quantitative estimates of its immediacy. Going farther, the methods might also lead to faster control procedures that are based on simulations of reduced models.

Consider a control problem in which parameters of a dynamical system must be optimized to achieve maximal performance, but the system is so large that detailed simulation is too slow for most purposes. Even in these circumstances, the system may behave as if it has much lower dimension, either as an "emergent" property of the network or because it was engineered to be that way. Then, a machine learning algorithm may be able to learn the relationship between the control variable settings and the performance level for the dynamical system. In particular, using data collected about the inputs and corresponding performance levels for various control parameter settings, a machine learning method should be able to create a function that predicts what the performance level is for a new control parameter setting. The expert system does not need to know anything about the dynamical system in order to make these predictions, it simply uses past data from the simulations. This way, the machine learning model is interpolating between known solutions from the dynamical system simulation. It may also be feasible to experimentally achieve the same end by stimulating the system itself. By injecting small disturbances into the grid from a known stable state, the transient response of the system gives information about the decay rates of the slowest modes in the system. For optimal control, the modes that are most vulnerable can be targeted. When known solutions are unavailable for interpolation, one can resort to additional large-scale simulation to plan responses when similar conditions occur in the future. The frequency with which observed states recur is closely related to attractor dimensions, so creating a low-dimensional grid will facilitate successful monitoring and future analysis.

Nonlinear time-series analysis methods can be applied to data without using a dynamical model; however, the goal is clearly to fit models to the data. Experience with other systems indicates that initial attempts to do so will produce disappointing results. Even when trajectories are obtained from simulation of a deterministic vector field, identifying the parameters that gave rise to the trajectory is a problem that does not yet have a good solution. One approach is to use optimization algorithms to minimize a residual between the trajectory and data as initial conditions and parameters are varied. Sensitivity to initial conditions within chaotic attractors is one reason why the problem is hard: Nearby trajectories separate at exponential rates, so identifying a trajectory by its initial condition requires exponentially increasing accuracy as the length of the trajectory grows. One consequence is that the residual function has very large gradients and higher derivatives, so its minima are not easily found with methods that are based on fitting quadratic functions (Phipps et al., 2006). This is an additional reason to set a time horizon on quantitative efforts at grid forecasting.

The analogy between grid simulation and weather prediction appears appropriate. The atmosphere is turbulent, and atmospheric models have chaotic attractors. Data assimilation is used with operational weather forecast models to repeatedly update initial conditions; these updates will be used in subsequent forecasts, because model trajectories will increasingly depart from observations as time goes on. The dimension of the atmospheric attractor appears to be large, so highly reduced models do not appear to work well for forecasting.

Are these also characteristics of the transmission grid? It is not fully known, but the future grid can be engineered so that it is more predictable and amenable to this kind of analysis. Machine learning techniques, together with dynamical systems analysis and control theory, provide foundations for pursuing high-fidelity, real-time simulations of the grid that can be usefully fit to large volumes of phasor measurement unit (PMU) data. In particular, machine learning can be applied to develop models for aspects of the systems like time-varying loads for which physical models are lacking. An additional research challenge with this approach is to correctly incorporate the many discrete events that can occur during a large-scale disturbance, such as those due to protection systems or to motors stalling.

Recommendation 5: Integration of theory and computational methods from machine learning, dynamical systems, and control theory should be a high-priority research area. The Department of Energy should support such research, encouraging the use of real and synthetic phasor measurement unit data to facilitate applications to the power grid. Establishment of experimental test-beds would be a valuable additional resource.

THE ROLE OF CONTROL THEORY IN THE CHANGING ELECTRIC ENERGY SYSTEMS

Automated controllers are critical to reliable power system operation over a wide range of operating conditions (equipment status and input/output ranges). The control architecture is one of the most uncertain aspects of evolution of the power system. At one extreme, system operators will adapt to increasingtly distributed, intermittent energy producers without relinquishing centralized control. At the other extreme, microgrids[4] will be created that can both operate independently and be attached to a regional grid. In normal operation, the microgrids would be responsible for their own control. For emergencies and during periods when low-cost power is available, there would be protocols for how power is delivered from the regional grid to the microgrid. Intermediate scenarios that rely on distributed control architectures are also possible. Whichever control architectures are adopted, they are expected to rely more on automation and to make more use of sensors and advanced communication technologies.

As conditions vary, set points of controllable equipment are adjusted by combining an operator's insights about the grid response and the results of optimization given an assumed forecast. If done right, system operators do not have to interfere with the automation: Their main task is to schedule set points given the forecasts. Fast dynamic transitions between new equilibria are stabilized and regulated by the primary controllers. Beyond this primary control of individual machines, there are two qualitatively different approaches to ensuring stable and acceptable dynamics in the changing power industry:

- The first approach meets this goal of ensuring stable and acceptable dynamics via coordinated action of the system operators. Planners will attempt to embed sensing, communications, and controllers sufficient to guarantee system stability for the range of operating conditions of interest. This is an ambitious goal that faces theoretical challenges. For example, maintaining controllability and observability with increased numbers of sensors and controllers is a challenge given the current state of primary control. It seems feasible that current technologies will allow meeting performance objectives, which are now constrained by requirements for synchronization and voltage stabilization/regulation. As mechanically switched transmission and distribution equipment (phase angle regulators, online tap changers, and so

[4] According to the Department of Energy (DOE), "A microgrid is a local energy grid with control capability, which means it can disconnect from the traditional grid and operate autonomously" (see DOE, "How Microgrids Work," June 17, 2014, http://www.energy.gov/articles/how-microgrids-work).

forth) is replaced by electronic devices—flexible ac transmission systems, high-voltage dc transmission lines, and the like—the complexity of the control infrastructure for provable performance in a top-down manner is likely to become overwhelming. In particular, variable-speed drives for efficient utilization of power are likely to interfere with the natural grid response and the existing control of generators, transmission, and distribution equipment.

- The second approach is the design of distributed intelligent Balancing Authorities (iBAs) and protocols/standards for their interactions. As discussed in Chapter 1, automatic generation control is a powerful automated control scheme and, at the same time, one of the simplest. Each area is responsible for coordinating its resources so that its level frequency is regulated within acceptable limits and deviations from the scheduled net power exchange with the neighboring control areas are regulated accordingly. A closer look into this scheme reveals that it is intended to regulate frequency in response to relatively slow disturbances, under the assumption that primary control of power plants has done its job in stabilizing the transients.

  It is possible to generalize this notion into something that may be referred to as an iBA, which has full responsibility for stabilization and regulation of its own area. Microgrids, distribution networks, portfolios (aggregates) of consumers, portfolios of renewable resources, and storage are examples of such areas. It is up to the grid users to select or form an iBA so that it meets stability and regulation objectives on behalf of its members. The operator of a microgrid is responsible for the distributed energy resources belonging to an area: The microgrid must have sufficient sensing, communications, and control so that it meets the performance standard. This is much more doable in a bottom-up way, and it would resemble the enormously successful Transmission Control Protocol/Internet Protocol (TCP/IP). Many open questions remain about creating a more streamlined approach to ensuring that the emerging grid has acceptable dynamics. For example, there is a need for algorithms to support iBAs by assessing how to change control logic and communications of the existing controllers to integrate new grid members.

The contrast between these two approaches reflects the tension between centralized and distributed control. Because experiments cannot be performed regularly on the entire grid, computer models and simulations are used to test different potential architectures. One goal is to design the system to be very, very reliable to minimize both the number and size of power outages. The problem of cascading failures looms large here. The large blackouts across the northeastern United States in 1965, 2003, and 2007 are historical reminders that this is a real problem. Since protective devices are designed to disconnect buses of the transmission network in the event of large fault currents, an event at one bus affects others, especially those connected directly to the first bus. If this disturbance is large enough, it may trigger additional faults, which in turn can trigger still more. The N – 1 stability mandate has been the main strategy to ensure that this does not happen, but it has not been sufficient as a safeguard against cascading failures. The hierarchy of control for the future grid should include barriers that limit the spread of outages to small regions.

PHYSICS-BASED SIMULATIONS FOR THE GRID

How can mathematics research best contribute to simulation technology for the grid? Data-driven models, described in "Data-Driven Models of the Electric Grid" earlier in this chapter, begin with a functioning network. Moreover, they cannot address questions of how the grid will respond when subjected to conditions that have never been encountered. What will be the effects of installing new equipment? Will the control systems be capable of maintaining stability when steam-driven generators are replaced by intermittent renewable energy resources? Simulation of physics-based models is the primary means for answering such questions, and dynamical systems theory provides a conceptual framework for understanding the time-dependent behavior of these models and the real grid. Simulation is an essential tool for grid planning, and its design requires extensive control, as discussed in the preceding section. In normal steady-state operating conditions, these simulations may fade into the background, replaced by a focus on optimization that incorporates constraints based on the time-dependent analysis. Within power systems engineering, this type of modeling and simulation includes TS analysis.

Effective physics-based modeling of the grid has three components: (1) the formulation of models, (2) the implementation of algorithms and software that produce simulations of the models, and (3) a theoretical and conceptual context for determining which simulations should be performed and for interpreting the results of these simulations. All three components pose mathematical and computational challenges, which the committee discusses in this section.

The modeling challenges begin with reducing models for the different devices that sit at the buses (nodes) of the networks by amalgamating multiple devices. The models can be at different scales, and the derivation of coarser-scaled models from ones at finer scales is an important part of the modeling. As new types of electronic devices are incorporated into the network, new device models are needed. Responsibility for creating, testing, and disseminating these models is diffuse and contributes to uncertainty about the fidelity of grid models. The committee returns to this issue in Chapter 8.

A second modeling issue that requires further attention is the mixture of discrete and continuous components in the grid. One modeling challenge for such hybrid systems is to incorporate the discrete logic into simulation algorithms in a seamless way. Numerical integration algorithms for continuous time systems approximate solutions in discrete steps. When discrete components like circuit breakers trip, it is necessary to detect that this happened within a step, compute more precisely the time at which it occurred, and compute the state of the system after it occurred. The effect of the event can even change the model in significant ways, as when a short circuit changes network topology. When multiple events occur at almost the same time, the model must resolve whether there are interactions between the events, minimally determining which occurred first, and whether this prevented the subsequent events from happening. In systems with the possibility of cascading failures, there is a trade-off between the complexity of the model logic and computational efficiency. In commercial transient stability packages this relatively common situation is often handled by splitting the integration time step. Increasing reliance on electronic control components makes this issue of hybrid simulation even more common.

A third modeling issue that requires more attention is the way stochastic and uncertain phenomena are represented in the models. The theory of stochastic differential equations (SDEs) treats models in which stochastic fluctuations are always present and their main effects come from the accumulation of large numbers of tiny events. The classical problem for these methods is Brownian motion of a particle responding to the impacts of molecular motions in a fluid. The grid manifestly sees unpredictable events that do not fit the context of SDEs. Lightning strikes, weather changes that affect wind and solar generation, and the switching of loads all have sufficient magnitude that it may be more appropriate to treat them as additional discrete events, albeit events whose magnitude and timing are unpredictable. Measuring probability distributions that characterize their uncertainty and incorporating these distributions into models is a significant modeling challenge.

Numerical methods for solving initial value problems for ordinary differential equations (ODEs) provide the algorithmic foundation for grid simulations. The theory and implementation of these methods is a mature subject that, like optimization, comes with many different algorithms that are the algorithms of choice for different subclasses of problems. However, grid simulations involve additional issues that create still further challenges. One immediate issue is that some grid problems like N – 1 stability require prodigious numbers of simulations, so the state of computational performance results in discouraging limits on what can be accomplished relative to our goals. This issue is even more concerning when considering N – 1 – 1 and N – k analysis. Especially in software packages that make it easy to assemble models from standard components, inadequate attention is often given to computational speed.

Models of even "standard" resistance, inductance, and capacitance electric circuits are differential algebraic equations (DAEs) rather than ODEs. In DAEs, some of the equations express fixed relationships among the variables. Ohm's law for a linear resistor (V = IR) is a familiar example of an algebraic equation. Sometimes, but not always, the algebraic equations can be "solved" to eliminate variables and produce a reduced system of ODEs. When reduction fails, the fundamental existence and uniqueness theorem for solutions of initial value problems may also fail. One remedy for this failure is to formulate more elaborate "multiphysics" models, which greatly increase the computational resources needed to solve the systems. Yet another difficulty is that reduction of a system may require differentiation of the algebraic equations.[5] Numerically, special algorithms are required to

[5] Such DAEs with higher indexes appear frequently in mechanics.

ensure that the algebraic equations themselves continue to be satisfied rather than just their derivatives (Hairer and Wanner, 1996). Research is needed to automatically detect when DAEs encounter these issues and take corrective action suited to the physics of the models. Failure of existence and uniqueness of DAEs introduces a new kind of uncertainty into debugging models for the grid because unexpected results can come from either model properties or programming errors.

Figure 1.4 in Chapter 1 depicts the time scales relevant to the grid. Many initial value solvers are limited in the time steps they can use by the fastest time scales present in a system. This prompts the creation of reduced models that retain only the time scales of immediate interest. When there are multiple time scales of interest in a problem, these time scales can interact to produce qualitative phenomena on intermediate time scales not represented explicitly in a model. There is a large literature on specialized "stiff" solvers that can operate on slower time scales of a system when the fast time scales are at quasi-steady states that form "slow manifolds" for the system. However, much less is known about the numerical analysis of systems with interacting time scales, and software libraries may lack algorithms that deal with special problems that arise in these circumstances. Transient stability and longer- term dynamics simulations have largely avoided such issues by ignoring these faster dynamics or by using multirate methods (Crow and Chen, 1994). Electromagnetic transient packages using implicit integration methods with time steps of microseconds can directly handle these situations, albeit with more stringent computational requirements.

Fast switching behaviors of digital devices further complicate accurate simulation of the grid. On the time scales of interest, these devices yield models that are not smooth and fall outside the domain for which most of the analysis of numerical solvers is performed. Coupled with the challenges of incorporating stochastic phenomena, the mathematical foundations for simulation algorithms become problematic. Nonetheless, simulation remains important, and research to improve these foundations for complex systems like the electric grid is needed. It is also important to develop further tools that enable more effective use of grid models and the interpretation of simulation results. The example of transient stability analysis will illustrate some of the issues.

A system cannot be expected to be at a desired operating state immediately following a fault. However, if the postfault state lies in the basin of attraction for this desired operating state, no additional control action is needed to clear the fault. This can be tested for individual postfault states, but the entire basin of attraction needs to be characterized for design purposes and making real-time control decisions. Basin boundaries may well be complicated fractal sets, so locating them is a problem. The exhaustive search of phase spaces that compute trajectories from a fine mesh of initial conditions is the obvious strategy to locate the basins and their boundaries. Such searches are feasible only for phase spaces of very small dimension, because the number of points in a mesh grows exponentially with dimension. Similarly, exploring the effects of varying parameters in a model by simulations using meshes of different parameter values is possible only when varying a small number of parameters. When studying models that could have more than a hundred thousand variables, different strategies are called for.

Progress in dynamical systems theory has repeatedly begun with the distillation of mathematical questions and conjectures that isolate critical issues in their simplest manifestations from the more complicated contexts in which they arose. An important example for the power grid that illustrates this strategy draws on the classification of bifurcations of vector fields. For equilibria of dynamical systems that depend on a single parameter, bifurcation theory identifies saddle-node and Hopf bifurcations as the only types of generic bifurcations. The theory yields normal forms, which are studied to discover the qualitative properties of each type of bifurcation. Further, the theory gives explicit procedures for locating the bifurcations within more complicated models. Identifiable features of the dynamics can be used to infer the occurrence of these bifurcations from time-series data. Computational tools based on this theory locate parameters where voltage collapse occurs in models of electrical grids and identify warning signs of incipient collapse in real-time monitoring of the grid. This example illustrates how advances in dynamical systems theory have resulted from intensive study of simple examples.

Many of the issues described above have been dealt with individually in varied contexts, but the prospect of developing software that treats them in an integrated manner on systems of the scale of the power grid is daunting. It may never be feasible to implement detailed real-time simulations of the entire grid that capture its dynamics accurately. Nonetheless, smaller models, whether they come as reductions of larger models, coarse graining, or selection of subgrids, will give important insights that can be further tested through the investigation of data-driven models described in the section "Data-Driven Models of the Electric Grid," earlier in this chapter.

Recommendation 6: The Department of Energy should support research to extend dynamical systems theory and associated numerical methods to encompass classes of systems that include electric grids. In addition to simulation of realistic grid models, one goal of this research should be to identify problems that exemplify in their simplest forms the mathematical issues encountered in simulating nonlinear, discontinuous, and stochastic time-dependent dynamics of the power system. The problems should be implemented in computer models and archived in a freely available database, accompanied by thorough documentation written for both mathematicians and engineers. Large grid-sized problems that exemplify the difficulty in scaling the methods should be presented as well.

DATA-DRIVEN APPROACHES FOR IMPROVING PLANNING, OPERATIONS, AND MAINTENANCE AND FOR INFORMING OTHER TYPES OF DECISION MAKING

The amount of data about the power system has been growing at a staggering rate, including PMU data, data from other types of sensors in the transmission and distribution systems and from sensors at the generators, and customer behavior information that includes smart meters, data about events that might influence demand (including media and social media data), weather data, and so forth. While there is certainly a growing interest in the use of these data, they are clearly not being used to their full potential, and their further application would be extremely valuable. Several examples are provided below, in addition to those listed in Chapter 4.

Creating Hybrid Data/Human Expert Systems for Operations

When a serious problem occurs on the power grid, operators might be overloaded with alarms, and it is not always clear what the highest priority action items should be. For example, a major disturbance could generate thousands of alarms (Kezunoic and Guan, 2009). Certainly much work has been done over the years in helping operators handle these alarms and more generally maintain situation awareness, with Panteli and Kirschen (2015) providing a good overview of past work and the current challenges. However, still more work needs to be done. The operators need to quickly find the root cause of the alarms. Sometimes "expert systems" are used, whereby experts write down a list of handcrafted rules for the operators to follow. These could be useful, but a data-driven approach could easily provide a better set of guidelines. In particular, data about past alarms can be used to automatically learn what the optimal set of rules should be for the operators to follow in order to find the root cause of the problem. The data can also be combined with expert knowledge to produce a data-driven model that looks as much as possible like the expert model. One way to do this is to use the expert model within a Bayesian prior; if enough data disagree with the expert knowledge, the expert knowledge will be overwhelmed and the model will agree with the data. Tools from machine learning, particularly tools for interpretable or comprehensible machine learning, could naturally be employed for these tasks. For instance, expert systems are often written in a logical structure of IF-THEN rules (IF Alarm 1 activates THEN do action A, ELSE IF Alarms 2 and 3 activate THEN do action B, and so on). There are ways to learn these types of logical structures from data, called "decision tree," "decision list," or "rule list" machine-learning algorithms.

Machine-Learning Models for Hazard Modeling and Reliability

Sensor data from power equipment are not being used to their fullest extent in ensuring reliability. Consider, for instance, a wind turbine or a transformer with several sensors on it, each emitting signals every few seconds or minutes. Most often, any signal going outside its design range will issue an alarm. This considers a single threshold for each sensor separately and is an incredibly limited way to use data. For instance, if there are trends in the sensor data, the trends will be invisible until the signal reaches the threshold. Worse, patterns that are coupled across the signals are completely ignored by considering each signal separately. For instance, several signals increasing sharply in value should potentially trigger an alarm but would not do so under the current system unless one of the signals crosses a threshold—and by then it may be too late to stop the problem. These patterns can come in many forms, and detection of these patterns can help in predicting future problems. This was discussed in Chapter 4.

Visualization Tools for Understanding Data

The power grid generates a tremendous amount of data, a trend that is growing with the rapid deployment of PMUs, smart meters, and other intelligent electronic devices. Trying to understand these potentially high-dimensional data manually is next to impossible. Power system visualization seeks to help people make sense of these data both by applying algorithms to better extract the information contained in the data, and then to display it graphically. One example was the use of clustering, presented in Chapter 4, associated with the analysis and visualization of transient stability results, a technique that could be readily extended to PMU data.

Much work has been done in this field; an overview is given in EPRI (2009). However, there is still much to be done. For example, it would be useful to have visualization tools that can project high-dimensional data to a relevant two- or three-dimensional subspace. Examples of such techniques include graph projection, where nodes of a graph in high dimensions are projected onto a plane in a way that aims to preserve distances between nodes. These techniques could be useful for knowledge discovery in high-dimensional data from equipment such as sensor data, but also for data about customers. One might like to cluster customers by usage patterns, or by factors that may be relevant for offering incentives.

Detecting Who Has No Power

Power companies do not always know who is out of power if there are no sensors at the locations of individual customers. People may assume the power company already knows about the outage and do not report the problem. It may be possible to deduce who is out of power based on the data from customers who are reporting outages, coupled with other data—for instance, from social media (Twitter, Facebook). It may also be possible to send one or more drones to gather information on outages. This may require an algorithm for image segmentation or image clustering.

Machine Learning for Long-Term Planning

Some of the key problems in the future power grid are behavioral: Will customers adopt certain energy-efficient behaviors if they are given rebates? What will be the demand tomorrow given the weather forecast and in light of, say, that tomorrow is Superbowl Sunday? Can we forecast whether a given customer is likely to purchase an electric vehicle? Knowing this would be useful for long-term planning of needs, like charging stations or generation. Similarly, who is likely to purchase photovoltaics or wind turbines? These can be cast as classic machine learning classification, or regression, problems, as described in the section on hybrid human/data expert systems. Machine learning techniques are natural for situations where there is no physical model (no dynamical system, for instance). Machine learning methods make very few assumptions about the data (usually the assumption is that the data are all drawn randomly from an unknown distribution) but are nonetheless able to predict well future data drawn from the same distribution. These methods can handle high-dimensional data and are easy to implement since there are many publicly available software packages.

Unfortunately all of these very clear ways that power companies could immediately see value from their data are not generally being realized: The data are becoming "data tombs." There could be many reasons for this. Perhaps it is difficult to compute the value of predictive modeling tools, as opposed to optimization tools. If a better solution to the day-ahead unit commitment problem is found, the value of the optimization improvement would be apparent. On the other hand, the value of data showing improvements in detecting who is out of power, determining which pieces of equipment are more likely to fail, or predicting who will purchase an electrical vehicle is not as easy to measure. (However, determining the value of such methods is not difficult.) As another possible reason, power companies do not generally hire data scientists—that is, people with the expertise to process these data and harvest them for useful information. If it is unclear what the value would be of mining these data: A company may not be able to envision what it could be useful for.

Recommendation 7: The Department of Energy should support research on data-driven approaches applied to power systems, including operations, planning, and maintenance. This would include better machine

learning models for reliability; comprehensible classification and regression; low- dimensional projections; visualization tools; clustering; and data assimilation. A partial goal of this research would be to quantify the value of the associated data.

OPTIMIZATION

Convex Relaxation in Grid-Related Optimization

As noted in the subsection Grid-Related Continuous Optimization in Chapter 4, a wish to avoid or finesse nonconvexity has led recently to enormous interest in the idea of solving the ACOPF problem through semidefinite (convex) relaxation; this approach is explored in Lavaei and Low (2012) and Low (2014a,b). Semidefinite programming (SDP) may be viewed as a generalization of linear programming: In SDP, a symmetric positive semidefinite matrix is sought that minimizes an affine function. SDP is of great interest for two related reasons: An SDP problem can be solved in polynomial time, and extensive work beginning in the 1990s showed that some NP-hard problems can be solved approximately via a sequence of SDP problems. For a survey of SDP and other convex optimization problems, see Anjos and Lasserre (2012). A frequently used formulation of the ACOPF is a quadratically constrained quadratic program (QCQP), generically expressed as

$$\begin{aligned} &\min\ v^T F v + c^v v \\ &\text{such that } v^T M_k v + d_k^T v + g_k \leq 0 \qquad\qquad k = 1, \ldots, K \end{aligned} \tag{1}$$

where $v \in \mathbf{R}^n$. In this formulation, v is the vector of variables, F is an $n \times n$ symmetric matrix, c is a vector, and, for $1 \leq k \leq K$, M_k, d_k, and g_k are, respectively, an $n \times n$ symmetric matrix, a vector, and a scalar. Provably polynomial-time algorithms exist for this problem only in the all-convex case, when F and all M_k are positive-semidefinite.

Problem (1) is difficult in large part because of its quadratic nature, that is, the appearance of terms of the form $v_i v_j$. One way to simplify (1) is to define a matrix W as $W_{ij} = v_i v_j$, which leads to a new optimization problem that is linear in the matrix W and the vector v—namely, a computationally tractable linear program, albeit with a larger number of variables. For example, the objective function of (1) becomes

$$\sum_{i,j} F_{ij} W_{ij} + c^T v.$$

However, the new problem with variables v and W is not equivalent to (1) unless the relation $W_{ij} = v_i v_j$ is enforced for every i and j or, in matrix notation, if W is the rank one matrix $W = vv^T$. Unfortunately and not surprisingly, this constraint makes the problem nonconvex. So this constraint is relaxed into the less restrictive condition that W is symmetric and positive semidefinite. This condition is a relaxation of $W = vv^T$ because it is weaker: Every W of the form $W = vv^T$ is rank one and positive semidefinite, but every symmetric positive-definite matrix is not rank one.

If $\hat{W}$ is a solution to the relaxed problem, in general this is not a solution of the original problem (1), unless $\hat{W}$ has rank one. In this case, $\hat{W}$ is an optimal solution of (1), and the objective value corresponding to $\hat{W}$ is the exact optimal value for (1). Even when $\hat{W}$ is not rank one, the semidefinite relaxation always provides a lower bound on the value of the objective in (1).

In Lavaei and Low (2012) and Low (2014a,b), justification is given for expecting that the SDP relaxations of many practical ACOPF problem instances will have rank one optimal solutions. As of 2015, community opinion is not unanimous; see Bienstock (2013) for additional discussion.

The SDP relaxation just given is not the only convex relaxation for ACOPF and may not necessarily be the tightest (or, especially, the fastest: large SDPs can be challenging). Numerous authors have proposed a variety of other relaxations, such as a QC (quadratic convex) relaxation (Coffrin et al., 2015).

A related set of ideas that may also be useful for grid-based optimization problems is polynomial optimization, which is just what its name suggests: minimizing the value of a polynomial subject to inequality constraints involving other polynomials. Like QCQP, this problem is known to be NP-hard, but, in the spirit of SDP-based approximation algorithms, it can be solved by developing a hierarchy of convex (semidefinite) relaxations, in which nonnegative polynomials are represented as sums of squares of polynomials. This representation allows efficient tests based on semidefinite programming of whether there is a unique global optimum.

Because of the substantial potential gains from convex relaxation and polynomial optimization formulations, these topics have recently become hot. Numerical evidence thus far suggests that relaxations are exact in many cases, and that network topology plays a critical role in determining whether a relaxation is exact. Further research is needed before these issues are well understood.

A key usage of ACOPF is in market operation. One of the challenges in the current optimal power flow solution techniques is that their dc-only solutions, where voltage constraints are linearized, do not always accurately represent the future voltage-constrained systems. As the development work on a full-blown ACOPF proceeds, it may be determined that a combination of ACOPF and a dc approximation will be required to converge to an optimal solution. Not only would this provide a better representation of the physics of the system, but it should be included in considerations of market design, particularly related to the consideration of uncertainty. Electricity markets can be quite complex, including an interplay between physics, economics, and uncertainties of many forms. Foundational research is needed to develop enhanced market structures to “future-proof” the grid to these uncertainties.

A final challenge in accelerating progress in grid-related optimization is the very serious lack of realistic test problems. Because service operators compete with one another, they are not willing to make data available that would allow more complete testing of alternative methods. This concern, in the form of proposing the development of synthetic data, appears in the preceding section, “Data-Driven Approaches for Improving Planning, Operations, and Maintenance and for Informing Other Types of Decision Making.” Without realistic and representative problems for testing, it will be impossible to develop reliable optimization methods and software.

Recommendation 8: Orders-of-magnitude improvement in nonlinear, nonconvex optimization algorithms are needed to enable their use in wholesale electricity market analysis and design for solving the ac optimal power flow (ACOPF) problem. Such algorithms are essential to determine voltage magnitudes. Therefore the Department of Energy should provide enhanced support for fundamental research on nonlinear, nonconvex optimization algorithms.

Robust and Chance-Constrained Optimization

A valid criticism of the approach implicit in stochastic optimization (see the subsection on stochastic optimization in Chapter 4) focuses on the generation of the set of scenarios S to be studied. In order for the formulation to prove effective, a very large number of scenarios might be needed; furthermore, their respective probabilities might have to be accurately estimated. In following this methodology, however, there is the chance that the computed solution might prove *nonrobust* in the sense that a realization of the second-stage data not captured by the scenarios could entail high cost or, worse, infeasibility, even if the realization is “close” to a scenario. As an alternative, one could model the second-stage loads d_k (the subscript indicates the bus, and the variables d_k are uncertain quantities).

(1r) For each bus k, there is an estimate $\overline{d}_k$ for the “average” second-stage load at k and an estimate δ_k for the maximum error in estimation. Thus, it is assumed $|d_k - \overline{d}_k| \le \delta_k$ for each k.

(2r) A maximum total error is assumed, e.g., a constraint of the form $\sum_k \left(d_k - \overline{d}_k\right)^2 \le \Gamma$, for some given parameter $\Gamma > 0$.

Let D be the (infinite) set of second-stage load vectors indicated by (1r), (2r). The analogue of the formulation in the subsection on stochastic optimization of Chapter 4 now becomes

$$\min \sum_{g \in G^L} \left[K_g^{(1)} y_g + f_g^{(1)}\left(p_g\right)\right] + \rho\left(p\right) \tag{2}$$

so that p is feasible during the first stage; $L_g y_g \leq p_g \leq U_g y_g$ for all g; and all variables y_g take value 0 or 1. Here, given a vector $p = \{p_g : g \in G\}$, $\rho(p)$ is defined as the max-min second-stage cost assuming that each generator g is used to output p_g in the first stage. Formally, for a vector $\hat{d}$ of second-stage loads, let

$$R\left(p,\hat{d}\right) \doteq \min \sum_{g \in G^F} \left[K_g^{(2)} \tilde{y}_g + f_g^{(2)}\left(\tilde{p}_g\right)\right] \tag{3}$$

so that $(p,\ \tilde{p})$ is feasible in the second stage under loads $\hat{d}$; $L_g\,\tilde{y}_g \leq \tilde{p} \leq U_g\,\tilde{y}_g$ for all $g \in 2^{GF}$; and all variables $\tilde{y}_g$ take value 0 or 1, the minimum cost that can be attained in the second stage under first-stage outputs p and second-stage loads $\hat{d}$.

Then

$$\rho(p) = \max_{\hat{d} \in D} R\left(p,\hat{d}\right). \tag{4}$$

In other words, if outputs p are chosen during the first stage, $\rho(p)$ is the worst-case cost expected in the second stage under all possible loads of the form (1r), (2r), assuming that at the start of the second stage one observes the realization of loads and reacts optimally. Note that the function $\rho(p)$ is implicitly defined through (3) and (4). These definitions further highlight that, as was the case with the formulation equation (3) of Chapter 4, (2) is a bilevel optimization problem. Solution techniques for this problem would likely entail, again, some form of decomposition and an indirect use of duality (if the functions $f^{(k)}$ are convex), such as Benders' decomposition, so as to approximate the lower envelope of the function $\rho(p)$. Sampling from the set D could prove useful in this regard—vectors are sampled satisfying (1r), (2r) and used to generate valid cutting planes, in advance of a formal application of Benders' decomposition.

The above example highlights various features of robust optimization: The data error model is fairly agnostic, and a fictitious adversary is assumed to be able to draw data from that model so as to exploit any weakness in the choice of decision variables. Hence the problem acquires a min-max nature (in the case above, a third "min" layer is used to model the recourse decisions in stage 2). The conservatism of the model is managed through appropriate choices of the parameters Γ and δ_k. In more accessible applications of robust optimization it is possible to pose the min-max problem as a single "compact" (polynomial-size) convex optimization problem. Robust optimization has become a mature methodology with significant theoretical underpinnings (see Ben-Tal and Nemirovski, 2000, and Bertsimas et al., 2011) and computational successes; its strongest suit is the agnostic nature of the models.

A particular version of robust optimization worth pointing out arises in the context of chance constraints. Returning to equation (2), note that an implicit constraint in the choice of the first-stage outputs p is that they must prove feasible for the second stage regardless of the second-stage loads. If the δ_k quantities are chosen very large, the solution might end up being overly conservative. If, on the other hand, they are chosen too small, the computed solution p might entail an unacceptable risk of infeasibility. An alternative would be to use "reasonably small" δ_k values and, in addition, to impose a constraint of the form

$$\text{Prob (the second-stage problem is infeasible)} < \varepsilon \tag{5}$$

where the probabilities are computed assuming that second-stage demands have a probability distribution with means $\bar{d}_k$ and standard deviations δ_k. Here ε is a small parameter. This model allows demand errors to exceed the limit δ_k; however, constraint (5) states that outright infeasibility is rare. One could impose (5) either in addition to the robust model (1r), (2r) for demands or as a substitute. This is an example of a chance constraint; see Nemirovski and Shapiro (2006) for background. As a further elaboration of chance constraints one can generate a robust variant of (5) so as to lessen sensitivity to parameter choices. This yields so-called ambiguous chance-constrained models; see Erdogan and Iyengar (2006).

To illustrate the application of these models in this setting, consider constraint (5). A reasonable version for this constraint would be one where second-stage demands are independently and normally distributed. However, this assumption introduces two kinds of possible errors: (1) errors in estimating the means and standard deviations and (2) model errors (in particular, the normality assumption could be incorrect). In order to guard against such errors, the normality assumption is maintained but now it is assumed that an adversary picks the means $\bar{d}_k$ and the standard deviations δ_k (in other words, these are no longer fixed values that are part of the problem). To formalize this approach, suppose that the total number of second-stage loads is N and let $\hat{D}$ and $\hat{\Delta}$ denote two sets contained in $\mathbf{R}^N$. Then replace (5) with the stronger requirement

$$\max_{\ddot{d}\in\hat{D},\delta\in\Delta} \text{Prob}_{\hat{d},\delta} \text{ (the second-stage problem is infeasible)} < \varepsilon \tag{6}$$

In this constraint, $\text{Prob}_{\hat{d},\delta}$ is the probability under the multinormal distribution where load k is normally distributed, with mean $\bar{d}_k$ and standard deviation δ_k. Thus, the adversary's power is somewhat constrained, but the planner must still choose a sufficiently robust solution. Clearly, using (6) guards against misestimation of the means and standard deviations of demand. In fact, experimental evidence suggests that (6) also protects against model error, in particular the normality assumption and possibly even the independence assumption (provided that correlations are weak). A recent use of this methodology in the power transmission setting is found in Bienstock et al. (2014).

CHALLENGES IN MODELING THE ELECTRIC GRID'S COUPLING WITH OTHER INFRASTRUCTURES

A reliable electric grid is crucial to modern society in part because it is crucial to so many other critical infrastructures. These include natural gas, water, oil, telecommunications, transportation, emergency services, and banking and finance (Rinaldi et al., 2001). Without a reliable grid many of these other infrastructures would degrade, if not immediately then within hours or days as their backup generators fail or run out of fuel. However, this coupling goes both ways, with the reliable operation of the grid dependent on just about every other infrastructure, with the strength of this interdependency often increasing. For example, PNNL (2015) gives a quite comprehensive coverage of the couplings between the grid and the information and communication technology (ICT) infrastructure. The coupling between the grid and natural gas systems, including requirements for joint expansion planning, is presented in Borraz-Sanchez et al. (2016). The interdependencies between the electric and water infrastructures are shown in Sanders (2014) with a case study for the Texas grid presented in Stillwell et al. (2011). While some of these couplings are quite obvious, others are not, such as interrelationships between the grid and health care systems in considering the vulnerability of the grid to pandemics (NERC, 2010). The rapidly growing coupling between electricity and electric vehicle transportation is presented in Kelly et al. (2015).

From the perspective of this report, focusing on the analytic foundations for the next-generation electric grid, it is perhaps best to present these infrastructure interdependencies utilizing the framework of complex adaptive systems (CASs). "Seen from this perspective, which has important benefits for modeling and analysis, each component of an infrastructure constitutes a small part of the intricate web that forms the overall infrastructures. All components are influenced by past experiences" (Rinaldi et al., 2001). The grid itself is a CAS, with the bulk of the report focused on its specific research needs. Including other infrastructures could increase complexity, but would remain a CAS. The degree to which the coupled infrastructures need to be accounted for in the foundational electric grid research considered here is problem specific and depends on many different factors. Hence consideration of coupled infrastructures is implicit in what is proposed in this report. For example, as shown in Kelly et al. (2015), electric vehicles can interact with the grid on time scales ranging from short-term transient stability (referring back to Figure 1.4), when the impact of battery dynamics is considered to decades, when their impact is factored into long-range generation and transmission system planning. The same could be said for the ICT (PNNL, 2015). In contrast, the coupling with water would not apply on the millisecond time frame but would be important over days if generation dispatches needed to be curtailed and over decades for planning (Sanders, 2014, and Stillwell et al., 2011). Hence for some problems, coupled infrastructure consideration would relate directly to Recommendations 5, 6, and 7. An example of how the natural gas infrastructure couples to the electric infrastructure in optimizations is given in Borraz-Sanchez et al. (2016), directly relating to Recommendation 8. Coupled infrastructure simulations will benefit from the synthetic data libraries called for in Recommendation 9, and from the software called for in Recommendation 10, both of which are in Chapter 8. Finally, the consideration of the coupling of the grid with other infrastructures plays a large role in Recommendation 11, also from Chapter 8.

Understanding the effects of these connections between different infrastructures will require answering even more complex versions of the mathematical and computational questions discussed throughout this report, whose applicability goes far beyond the grid. Chapter 3, "Existing Analytic Methods and Tools," and Chapter 4, "Mathematical Research Areas Important for the Grid," summarize some of the most important mathematical and analytical tools that can be applied to such systems, where progress would carry over to other infrastructures.

For example, models that represent coupling between the grid and gas, water, transportation, or communication will almost certainly include hierarchical structures characterized by a mixture of discrete and continuous variables whose behavior follows nonlinear, nonconvex functions at widely varying time scales. This implies that new approaches for effectively modeling nonlinearities, formulating nonconvex optimization problems, and defining convex subproblems would be immediately relevant when combining different infrastructures.

One of the foundations of this report is that mathematics can be used to describe and solve problems from significantly different application domains. In this spirit, the discussion in Chapter 7, "Case Studies," concerning high-impact, low-frequency events stresses the importance of interdisciplinary modeling, noting that these events have commonalities that could be addressed by research in the mathematical and computational areas discussed in this report.

REFERENCES

Agrawal, R., and R. Srikant. 1994. Fast algorithms for mining association rules in large databases. Pp. 487-499 in *VLDB '94 Proceedings of the 20th International Conference on Very Large Data Bases*. Morgan Kaufmann Publishers, San Francisco, Calif.

Anjos, M.F., and J.B. Lasserre (eds.). 2012. *Handbook on Semidefinite, Conic, and Polynomial Optimization*. International Series in Operations Research and Management Science. Springer, New York, N.Y. http://www.springer.com/us/book/9781461407683.

Ben-Tal, A., and A. Nemirovski. 2000. Robust solutions of linear programming problems contaminated with uncertain data. *Mathematical Programming* 88(3):411-424.

Bertsimas, D., D.B. Brown, and C. Caramanis. 2011. Theory and applications of robust optimization. *SIAM Review* 53(3):464-501.

Bienstock, D. 2013. Progress on solving power flow problems. *OPTIMA Mathematical Optimization Society Newsletter* 93, December. http://www.mathopt.org/Optima-Issues/optima93.pdf.

Bienstock, D., M. Chertkov, and S. Harnett. 2014. Chance-constrained DC-OPF. *SIAM Review* 56(3):461-495.

Borden, A.R., D.K. Molzahn, P. Ramanathan, and B.C. Lesieutre. 2012. Confidentiality-preserving optimal power flow for cloud computing. Pp. 1300-1307 in *2012 50th Annual Allerton Conference on Communication, Control, and Computing*. doi:10.1109/Allerton.2012.6483368.

Borraz-Sanchez, C., R. Bent, S. Backhaus, S. Blumsack, H. Hijazi, and P. van Hentenryck. 2016. Convex optimization for joint expansion planning of natural gas and power systems. *49th Hawaii International Conference on System Sciences (HICSS)*. http://public.lanl.gov/rbent/hicss2016.pdf.

Coffrin, C., H.L. Hijazi, and P. Van Hentenryck. 2015. "The QC Relaxation: Theoretical and Computational Results on Optimal Power Flow." Optimization Research Group, Australian National University. http://arxiv.org/pdf/1502.07847.pdf.

Crow, M.L., and J.G. Chen. 1994. The multirate method for simulation of power system dynamics. *IEEE Transactions on Power Systems* 9(3):1684-1690.

Dobson, I., and H.-D. Chiang. 1989. Towards a theory of voltage collapse in electric power systems. *Systems and Control Letters* 13(3):253-262.

EPRI (Electric Power Research Institute). 2009. *Technology Assessment of Power System Visualization*. Product Abstract 1017795. December. http://www.epri.com.

Erdogan, E., and G. Iyengar. 2006. Ambiguous chance constrained problems and robust optimization. *Mathematical Programming Series B* 107:37-61.

Fienberg, S.E. 2006. Privacy and confidentiality in an e-commerce world: Data mining, data warehousing, matching and disclosure limitation. *Statistical Science* 21(2):143-154.

Hines, P., S. Blumsack, E. Cotilla Sanchez, and C. Barrows. 2010. The topological and electrical structure of power grids. Pp. 1-10 in *2010 43rd Hawaii International Conference on System Sciences (HICSS)*. doi:10.1109/HICSS.2010.398.

ICSEG (Illinois Center for a Smarter Electric Grid). 2015. "Power Flow Cases." http://publish.illinois.edu/smartergrid/power-cases/. Accessed September 15, 2015.

ICSU (International Council for Science). 2014. "Open Access to Scientific Data and Literature and the Assessment of Research by Metrics" (ICSU Report on Open Access). News from the 31st ICSU General Assembly. http://www.icsu.org/general-assembly/news/.

Kantz, H., and T. Schreiber. 2004. *Nonlinear Time Series Analysis*. 2nd ed. Cambridge University Press, U.K.

Kelly, J.C.,T. Ersal, C.-T. Li, B.M. Marshall, K. Soumya, G.A. Keoleian, P. Huei, I.A. Hiskins, and J.L. Stein. 2015. Sustainability, resiliency, and grid stability of the coupled electricity and transportation infrastructures: Case for an integrated analysis. *Journal of Infrastructure Systems* 21(4).

Kezunoic, M., and Y. Guan. 2009. Intelligent alarm processing: From data intensive to information rich. *42nd Hawaii International Conference on System Sciences, 2009. HICSS '09*. doi:10.1109/HICSS.2009.268.

Lavaei, J., and S.H. Low. 2012. Zero duality gap in optimal power flow problem. *IEEE Transactions on Power Systems* 27(1):92-107.

Low, S. 2014a. Convex relaxation of optimal power flow, Part I: Formulations and equivalence. *IEEE Transactions on Control of Network Systems* 1:15-27.

Low, S. 2014b. Convex relaxation of optimal power flow, Part II: Exactness. *IEEE Transactions on Optimal Control of Network Systems* 1:177-189.

Nemirovski, A., and A. Shapiro. 2006. Convex approximations of chance constrained programs. *SIAM Journal on Optimization* 17:969-996.

NERC (North American Electric Reliability Corporation). 2010. *High-Impact, Low-Frequency Event Risk to the North American Bulk Power 2992 System*. June. http://www.nerc.com/pa/CI/Resources/Documents/HILF%20Report.pdf.

Panteli, M., and D.S. Kirschen. 2015. Situation awareness in power systems: Theory, challenges and applications. *Electric Power Systems Research* 122:140-151.

Phipps, E., R. Casey, and J. Guckenheimer. 2006. Periodic orbits of hybrid systems and parameter estimation via AD. *Lecture Notes in Computational Science and Engineering* 50:211-224.

PNNL (Pacific Northwest National Laboratory). 2015. *The Emerging Interdependence of the Electric Power Grid and Information and Communication Technology*. PNNL-24643. August. http://www.pnnl.gov/main/publications/external/technical_reports/PNNL-24643.pdf.

Rinaldi, S.M., J.P. Peerenboom, and T.K. Kelly. 2001. Identifying, understanding and analyzing critical infrastructure interdependencies. *IEEE Control Systems Magazine*, December. pp. 11-25.

Sanders, K.T. 2014. Critical review: Uncharted waters? The future of the electricity-water nexus. *Environmental Science and Technology* 49:51-66.

Stillwell, A.S., C.W. King, M.E. Webber, I.J. Duncan, and A. Hardberger. 2011. The energy-water nexus in Texas. *Ecology and Society* 16(1).

Wang, Z., A. Scaglione, and R.J. Thomas. 2010a. Generating statistically correct random topologies for testing smart grid communication and control networks. *IEEE Transactions on Smart Grid* 1(1):28-39.

Wang, Z., A. Scaglione, and R.J. Thomas. 2010b. The node degree distribution in power grid and its topology robustness under random and selective node removals. *2010 IEEE International Conference on Communications Workshops (ICC)*. doi:10.1109/ICCW.2010.5503926.

Wang, Z., A. Scaglione, and R.J. Thomas. 2012. Electrical centrality measures for power grids. Chapter 12 in *Control and Optimization Theory for Electric Smart Grids* (M. Ilic and A. Chakrab, eds.). Springer,New York, N.Y.

Zimmerman, R.D., C.E. Murillo-Sanchez, and D. Gan. 2011. MATPOWER: Steady-state operations, planning and analysis tools for power systems research and education. *IEEE Transactions on Power Systems* 26(1):12-19.

7

Case Studies

INTRODUCTION

This chapter presents several case studies, each of which connects power grid problems to mathematical and computational challenges. The chapter's overall goal is to illustrate some current mathematical and computational techniques in greater detail than could be captured in earlier chapters. The first section provides an overview of some of the key optimization software used at one of the electricity markets mentioned in Chapter 2 (PJM) and discusses how solving the mathematical challenges would improve its capabilities. That is followed by a case study addressing how to predict and handle high-impact, low-frequency events that could threaten our critical infrastructure. The section "Case Study in Data-Centered Asset Maintenance: Predicting Failures in Underground Power Distribution Networks" discusses the prediction of failures that occur more commonly in which a single piece of equipment fails. This ties into the problem of data-driven asset maintenance, where each asset is a physical component of the grid (e.g., a cable or a transformer) that needs to be maintained before it fails. The section "Case Study in Synchrophasors" discusses synchrophasors, which utilize sensors that can determine both the magnitude and phase angles of power system voltages at rates of 30 to 60 samples per second. The final section presents a case study on real-time, inverter-based control, where potential problems are not only detected, but fast calculations and controls also are utilized to push signals back toward their reference settings.

CASE STUDY IN OPTIMIZATION: PJM'S DAILY OPERATIONS

PJM is a regional transmission organization that coordinates the movement of wholesale electricity in all or parts of Delaware, Illinois, Indiana, Kentucky, Maryland, Michigan, New Jersey, North Carolina, Ohio, Pennsylvania, Tennessee, Virginia, West Virginia, and the District of Columbia. Acting as a neutral and independent party, PJM operates a competitive wholesale electricity market and manages the high-voltage electricity grid to ensure reliability for more than 61 million people. PJM Market Operations coordinates the continuous buying, selling, and delivery of wholesale electricity through the energy market. In its role as market operator, PJM balances the needs of suppliers, wholesale customers, and other market participants, and monitors market activities to ensure open, fair, and equitable access. The operation of PJM's various markets requires the use of many software applications, which vary in purpose and complexity. The next subsection contains a high-level description of applications that are used to support the operation of PJM, which show how important optimization tools are to the power grid in general.

Day-Ahead Market

As covered in Chapter 2, the purpose of the day-ahead market is to make the generator commitment decisions a day ahead of time so the generators have sufficient time to start up or shut down. This market utilizes several different key applications, which are discussed in this subsection.

The Resource Scheduling and Commitment application is a mixed-integer program responsible for committing the bulk—more than 90 percent—of the resource commitments for the PJM system. The following equation presents a simplified version of the unit commitment problem that PJM solves every day to commit resources in the day-ahead market. The objective function of day-ahead unit commitment is to minimize the total production cost of the system while adhering to the enforced transmission limitations. That is,

Minimize

$$Z = \sum_{t=1}^{T}\Bigg[\sum_{i=1}^{I}\Big(C_{i,t}P_{i,t}U_{i,t} + NL_{i,t}U_{i,t} + S_{i,t}\big(U_{i,t},U_{i,t-1}\big)\Big) + \sum_{i=1}^{I}RC_{i,t}ASMW_{i,t} + \sum_{j=1}^{J}C_{j,t}INC_{j,t} - \sum_{k=1}^{K}C_{k,t}DEC_{k,t}$$
$$+\sum_{l=1}^{L}C_{l,t}UTC_{l,t} + \sum_{m=1}^{M}\Big(C_{m,t}ELRP_{m,t}U_{m,t} + SD_{m,t}\big(U_{m,t},U_{m,t-1}\big)\Big) - \sum_{q=1}^{Q}C_{q,t}PD_{q,t}\Bigg]$$

where

$C_{i,t}$ = cost of generating unit i at time t
$P_{i,t}$ = power (MW) generation of unit i at time t
$U_{i,t}$ = commitment status of unit i at time t (1 or 0)
$U_{i,t-1}$ = commitment status of unit i at time $t - 1$ (1 or 0)
$NL_{i,t}$ = no-load cost of unit i at time t
$S_{i,t}$ = start-up cost of unit i at time t
$RC_{i,t}$ = reserve cost of unit i at time t
$ASMW_{i,t}$ = ancillary service (MW) of unit i at time t
$C_{j,t}$ = offer price of increment offer j at time t
$C_{k,t}$ = bid price of decrement bid k at time t
$INC_{j,t}$ = MW for increment offer j at time t
$DEC_{k,t}$ = MW for decrement bid k at time t
$C_{l,t}$ = offer price or up-to-congestion transaction l at time t
$UTC_{l,t}$ = MW of up-to-congestion transaction bid l at period t
$C_{m,t}$ = cost of economic load response resource m at time t
$ELRP_{m,t}$ = MW of economic load response resource m at time t
$U_{m,t}$ = commitment status of economic load response resource m at time t (0 or 1)
$U_{m,t}$ = commitment status of economic load response resource m at time $t - 1$ (0 or 1)
$SD_{m,t}$ = shutdown cost of economic load response resource m at time t
$C_{q,t}$ = bid price of price sensitive demand bid q at time t
$PD_{q,t}$ = MW of price-sensitive demand bid q at time t

Subject to the following constraints:

1. Power balance constraint

$$\sum_{i=1}^{I}P_{i,t}U_{i,t} + \sum_{m=1}^{M}ELRP_{m,t}U_{m,t} + \sum_{j=1}^{J}INC_{j,t} - \sum_{k=1}^{K}DEC_{k,t} - \sum_{q=1}^{Q}PD_{q,t} = \text{Fixed demand}_t + \text{Losses}_t \qquad \text{for } t = 1, \ldots, T$$

2. Ancillary reserve constraint

$$\sum_{i-1}^{I} \mathrm{ASMW}_{i,t} \geq \mathrm{R}_i \qquad \text{for } t = 1, \ldots, T$$

3. Capacity constraints

$$\mathrm{P}_{i,t}^{\min}\mathrm{U}_{i,t} \leq \mathrm{P}_{i,t} \leq \mathrm{P}_{i,t}^{\max}\mathrm{U}_{i,t} \qquad \text{for } i = 1, \ldots, I$$
$$0 \leq \mathrm{INC}_{j,t} \leq \mathrm{INC}_{j,t}^{\max} \qquad \text{for } j = 1, \ldots, J$$
$$0 \leq \mathrm{DEC}_{k,t} \leq \mathrm{DEC}_{k,t}^{\max} \qquad \text{for } k = 1, \ldots, K$$
$$0 \leq \mathrm{UTC}_{l,t} \leq \mathrm{UTC}_{l,t}^{\max} \qquad \text{for } l = 1, \ldots, L$$
$$\mathrm{ELRP}_{m,t}^{\min}\mathrm{U}_{m,t} \leq \mathrm{ELRP}_{m,t} \leq \mathrm{ELRP}_{m,t}^{\max}\mathrm{U}_{m,t} \qquad \text{for } m = 1, \ldots, M$$

where

R_t = reserve requirement at time t
$\mathrm{P}_{i,t}^{\max}$ = maximum output limit of unit i at time t
$\mathrm{P}_{i,t}^{\min}$ = minimum output limit of unit i and time t
$\mathrm{INC}_{j,t}^{\min}$ = maximum MW of increment offer j at time t
$\mathrm{DEC}_{k,t}^{\max}$ = maximum MW of decrement offer k at time t
$\mathrm{UTC}_{l,t}^{\max}$ = maximum MW or up-to-congestion offer l at time t
$\mathrm{ELRP}_{m,t}^{\max}$ = maximum output limit of economic load response m at time t
$\mathrm{ELRP}_{m,t}^{\min}$ = minimum output limit of economic load response m at time t

For simplicity, neither the objective function nor the constraints are shown in the above unit commitment problem formulation, but they are included in the actual day-ahead market clearing software. Some elements that are in the actual formulation but omitted here for simplicity are transmission limitations enforced in the day-ahead market; temporal constraints of units such as start-up times and minimum run times; and the pumped storage hydro-optimization model that PJM currently uses.

A second piece of software used in the day-ahead market is the scheduling, pricing, and dispatch (SPD) application, a linear program that dispatches physical generation and demand resources already committed by resource scheduling and commitment. It can also dispatch virtual bids, including increment offers, decrement bids, and up-to-congestion transactions. Virtual bids are fundamental components of two-settlement markets in every independent system operator (ISO) /regional transmission organization (RTO) in the United States. They are financial instruments bid in by market participants to arbitrage differences between the day-ahead markets and real-time markets. The main benefits of virtual bids are mitigating the unbalance in supply and demand of market power and facilitating the convergence of price and unit commitment.

The third package is known as the simultaneous feasibility test (SFT), which is a contingency analysis program that performs a security analysis of the day-ahead market (details on contingency analysis are covered in Chapters 1 and 3). The SFT screens each dispatch hour for N – 1 overloads. If one is encountered, the SFT application passes information back to the SPD application regarding the N – 1 overload, and the SPD application enforces a specific transmission constraint to mitigate the overload and dispatches resources and calculated prices to appropriately reflect this limitation.

Real-Time Markets

The set of applications described in this section is part of the suite of applications that works simultaneously to control and price the PJM system in real time. The suite of applications includes tools that procure the ancillary services discussed in Chapter 2 and that provide resource commitment and dispatch functionality and, ultimately, the calculation of 5-minute locational marginal costs (LMPs) across the system (LMPs are also discussed in Chapter 2). In the real-time market tools there is no equivalent of the SFT application that exists in the day-ahead market. This is because N – 1 security constraints are identified by the security analysis package in PJM's Energy Management System and are passed right into the dispatch tools listed below. A block diagram of these applications is given in Figure 7.1, with each briefly discussed.

The Ancillary Services Optimizer is software that solves a mixed-integer program to optimize PJM's hour-ahead ancillary services. This application jointly optimizes energy and reserves.

The Intermediate Term Security Constrained Economic Dispatch is a mixed-integer program that provides a time-coupled 2-hour forecast and unit commitment. This application uses forecast data and generator offer parameters to create a dispatch trajectory and unit commitment plans for the next 2 hours. The generator dispatch points calculated by this application are not used for system control. The main purpose of this application is to provide intraday unit commitment information to the system operator.

The Real Time Security Constrained Economic Dispatch (RT SCED) is a 10-min forward linear program that produces the economic dispatch points for all resources on the PJM system. PJM uses this application to dispatch all online generation resources from their current operating point to their most economic operating point based on a 10-min-ahead forecast of system conditions. For example, an RT SCED solution that is executed at 7:45 a.m. uses the current operating state of the system provided by the state estimator as a set of initial conditions. The application then uses load and constraint forecast information for 7:55 a.m., in addition to generator offer information such as ramp rates and the real power minimum/maximum limits, to dispatch the set of online generation resources of PJM in a least-cost fashion to meet system expectations 10 min into the future. This application runs every 5 min or on command by the PJM system operator.

The Locational Pricing Calculator is an application that is identical to the RT SCED application except that the market prices calculated in this application are for the entire network model as opposed to just for generation buses.

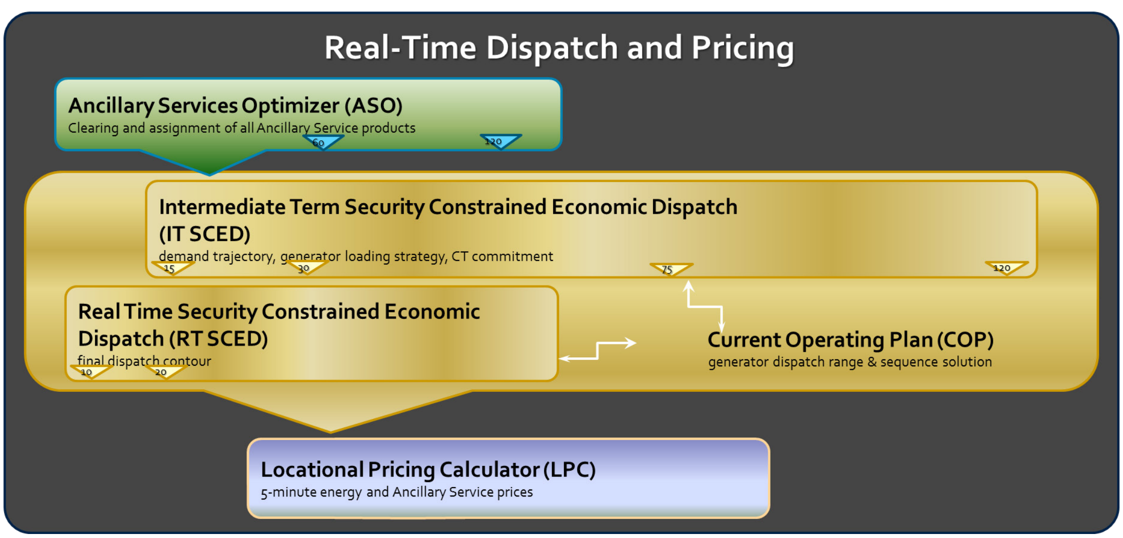

FIGURE 7.1 PJM real-time markets. SOURCE: Courtesy of PJM Interconnection.

Capacity Market—Reliability Pricing Model Optimization

This is the market-clearing engine that clears the PJM capacity markets' base residual and incremental capacity auctions. This application is a mixed-integer program that is used to clear PJM's 3-year forward capacity auction. The main capacity auction, the Base Residual Auction, is run annually 3 years before the actual year for which the capacity is committed. This application uses demand curves to express the willingness to pay for capacity and supply offers to clear the market.

Financial Transmission Rights

Financial transmission rights (FTRs) provide a mechanism by which market participants can hedge against potential losses in the LMP market by providing a stream of revenue when there are price differences in the LMPs between different locations in the system, along what is known as an energy path. FTRs are acquired through auctions. Associated with FTR auctions is the SPD application, which is a linear program that dispatches FTR bids up to cleared quantities. The clearing of an FTR auction is similar to the clearing of point-to-point transactions like up-to-congestion transactions in the day-ahead market. These bid types are described by source and sink locations, as well as a maximum willingness to pay for the price spread between the locations. If the transaction clears, it imposes a flow on the transmission system that is based on the source and sink location and the topology of the system.

Challenges for the Day-Ahead Unit Commitment Formulation

The day-ahead market unit commitment problem is the most complex problem solved by most ISO/RTOs that operate power markets. Building on what was presented in the section "Day-Ahead Markets," the problem could also be formulated using a Lagrangian relaxation where commitment decisions are approximated. The section on Day-Ahead Markets presents a mixed-integer program (MIP) formulation, where binary variables are used to more precisely model discrete decisions. While the MIP provides a more precise solution, it also takes longer to solve than the approximated Lagrangian relaxation solution. The MIP formulation that PJM utilizes to solve the day-ahead market unit commitment problem produces an efficient, reliable unit commitment that is the basis for the next operating day. Like anything else, however, it can be improved with the proper direction and investment.

ISOs and RTOs solve many other optimization problems to schedule and dispatch the system and clear power markets, but all can be derived by simplifying the day-ahead market unit commitment problem. Therefore, typically any challenges encountered in the solution process will be evident somewhere in the day-ahead market. Below is a brief summary of some of the common challenges PJM encounters:

- Significant increases in bid and offer volumes will increase the MIP solution time because of an increased number of binary and continuous variables.
- Large numbers of transmission constraints combined with continuous variables can cause a very dense matrix, which limits the ability to use more efficient sparse matrix solution techniques. Additionally, large numbers of continuous variables increase the time to solve each linear program (LP) in the search tree during the MIP searching process.
- Increasing the MIP gap to improve convergence tolerance and consistency between the LP and MIP solutions degrades performance exponentially. Decreasing the MIP gap to improve performance may result in nonunique MIP solutions.

The above challenges are in some way related to the size and scalability of the general unit commitment problem that exists today. The challenges in solution time presented by these issues typically have been addressed by increasing computer processing capability. If Moore's law continues to hold true, the increases in computer capability may be able to meet the needs of the current unit commitment problem PJM solves. This does not change the need for mathematical work in the short term, however, nor does it change the fact that the problem is likely to become substantially larger as the power grid changes.

In order to make a step change in the size and complexity of the unit commitment problem being solved, there likely needs to be a significant increase in processing capability or a reformulation of the problem. For example, the ability to solve an ac unit commitment problem would be a significant breakthrough for ISO/RTOs in terms of unit commitment accuracy and efficiency. In today's dc models, voltage and reactive constraints are linearized into dc approximations that attempt to model voltage restrictions that are real power flow limitations. This practice has been in place since the inception of power markets in the United States in the late 1990s; however, the practice still results in some unit commitment and market inefficiencies that a better model of ac constraints during the commitment, dispatch, and pricing process could improve.

An example of a simplification that is widely used is the modeling of a reactive limit in a dc model. Currently, reactive limits are an input into the dc problem based on offline studies and a predefined local area unit commitment, as opposed to being optimized as part of the unit commitment problem itself. In reality, the level of the reactive limit will vary based not only on the actual units committed but also on where they are dispatched, because of the relationship between active power and reactive power on generators. Currently, this level of granularity cannot be modeled efficiently enough to solve the problem within the time frame of the day-ahead market; therefore, the outcome of that market may be less efficient than it could have been. The general result is less transparent market prices and out-of-market uplift payments.

Approximated voltage constraints can also be problematic. From a market efficiency perspective, dispatching to a dc approximation of a voltage constraint can create some undesirable outcomes. For example, suppose 100 MW of FTR are sold on an energy path based on the thermal limit of the facility. If that path is then constrained in the day-ahead market or in real time to a flow less than the 100 MW of the FTRs sold because it is being used as a thermal proxy for a voltage constraint, the result will be underfunded FTRs on that path. The level of underfunding will vary depending on the difference between the FTR and day-ahead market and real-time market flows, as well as the shadow price to control the thermal surrogate.

In the dc-only solution in use today, voltage constraints are linearized so that they can be enforced in a linear program. This solution has its shortcomings; however, it is likely that there is a point of diminishing returns with the full ac model such that expansion of the problem beyond a certain point would yield little or no discernable benefit. The most efficient solution might be a blend of the two; the efforts focused on improving the model should consider the benefits and drawbacks of each.

For example, the breakpoint for gaining accuracy by implementing additional ac constraints in the model may stop at a certain voltage level (or in a certain geographic area surrounding a reactive or voltage constraint), such that those constraints would only need to be implemented selectively. This would cut down on the complexity added to the model, while adding the information needed to resolve these types of constraints more efficiently.

CASE STUDY IN MATHEMATICAL NEEDS FOR THE MODELING AND MITIGATION OF HIGH-IMPACT, LOW-FREQUENCY EVENTS

Worldwide, the bulk power system is one of the most critical infrastructures, vital to society in many ways, but it is not immune to severe disruptions that could threaten the health, safety, or economic well-being of the citizens it serves. The electric power industry has well-established planning and operating procedures in place to address "normal" emergency events (such as hurricanes, tornadoes, and ice storms) that occur from time to time and disrupt the supply of electricity. However, the industry has much less experience with planning for, and responding to, what the North American Electric Reliability Corporation (NERC) calls high-impact, low-frequency (HILF) events (NERC, 2010).

The events that fall into this category must meet two criteria. First, they need to be extremely rare or they may never have actually occurred but are plausible. Second, their impact must be potentially catastrophic across a broad portion of the power system. These are events that if they occurred, could bring prolonged blackouts on a large scale, have an adverse economic impact reaching into the trillions of dollars, and kill millions of people. Our modern, just-in-time economy is becoming increasing fragile with respect to disruptions to critical infrastructures in which even short-time, localized blackouts are quite disruptive. Imagine if the power went out for many millions of people and would not be coming back on for weeks or months!

NERC identified several events that fall into the HILF category, including (1) coordinated physical attacks or cyberattacks, (2) pandemics, (3) high-altitude electromagnetic pulses (HEMPs), and (4) large-scale geomagnetic disturbances (GMDs). One such disturbance, a solar corona mass ejection, is shown in Figure 7.2. The identification of these risks was not new with the 2010 report (NERC, 2010), and some work has been done over the years to try to mitigate their impacts. One example is the recently published *Electric Grid Protection (E-Pro) Handbook* (Stockton, 2014). Yet, collectively, HILF events present an interesting case study on the mathematical and computational challenges needed for the next-generation electric grid.

The existing power grid is certainly resilient, often able to operate reliably with a number of devices unexpectedly out of service. While blackouts are not rare, most are small in scale and short term, caused by local weather (e.g., thunderstorms), animals, vegetation, and equipment failures. Regional blackouts, affecting up to several million people for potentially a week or two, occur less frequently. Such events are usually due to ice storms, tornados, hurricanes, earthquakes, severe storms, and, occasionally, equipment failure.

As an example, the derecho that happened in late June 2012 in the U.S. Mid-Atlantic and Midwest was one of the most destructive and deadly, fast-moving, and severe thunderstorm complexes in North American history. It was 200 miles wide, 600 miles long and registered winds as high as 100 mph as it tracked across the region. The morning after the event approximately 4.2 million customers were without electricity across 11 states and the District of Columbia, and restoration took 7 to 10 days (DOE, 2012). A second example that same year was Superstorm Sandy, which caused 8.5 million customer power outages across 24 states, causing damage estimated at $65 billion (Abi-Samra et al., 2014).

While tragic for those affected, aid from unaffected utilities helps to speed the recovery, and electric utility control centers have long experience in dealing with weather-related events. For example, during Superstorm Sandy utilities conducted the largest movement of restoration crews in history, with more than 70,000 utility personnel from across the United States and Canada deploying to support power restoration, and power restoration was an overriding priority for all U.S. federal departments, including the Department of Defense (Stockton, 2014).

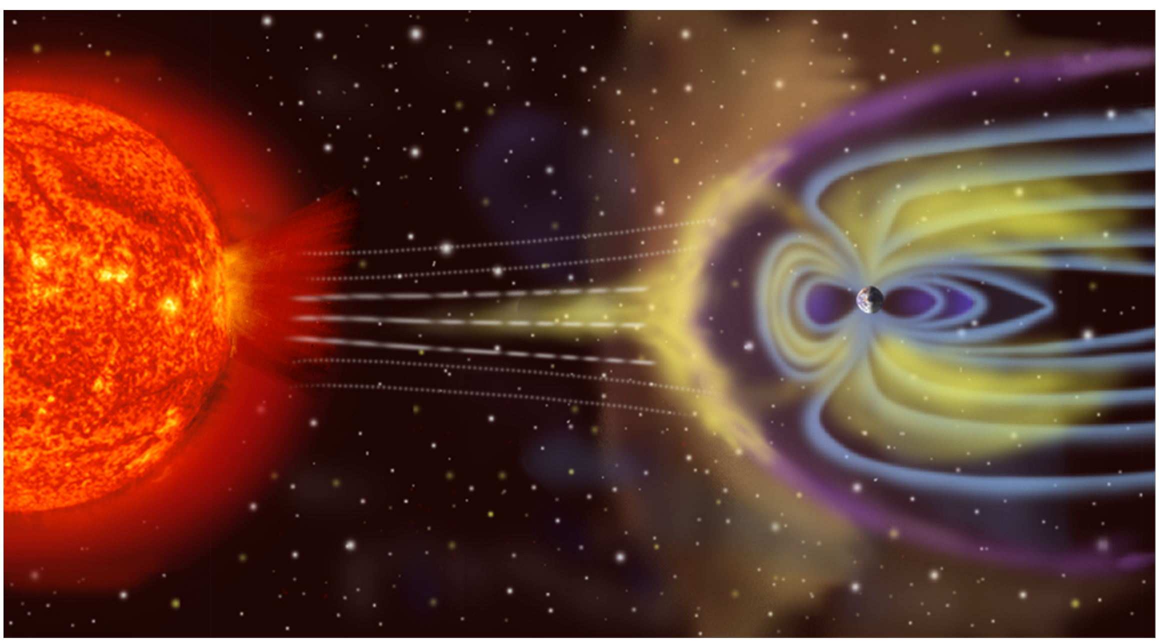

FIGURE 7.2 Image visualizing impact of solar corona mass ejection causing a GMD (not to scale). SOURCE: NASA Goddard Space Flight Center, http://sec.gsfc.nasa.gov/popscise.jpg.

HILF events are in another, almost unthinkable category in which outages could affect tens of millions for potentially months. But ignoring these threats will not make them go away. HILF events are a category where fundamental research in the mathematical sciences could yield good dividends. The event types in this category are different and they require unique solutions. However, they also have commonalities that the committee describes here in presenting some of the relevant mathematical and computational research challenges.

Interdisciplinary Modeling

The HILF events are all interdisciplinary and hence cannot be solved by experts from any single domain. GMDs start at the Sun, travel through space, interact with Earth's magnetic fields to induce electric fields at the surface that are dependent on the conductivity of Earth's crust going down hundreds of kilometers and that ultimately cause quasi-dc currents to flow in the high-voltage transmission grid, saturating the transformers, causing increased power system harmonics, heating in the transformers, and higher reactive power loss and resulting in a potential voltage collapse (NERC, 2012). In March 1989 a GMD estimated to have a magnetic field variation of up to 500 nT/min caused the collapse of Hydro-Québec's electricity transmission system and damaged equipment, including a generator step-up transformer at the Salem Nuclear Plant in New Jersey. More concerning is the potential for much larger GMDs, such as the ones that occurred in 1921 and 1859, before the development of large-scale grids, with magnetic field variations estimated to have been as much as 5,000 nT/min; such GMDs could cause catastrophic damage to different infrastructures, including the electric grid (Kappenman, 2012).

HEMPs have time scales ranging from nanoseconds to minutes. On the longer time scale of minutes, HEMP E3[1] is similar to an extremely large GMD, except with a faster rise time, requiring power system transient stability (TS) and TS-level modeling. Hence HEMPs would involve not only the disciplines surrounding GMD but also those surrounding the dynamics of nuclear explosions. A pandemic could affect a huge number of people, simultaneously impacting a large number of coupled infrastructures, including health, water, natural gas, and police and fire services. To defend against coordinated physical attacks would require a combination of power system knowledge and knowledge associated with the protection of physical assets, whereas defense against coordinated cyberattacks would need a combination of power system and cybersecurity domain knowledge. In modeling across different domains, each with its own assumptions and biases, mathematicians would be well positioned to help bridge the gaps between disciplines.

Rare Event Modeling

There is a need for research associated with HILFs in the area of rare event modeling. HILF events can be thought of as extreme manifestations of often more common occurrences. For example, while extreme GMDs are quite rare, more modest GMDs occur regularly, resulting in increasing quantities of data associated with their impacts on the grid. The same could be said for pandemics, while a large-scale attack on the grid would be a more severe manifestation of the disturbances (either deliberate or weather-induced) that occur regularly. The research challenge is extrapolation from the data sets associated with the more benign events.

Resilience Control Center Design

HILFs will stress the power system's cyberinfrastructure. This could come about as a result of either a direct cyberattack or the stressing of computational infrastructure and algorithms in ways not envisioned by their design specifications. As an example, one impact of a GMD (or a HEMP E3) would be increased reactive power consumption on the high-voltage transformers. However, existing state estimator (SE) models do not provide for these reactive losses. Hence it is likely that during a moderate to severe GMD the SE would not converge, leaving the

[1] The E3 component (a designation of the International Electrotechnical Commission, or IEC) of the pulse is a very slow pulse, lasting tens to hundreds of seconds, that is caused by the nuclear detonation heaving the Earth's magnetic field out of the way, followed by the restoration of the magnetic field to its natural place.

control center without the benefit of the other advanced network analysis tools. Another issue is the potential inundation of data in either the communication infrastructure or in the application software. For example, during the blackout of August 14, 2003, operators in FirstEnergy Corp.'s control center were overwhelmed with phone calls, whereas the Midcontinent ISO real-time contingency analysis experienced hundreds of violations (U.S.-Canada Power System Outage Task Force, 2004). Resilient control center software design and testing is a key area for future research. Effective visualization of stressed system conditions is also an important area for computational research.

Resilience Power System Design

Ultimately the goal of HILF research is to either eliminate the risk or reduce its consequences. As such, there are a number of interesting research areas to pursue depending on the type of HILF. Of course, a starting point for this work is the ability to have reasonable models of the events, and the economic impacts of all mitigations need to be considered. One promising area is the extent to which the impact of GMDs and HEMP E3s can be mitigated through modified operating procedures, improved protection systems, or GMD blocking devices on transformer neutrals. Algorithms for GMD blocking device placement could leverage advances in mixed-integer programming algorithms. The impacts of cyberattacks or physical attacks could be mitigated by adaptive system islanding. The deployment of more distributed energy resources, such as solar photovoltaics (PV), could make the grid more resilient if they were enhanced by storage capabilities or coupled with other, less intermittent resources to allow more of the load to be satisfied by potentially autonomous microgrids.

CASE STUDY IN DATA-CENTERED ASSET MAINTENANCE: PREDICTING FAILURES IN UNDERGROUND POWER DISTRIBUTION NETWORKS

Figure 7.3 illustrates the genesis of a manhole fire and its results. The oldest and largest underground power distribution network in the world is that of New York City. A power failure in New York can be a catastrophic event, where several blocks of the city lose power simultaneously. In the low-voltage distribution network that traverses a whole city underground, these events are caused by the breakdown of insulation for the electrical cables. This causes a short circuit and burning of the insulation, a possible buildup of pressure, and an explosion of a manhole cover leading down to the electrical cables, with fire and/or smoke emanating from the manhole. The power company

(a) (b) (c)

FIGURE 7.3 (a) Explosions are typically caused when a spark from wiring ignites gas inside the manhole. (b) Putting out a manhole fire. (c) Inside a manhole. SOURCES: (a) Kevin Bonsor, "How Exploding Manholes Work," August 30, 2001, ©HowStuffWorks, http://www.howstuffworks.com; (b) courtesy of Noah Berger; (c) Amanda Little, "Our Old Electric Grid Is No Match for Our New Green Energy Plans," October 13, 2009, Grist.org, http://grist.org/article/2009-10-13-our-old-electric-grid-is-no-match-for-our-new-green-energy-plans/ courtesy of Gina LeVay photography.

would like to predict in advance which manholes are likely to have such an event and to prevent it. There can be problems beyond the low-voltage network, for instance in the feeder cables of the primary distribution network, or in the transformers that step down the power between high and low voltage, in the transmission system, or in any other part of the system. If reliability-centered asset maintenance can be effectively performed, the number of outages and failures that occur in the city could be substantially reduced.

In each borough of New York City, the power company, Consolidated Edison (ConEdison), has been collecting data about the power network since the power grid started, at the time of Thomas Edison. Back then, these data were collected for accounting purposes, but now ConEdison records data from many different sources so the data can be harnessed for better power grid operations. Some of the types of data sets that ConEdison collects are as follows:

- *Company assets*. Data tables of all electrical cables, cable sections, transformers, connectors, manholes, and service boxes (access points to the energy grid), including their connectivity and physical locations, physical properties (e.g., manufacturer of the copper cable), installation dates, and other relevant information.
- *Trouble tickets*. Records of past failures or outages, sometimes in the form of text documents.
- *Supervisory control and acquisition (SCADA)*. Real-time measurements of the performance of equipment from monitors.
- *Inspection reports*. Records of each equipment inspection and the inspection results.
- *Programs*. Records of other preemptive maintenance programs, such as the vented cover replacement program, where solid manhole covers are replaced with vented covers that mitigate explosions, and the stray voltage detection program, where a mobile device mounted on a truck drives around the city and records stray voltage from already electrified equipment.

Discussed briefly below are some of the serious challenges in harnessing data from the past to prevent power failures in the future. See Rudin et al. (2010, 2012, 2014) for more details.

Data Integration

Data integration is a pervasive and dangerous problem that haunts almost all business intelligence. This is the problem of matching data records from one table to data records from another table when the identifiers do not exactly match. For instance, if the aim is to determine which electrical cables enter into which access points (manholes, service boxes) in Manhattan, a raw match without additional processing would miss over half of the cable records. Given that there is enough electrical cable within Manhattan to go almost all the way around the world, this data integration problem could lead to severe misrepresentation of the state of the power system. Data integration can be severely problematic generally. For one thing, companies need to locate records that provide a full view of each entity. They would like to know, for instance, that inspection reports detailing a particular faulty cable in a particular manhole are connected to customer complaints in a particular building, but there are many ways that this can go wrong: A cable identifier, manhole identifier, or street address that is mistyped in any of the tables could cause this connection to be missed.

One way to handle this problem is to create a machine learning classification model for predicting high-quality matches between two records from different tables. Let $\mathbf{x}$ be a vector of a pair of entities, one from each of the two tables to be joined. For example, consider cables and manholes where the three manhole identifiers are (1) type (manhole or service box), (2) number (e.g., 1,624), and (3) mains and service plate (M&S) for a three-block region of New York City. Let $x_{i1} = 1$ if there is an exact match between all three identifier fields, let $x_{i2} = 1$ if there is an exact match between the manhole types and numbers and the M&S plates are physically close to each other, etc. Given a sample of labeled pairs, where $y_i = 1$ when the match is correct and $y_i = 0$ otherwise, a classification problem can be formed as described in Chapter 4.

Handling Unstructured Text

Much of the data generated by power companies is in the form of unstructured text. The data could include trouble tickets, inspection reports, and transcribed phone conversations with customers. The field of natural language processing involves using sophisticated clustering techniques, classification techniques, and language models to put unstructured text into structured tables that can be used for business intelligence applications. ConEdison, for instance, has generated over 140,000 free text documents describing power grid events over the last decade within Manhattan. These text documents contain the main descriptions of power grid failures on the low-voltage network and thus are a key source of data for power failure predictions. If these text documents can be translated into structured tables that can be used within a database, these text documents can become extremely valuable sources of data for studying and predicting power failures.

Rare Event Prediction

Many classification techniques (such as logistic regression) can fail badly when the data are severely imbalanced, meaning there are very few observations of one class. Power failures are rare events, so it can be difficult to characterize the class of rare events if very few (or none) have been observed. If failures happen only 1 percent of the time, a classification method that always predicts no failure is right 99 percent of the time, but it is completely useless in practice. This problem of imbalanced classification is discussed next.

Causal Inference

Many power companies are starting to take preemptive actions to reduce the risks of failure. These actions could include, for instance, equipment inspections or preemptive repairs. To justify the expenses of these programs, one must estimate the benefits they provide. Without such estimates, it is unclear how much benefit each program creates or indeed whether there is a benefit. For instance, on the New York City power grid, a study (Rudin et al., 2012) called into question the practice of high potential (hipot) testing on live primary distribution cables. Hipot testing is where a live cable is given a much-higher-than-usual voltage, under the assumption that if the cable is weak it is more likely to fail during the test and can thus be replaced before it fails during normal operation. The problem is that the test itself can damage the cable. Other examples are manhole inspection programs and vented manhole cover replacement programs: To justify the costs of these programs, one needs to estimate their effectiveness. In this case, where the test itself does damage, predicting failures does not suffice; one needs to predict what would have happened to untreated cases had they been treated, and one needs to predict what would have happened to treated cases had they not been treated (the counterfactual).

Visualization and Interpretation of Results

Visualization of data is a key aspect of the knowledge discovery process. With ever more complex information arising from the power system, new ways of making sense of it are needed. For instance, for data from a distribution network such as New York's, it is useful to visualize aspects of the electrical cables, manholes, geocoded locations of trouble tickets where problems arise, inspections, and more. Modern visualization tools can be interactive: One can probe data about local areas of the power grid or explore data surrounding the most vulnerable parts of the grid. One particular type of tool designed for New York City is called the "report card" tool (Radeva et al., 2009). With this tool, an engineer can type in the identifier for a manhole and retrieve an automated report containing everything that must be known to judge the vulnerability of the manhole to future fires and explosions.

Machine-Learning Methods Comprehensible to Human Experts

Most of the top 10 algorithms in data mining (Wu et al., 2008) produce black-box models that are highly complicated transformations of the input variables. Despite the high prediction quality of these methods, they are often not useful for knowledge discovery because of their complexity, which can be a deal breaker for power grid engineers who will not trust a model they cannot understand.

It is possible that very interpretable yet accurate predictive models do exist (see Holt, 1993, for instance). However, interpretable models are often necessarily sparse, so finding them is computationally hard. There is a fundamental trade-off between accuracy, interpretability, and computation; current machine-learning methods are very accurate and computationally tractable, but with tractability trade-offs or statistical approximations to reduce computation, it may be possible to attain models that are more interpretable and even more accurate.

The challenges above are not specific to New York; they are grand challenges that almost every power company for a major city faces. Solutions to the problems discussed here can be abstracted and used in many different settings.

CASE STUDY IN SYNCHROPHASORS

Hurricane Gustav made landfall near Cocodrie, Louisiana, at 9:30 a.m. CDT on September 1, 2008, as a strong category 2 storm (based on 110 mph sustained winds) and a central pressure of 955 millibars.[2] As usually happens with these types of events, there was significant damage to both electric transmission and distribution infrastructure. An example of the devastation is shown in Figure 7.4.

For Entergy, the electric utility company operating in this area, Hurricane Gustav caused the second largest number of outages in company history, behind only Hurricane Katrina. Gustav restoration rivals the scale and difficulty of Hurricane Katrina restoration.[3] Unlike for previous storms, however, Entergy was able to utilize cutting-edge measurement technology to facilitate the restoration of its system. As the storm disrupted individual circuits, an electrical island was formed within Entergy's service territory. What this means is that some generators were serving load using infrastructure that was electrically separated from the remainder of the interconnected power grid. Historically, this situation would have been difficult to manage in the control room, and it would likely have required de-energizing the loads, connecting the generators to the remainder of the grid, then reconnecting the load in the restoration sequence of events. However, because Entergy had previously deployed synchrophasor technology in its control room, the system operators were able to better observe the operation of the electrical island and utilize this information to facilitate its reconnection with the remainder of the grid as an intact electrical island.

Overview of Synchrophasors

As discussed in earlier chapters, a synchrophasor is a time-synchronized measurement of an electrical quantity, such as voltage or current. In addition to measuring the magnitude of the quantity being measured, the accurate time reference also measures the phase angle of that quantity. The enabling technology underlying this measurement approach is an accurate time reference. One common and ubiquitous time reference is the Global Positioning System, which provides microsecond-class timing accuracy. This is sufficient to measure phase angles with better than 1° accuracy. (For example, if the user desires to measure the angle with 1° accuracy on a 60-Hz system, the time error must be less than 4.6 µsec.)

The phasor measurement unit (PMU) can also calculate derived parameters associated with other electrical quantities, including frequency, rate of change of frequency, power, reactive power, and symmetrical components, by processing the raw voltage and current information that is measured by the instrument. Widely adopted standards, such as IEEE C37.118.1, govern the definition of these measurements. There are also different classes of PMUs that have been defined based on whether speed or accuracy is the primary consideration, given different assumptions that can be made by the equipment vendor for sampling and filtering algorithms. The M-class, for measurement, emphasizes accuracy, while the P-class, for protection, emphasizes speed of detection, which may sacrifice steady-state accuracy. Future modifications to these standards are defining dynamic performance requirements.

[2] National Weather Service Weather Forecast Office, "Hurricane Gustav," last modified September 1, 2010, http://www.srh.noaa.gov/lch/?n=gustavmain.

[3] Entergy, "Hurricane Gustav," http://entergy.com/2008_hurricanes/gustav_video_2.aspx. Accessed December 15, 2015.

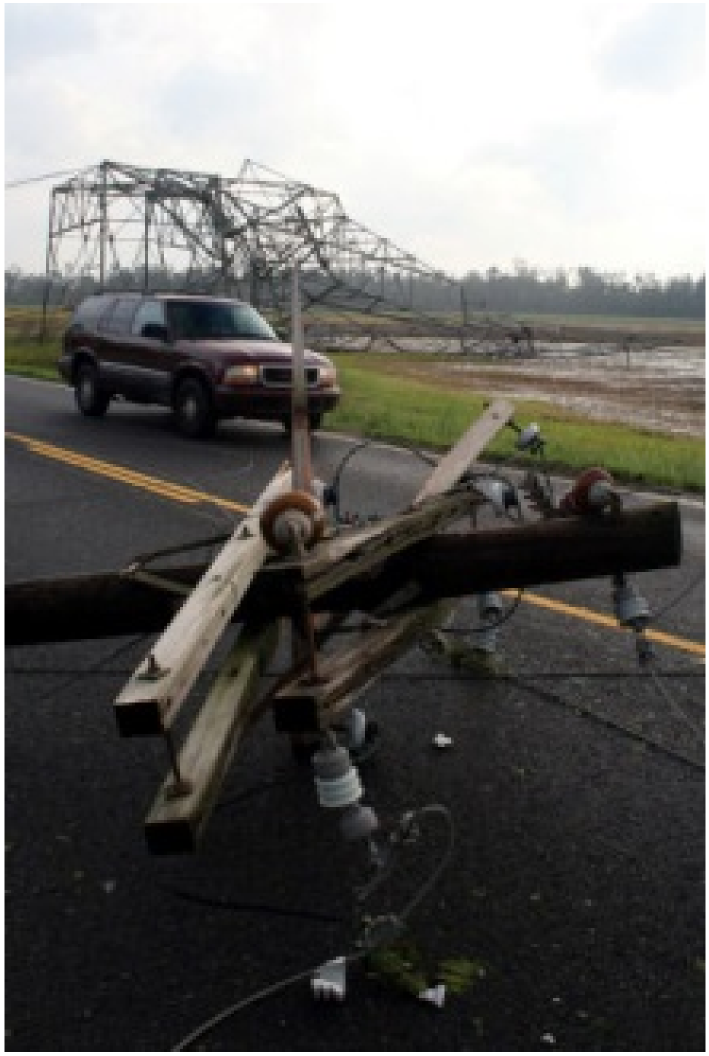

FIGURE 7.4 An example of transmission (background) and distribution (foreground) electrical infrastructure damage associated with the 2008 Hurricane Gustav in southern Louisiana. SOURCE: Entergy, “Images – Gustav Damage,” image gallery, http://www.entergy.com/2008_hurricanes/gustav_media.aspx. Courtesy of Entergy.

There are other benefits of synchrophasors beyond those achievable from traditional measurements that are provided by SCADA telemetry. Because PMUs provide data with multiple frames per second (a modern PMU is capable of measuring at least 30 samples per second), dynamic characteristics of the power system can be measured. This is a valuable data source to calibrate dynamic power system models. Furthermore, accurately time-stamping the measurements can aid in the investigation of system disturbances (blackouts).

Internationally, the use of synchrophasors has been increasing dramatically in the past several years. After the technology was adopted and proven by early adopters over the past few decades, and with the cost of the technology steadily decreasing, more and more operational entities have adopted the technology. Some applications are given next.

Application of Synchrophasors

One of the first applications of this technology was to support planning engineers. Having high-speed, time-stamped data was helpful for calibrating and validating dynamic models of the power system. New insights were gleaned concerning the dynamic behavior of the grid. Additionally, blackout investigations made extensive use of these measurements whenever they were available. The key attributes of the measurements sought for these applications were that they were high speed and time stamped.

One of the early applications in the power system control room was visualization to provide operators enhanced wide-area situational awareness. Because the relative phase angles between different regions of the power grid are directly proportional to the real power flowing across the network, displaying the phase angles across a wide-area power system depicts the power flowing across the network in a comprehensive and intuitive manner.

Also, because it is also affected by the net impedance between different points in the network, the phase angle can also serve as a proxy for system stress across critical boundaries. For example, given a constant power transfer across a corridor, if one of the lines is removed from service, the angle across the corridor will increase. Some utilities have adopted alarms and alerts for their operators based on measured phase angles.

Bringing synchrophasors directly into the state-estimation process can also improve the accuracy of those tools. Some utilities have deployed hybrid state estimation, where synchrophasor data are added to SCADA data in the state estimation, where others are evolving toward linear state estimators that are fed solely from PMUs. The linear state-estimation process can reduce measurement error by fitting the measured data to a real-time model of the power system.

More advanced applications are investigating the use of synchrophasors as inputs to Special Protection Systems. These schemes trigger automated responses based on real-time changes to system conditions. The synchrophasor data can arm the system and can also be used to trigger an automated response if that is appropriate.

Today PMUs are deployed primarily on the transmission system, but the industry is beginning to explore their use at the distribution level for power quality, demand response, microgrid operation, distributed generation integration, and enhanced distribution system visibility.

Mathematical Challenges to Improve Synchrophasor Measurements

Today's synchrophasor measurement systems are governed by industry standards that define their accuracy requirements.[4,5] However, these accuracy requirements are only defined for steady-state measurements. In an attempt to reconcile the different applications of the measurements and how different vendors would make trade-offs in their sampling and filtering algorithms associated with speed and accuracy of the measurements, different classes of synchrophasor measurements have been defined. The so-called M-class (measurement) provides a more accurate estimate of the measurement but is allowed to take longer to converge on the measured value. The P-class (protection) is designed to operate faster and is primarily intended to quickly assess the new state of the system after a change in conditions, such as would occur during a fault or other system change. However, neither aforementioned class of measurements will necessarily provide consistent results between different vendor products for continuously time-variant conditions, such as a persistent dynamic instability, or in the presence of other imperfections in the measured signal, such as harmonics. Part of the challenge is that the entire premise of defining what a synchrophasor is applies only to a steady-state representation of the power system, and the changes are neither consistently nor comprehensively well defined. For example, the relationship between phase angle and frequency is not clearly defined whenever either of these parameters is changing. In much the same way that advanced mathematical algorithms are used to extract weak signals from a noisy environment in the communications domain, there is an opportunity for algorithmic advancement to provide a better foundation for extracting meaningful signals from power system measurements, particularly those associated with dynamical systems.

---

[4] IEEE C37.118.1-2011 (IEEE Standard for Synchrophasor Measurements for Power Systems) and C37.118.1a-2014 (IEEE Standard for Synchrophasor Measurements for Power Systems—Amendment 1: Modification of Selected Performance Requirements).

[5] International Electrotechnical Commission (IEC) IEC 67850-90-5.

CASE STUDY IN INVERTER-BASED CONTROL FOR STABILIZING THE POWER SYSTEM

The committee considered two cases of power grid instability that could have been avoided with better analytical and mathematical tools. The first example is in Texas, where wind power farms in northwest Texas were producing power that is carried by weak transmission lines to the large load centers in east Texas (Dallas, Austin, Houston, San Antonio, and others). The turbines and the cables both have built-in controls to help dampen oscillations, in particular, in (1) the thyristor-controlled series capacitor (TCSC) transmission lines, which means that their line power flow can be directly controlled, and in (2) the doubly fed induction generators (DFIGs) of wind farms whose voltage is electronically, rather than mechanically, controlled. If any electrical signals vary from the control center's reference settings, this needs to be remedied very quickly. The cables and the wind farms are equipped with fast electronic inverter-based controls, which change the stored energy in the equipment whose power is electronically controlled to push the signal back toward its reference settings. However, the controls on the Texas equipment did not work properly, and this led to oscillatory dynamics between the controllers of wind power farms and line flow controllers of weak transmission lines delivering wind power to the faraway loads such as Dallas. The new technical term for these instabilities is subsynchronous control instabilities, which had not been experienced by any power grid before the situation in Texas. For details of operational problems related to large wind power transfer in Texas, see ERCOT System Planning (2014).

Similarly, in Germany, by government regulation, all of the wind power produced in the northwest of Germany must be delivered by the grid. However, the German power grid is not strong enough to handle this massive variability nor is it controlled online. Because of this, it is not always possible to deliver wind power to the major cities in the south of Germany (Munich in particular). Instead, power spills over to the Polish and Czech power systems, which complain about this and wish to build high-voltage dc tie links to block the German wind power from entering their grids. In addition, a serious problem of harmonic oscillations, similar to the problem observed in Texas, has been observed.

Situations like those in Texas and Germany could be avoided in the future if analytical capability in inverter-based control could be advanced—that is, the fast calculations performed in response to signals deviating from their reference settings. A lot of technology currently being developed will require inverter-based control. Mature versions of power inverter control are the automatic voltage regulators and power system stabilizers, both controls for exciters, of conventional power plants. More recent inverter control is being deployed for storage control of intermittent power plants, such as DFIGs and flywheels placed on wind power plants; for real and/or reactive power line flow and voltage control of series controllers, such as TCSCs and shunt capacitors (static var compensators); for control of storage placed on PVs; and for control of variable-speed drives ubiquitous to controllable loads, such as air conditioners, dryers, washers, and refrigerators. Recently there have been large investments in better switches, such as silicon nitride switches. For example, the National Science Foundation's Energy Research Center for Future Renewable Electric Energy Delivery and Management (FREEDM) system works on designing such switches and using them to control substation voltages and frequencies (http://www.freedm.ncsu.edu), and there are several efforts to design more durable and compact switches with higher voltage levels (ARPA-E's GENI program is one).

The basic role of inverter control is unique in the sense that it is capable of controlling very fast system dynamics; the cumulative effects of kilohertz rate switching are capable of stabilizing fast frequency and voltage dynamics that are not otherwise controllable by slower controllers, in particular power plant governors. EPRI has led the way to Flexible AC Transmission Systems (FACTS) design for several decades. Interestingly, the early work made the case for using FACTS to control line reactances, and, as such, being fundamental to increasing maximum power transfer possible by FACTS-equipped transmission lines. The decrease in line reactance directly increases power transfer by the line. More recently, there has been major research and development aimed at inverter-based control for microgrids, which is based on placing inverter controls on each PV and directly controlling reactive power-voltage (Q-V) and real power-angle (P-theta) transfer functions of closed-loop PVs (Consortium for Electric Reliability Technology Solutions-microgrid concept, http://certs.aeptechlab.com). Similarly, when modeling inverter-based storage control (flywheels, DFIGs) it is assumed that voltage/reactive power and real power/energy can be controlled directly by inverters so that the closed-loop model is effectively a steady-state droop characteristic. An emerging idea is that inverter-based control placed on direct-energy resources could be used to ensure stable response of power systems with massive deployment of intermittent resources; in effect, inverter-based control could replace inertial response of governor-controlled conventional power plants.

The approaches to stabilization in future power grids require careful new modeling and control design for guaranteed performance. As shown by the examples in Texas and Germany discussed above, at the lower distribution grid level, today's inverter control practice of maintaining the PV power factor at unity has been known to result in unacceptable deviations of voltages close to the end users.

The problems in Texas and Germany are only early examples of the problems that could be caused by poor tuning of inverter control. They point to the need to model the dynamics relevant for inverter control to the level of detail necessary so that controllers are designed for provably stable response to each given range and type of disturbance. Some challenges are as follows:

- *Modeling realistic fast dynamics*. Most models currently used in control centers do not even attempt to model the fast dynamics relevant for assessing the performance of power electronically switched automation embedded in different components throughout the complex power grids. This is a very difficult problem since it requires accurate modeling of fast nonlinear dynamics and control design, which are often close to bifurcation point conditions. Some recently reported theoretical results on this topic were derived under highly unrealistic assumptions, such as "real-reactive power decoupling"—that the grid is entirely inductive (which is not possible when one relies on capacitive storage for voltage/reactive power control)—and that the loads are simple constant impedance loads. Modeling the fast dynamics with realistic assumptions and in a computationally fast way would be a big step forward.
- *Aggregation of small inverter controllers*. Another problem in power grids still to be studied concerns modeling dynamical effects of aggregate small inverter controllers on closed-loop dynamics in the grid. Modeling and designing switching control to avoid the real-world problems described above in using power electronics represents a grand challenge for modeling and computational methods. This challenge must be addressed if benefits from hardware improvements in power electronic switching are to be realized without excessive cost.

REFERENCES

Abi-Samra, N., J. McConnach, S. Mukhopadhyay, and B. Wojszczyk. 2014. When the bough breaks: Managing extreme weather events affecting electrical power grids. *IEEE Power and Energy Magazine* 12(5):61-65.

DOE (U.S. Department of Energy). 2012. *A Review of Power Outages and Restoration Following the June 2012 Derecho*. August. http://www.oe.netl.doe.gov/docs/Derecho%202012_%20Review_080612.pdf.

ERCOT System Planning. 2014. *Panhandle Renewable Energy Zone (PREZ) Study Report*. http://www.ercot.com/content/news/presentations/2014/Panhandle%20REnewable%20Energy%20Zone%20Study%20Report.pdf.

Fink, L.H., and K. Carlsen. 1978. Operating under stress and strain. *IEEE Spectrum* 15(3):48-53.

Holt, R.C. 1993. Very simple classification rules perform well on most commonly used datasets. *Machine Learning* 11(1):63-90.

Kappenman, J. 2012. A perfect storm of planetary proportions. *IEEE Spectrum*, February. http://spectrum.ieee.org/magazine/2012/February.

NERC (North American Electric Reliability Corporation). 2010. *High-Impact, Low-Frequency Event Risk to the North American Bulk Power System*. A Jointly-Commissioned Summary Report of the North American Electric Reliability Corporation and the U.S. Department of Energy's November 2009 Workshop. June. http://www.nerc.com/pa/CI/Resources/Documents/HILF%20Report.pdf.

NERC. 2012. *Effects of Geomagnetic Disturbances on the Bulk Power System*. Special Reliability Assessment Interim Report. February. http://www.nerc.com/files/2012GMD.pdf.

Radeva, A., C. Rudin, R. Passonneau, and D. Isaac. 2009. Report cards for manholes: Eliciting expert feedback for a machine learning task. Pp. 719-724 in *International Conference on Machine Learning and Applications, 2009. ICMLA '09*. doi:10.1109/ICMLA.2009.72.

Rudin, C., R. Passonneau, A. Radeva, H. Dutta, S. Ierome, and D. Isaac. 2010. A process for predicting manhole events in Manhattan. *Machine Learning* 80.

Rudin, C., D. Waltz, R.N. Anderson, A. Boulanger, A. Salleb-Aouissi, M. Chow, and H. Dutta. 2012. Machine learning for the New York City power grid. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 34(2): 328-345.

Rudin, C., S. Ertekin, R. Passonneau, A. Radeva, A. Tomar, B. Xie, and S. Lewis. 2014. Analytics for the power grid distribution reliability in New York City. *Interfaces* 44(4).

Stockton, P. 2014. *E-Pro Handbook*. Executive Summary. Electric Infrastructure Security (EIS) Council. http://www.eissummit.com/images/upload/conf/media/EPRO%20Handbook%20Executive%20Summary.pdf.

U.S.-Canada Power System Outage Task Force. 2004. *Final Report on the August 14, 2003, Blackout in the United States and Canada: Causes and Recommendations*. http://energy.gov/oe/downloads/us-canada-power-system-outage-task-force-final-report-implementation-task-force.

Wu, X., V. Kumar, R. Quinlan, J. Ghosh, Q. Yang, H. Motoda, and G. Mclachlan. 2008. Top 10 algorithms in data mining. *Knowledge and Information Systems* 14:1-37.

8

Building a Multidisciplinary Research Community

INTRODUCTION

Developing the foundational capabilities for understanding and analyzing the next-generation electric power grid is a multidisciplinary endeavor. A sample of the disciplines required includes electric power engineering, mathematics, statistics, operations research, computer science, and economics. Each of these disciplines in turn could have subdisciplines that would be appropriate for specific problems—for example, optimization, nonlinear dynamics, machine learning, or databases. And while some multidisciplinary teams form naturally through mutual collaborations, a more strategic approach will be required to build a more effective multidisciplinary community to address the challenges described in this report. Interestingly, when these partnerships do form, it is often the case that the use of well-known techniques from one discipline (which may not perhaps be as well known in another discipline) has yielded a breakthrough—as, for example, the use of interior-point optimization methods for solving the large linear programming problems that arise in solving the optimal power flow problem.

However, adapting and disseminating state-of-the-art algorithms and methods from other disciplines to the electric power systems community and developing entirely new research areas that build on the joint strengths of two or more communities are challenges. Power engineers have to formulate the problems and help the rest of a multidisciplinary team understand the underlying issues and their nuances. This can be difficult, as power engineers frequently do not have the background or training that would allow them to articulate their problems in the language of another discipline. How would these engineers know that the other discipline has methods or expertise that could be brought to bear on the problem of interest? Conversely, experts from other disciplines such as mathematics need to acquire requisite background for understanding the language of power engineers. There are many examples where people with a core competency in a related area (e.g., control or linear systems) formulate problems in their own language that are consistent with what they know and can solve but that fail to capture the key issues or the subtleties that differentiate a useful, pragmatic approach from one that is of only theoretical interest. So, a common framework needs to be created and researchers educated across disciplines so they become fluent in one another's language and approaches to problem solving and become familiar with state-of-the-art methods in the other disciplines. One strategy is to use case studies to provide insights into problems of common interest, because such studies give concrete details that allow the understanding of heterogeneous groups of researchers to converge.

While it is clear that all the contributors to cross-disciplinary teams must understand the underlying problems, it is equally important that they be cognizant of a broader context. Take the example of wind forecasting. Power system engineers often cite the need for better forecasts so that the availability of wind power can be estimated for a particular service provider. Recent research has contributed to improving such forecasts using the techniques of data assimilation described in Chapter 4. However, the atmosphere has long been regarded as a chaotic system that is very sensitive to initial conditions. This suggests that there are inherent limits to the predictability of the weather. Ensemble forecasting based on this assumption produces probabilistic forecasts rather than specific predictions of wind velocity or cloud cover that would affect renewable energy resources. So while we may continue to work toward more precise forecasts, as desired by the power systems community and many others, multidisciplinary research teams can recognize that wind models for the grid must be treated probabilistically, and that they will produce results with uncertainty and can have larger variance than we desire.

How then should it be determined which disciplines and subdisciplines could contribute to the analytic foundations for the next-generation electric grid? As all the problems have multidisciplinary dimensions, it is important to determine which dimensions are more important or more difficult for each individual application. One can then determine the strategy appropriate for each problem. In other words, just because teams engaged in research on power systems have members from different disciplines does not ensure their ability to solve the specific issues that are facing the grid like those outlined in the previous chapters. For that multidisciplinary teams need to be brought together that have individuals with backgrounds in the specific subdisciplines that are needed for each application. One example from the 1990s can perhaps illustrate some of the key characteristics of a successful multidisciplinary team.

EXAMPLE OF A MULTIDISCIPLINARY TEAM: PSERC

By 1988 it became clear that fundamental and lasting changes were coming to the electricity business because of the restructuring that the Federal Energy Regulatory Commission was beginning to mandate. But at that time, there were very few power engineering programs left at major research universities, raising concerns about how new technologies necessary for the grid changes would be developed. In 1995 five universities (Cornell, University of Illinois, the University of California at Berkeley, Howard University, and the University of Wisconsin) were able to establish a National Science Foundation (NSF) industry/university cooperative research center, the Power Systems Engineering Research Center (PSERC).[1]

The fundamental premise of PSERC was that no single university had the breadth of expertise needed to address the research issues. In addition, it was recognized that multiple disciplines would have to be brought together. Quoting from the original 1995 proposal, "The Center's . . . basic premise is that engineering considerations should not be an afterthought but rather be a principal force in the planning process of a restructured Industry. Consequently, the center's agenda differs from the institutional research programs currently in operation in that it focuses on the technological needs (such as computational methods, information needs and protocols, and control schemes) that are fundamental to successful implementation of alternative economic paradigms for opening the power system to greater competitive forces. It also provides stronger organizational and personal structure that encourages interdisciplinary research and communication among engineers, economists, and computer scientists." When PSERC was established in 1995, it was well positioned to participate in work needed to implement the landmark transition set in motion by the April 1996 release of FERC orders 888 and 889.

EXAMPLE OF A MULTIDISCIPLINARY EFFORT: MARKETS

After PSERC was established it began to expand the number of its member schools in order to capture the expertise needed to address the broad array of problems envisioned by the founders. In addition it organized its programs into three "stems"—namely, markets, systems, and transmissions and distributions (T&D) technologies. Within the markets stem were engineers and economists, most of whom had not collaborated before. The first hurdle to be overcome by this collaboration was language. While mathematics is a universal language common

[1] PSERC now comprises some 13 schools, and it still serves as an education and research resource for the electric power community (see the PSERC website at http://www.pserc.wisc.edu/home/index.aspx).

to many disciplines, the context in which it is to be applied is not. Even nomenclature had to be settled on. For example, in the power literature it is common to label real and reactive power P and Q, whereas to an economist P and Q are price and quantity. In order to facilitate communication across disciplines, a test-bed platform called PowerWeb was created that could be used to coordinate the work across disciplinary boundaries as well as to test new market design concepts. The premise was that the experimental economic concepts pioneered by Nobel laureate Vernon Smith, together with a more realistic engineering representation of the power system, would reveal interesting and useful insights about market design. A substantial problem for economists was that there was very little experience with or insight into repeated auctions of the complexity that were being contemplated at the time. It was clear that, because the auctions were repeated, learning was possible and even necessary. If market participants can learn, they can then presumably optimize their position by learning the agendas of other market participants. It is not necessary to openly collude to learn what behavior is best for maximizing profits given tolerance for risk. Economists were of the opinion that the best-designed markets are those designed to reveal true costs of the participants. Because of the ever-changing operational equilibrium of the delivery network, the Nash equilibrium does not exist. The new techniques devised by the engineer/economist collaboration were able to address important tasks such as the following:

- Explain the origin of price spikes that were occurring,
- Identify where pockets of market power might be,
- Quantify the advantages of having a demand-side market,
- Determine the number of participants needed for a supply-side-only market to be competitive,
- Observe that a distributed unit commitment schedule may be as efficient as an optimization-based centralized commitment, and
- Explore a host of other interesting phenomena that would not have been possible without the cross-disciplinary collaboration.

One of the important findings from these multidisciplinary collaborations was that having a complete engineering model was crucial in determining the economic outcome of any of the many market designs being discussed. It is interesting to note that today, most Independent System Operators have market-testing platforms that are used to test new concepts before market participants experience them.

Another problem that surfaced in PSERC's cross-disciplinary research was how to efficiently collaborate across multiple institutions. Successful collaborations of this type involve overcoming organizational and social differences and establishing the same kind of trust that comes from working with someone down the hall who is more easily accessible. To some extent, the Internet communication technology that was then newly evolving helped ease some of these problems.

Today, students who graduated from programs that crossed the power engineering/economics boundary are professors in various departments around the country, teaching what is now an integrated discipline. Some professors are in electrical engineering, some in operations research, some in economics, and some in other departments. But they are in communication with each other and publish in the same journals—hallmarks of a well-functioning research community. The industry and the country as a whole have benefited from the interdisciplinary marriage of these once disparate disciplines. Most important, this community of power systems economists is focused on solving problems directly related to the grid.

While PSERC has contributed in important ways to building a multidisciplinary research community that supports the electric grid, its mix of expertise does not extend to many of the areas of importance to developing the analytical and mathematical capabilities that will be needed for the next-generation grid. For example, computational tools to address the ac optimal power flow problem discussed throughout this report are likely to require fundamental advances in optimization that are typically outside the domain boundaries of PSERC. As documented throughout this report, many of the analytical challenges require insights from the mathematical sciences, which in turn call for new cross-disciplinary connections that are currently scarce.

EXAMPLES FROM OTHER DISCIPLINES

Earlier chapters of this report identified the following main analytical challenges for the future grid:

- Making effective use of large data from improved measurements,
- Modeling the availability of uncertain renewable energy resources and their effect on grid reliability,
- Building and operating smart grids that incorporate demand response, and
- Improving optimization methods for nonlinear, nonconvex, and stochastic problems.

These challenges call for new multidisciplinary research communities that draw from mathematics, computational science, computer science, operations research, statistics, and control theory. Since these multidisciplinary research communities are still emerging, we should look broadly at examples from other disciplines for insight into how to best form these teams and enable their effective operations. Several types of existing models are relevant.

The Department of Energy (DOE) has long-standing experience in developing programs that span multidisciplinary groups, including the Scientific Discovery through Advanced Computing program and the Advanced Simulation and Computing Program. Both programs were designed to build the simulation capabilities needed by computational scientists and engineers to make effective use of the vast computational and data resources provided by the DOE laboratories. Finally, the Mathematical Multifaceted Integrated Capabilities Center is another good example of building multidisciplinary teams. These centers have a strong focus on the mathematical sciences, as described in DOE Program Announcement 12-698: "These science and engineering challenges must be abstracted into an interrelated set of mathematical research challenges that require new integrated, iterative processes across multiple mathematical disciplines."

The Mathematics Climate Research Network (MCRN), which started in 2010 with support from the NSF, provides another model for fostering multidisciplinary collaborations between mathematicians and scientists in another discipline. Climate science is a field that is already organized around large data, comprehensive computational models, and the use of high-performance computers. Thus, a large initial time investment has been required of individuals whose research is in the area. The goal of MCRN has been to reduce the barriers to engaging mathematical scientists in problems emerging from the study of the climate system. One of the motivations was to significantly increase the number of mathematical researchers working in this area. The strategy was to create a community of individuals who would support and inform one another in defining key mathematical directions and challenges while pushing the resulting research to a high scientific level. Part of the effort has focused on bringing junior people into the area through the design of effective training elements incorporated into existing mathematics graduate programs at MCRN member institutions. Since the collection of participating researchers is widely distributed, the network has utilized and further developed web-based tools for communication, conferencing, and collaboration. The network has grown to over 200 individual members and makes extensive use of web-based collaboration tools, computational sharing capabilities, and communal data and software storage. Although studying the electric grid was not a primary objective of the network, a research group emerged in 2014 with a focus on determining the mathematical challenges posed by the next-generation grid. This small group represents, in embryonic form, a web-based effort to increase the participation of mathematical scientists in this area.

While the analogy between power systems engineering and the atmospheric and climate sciences is a good one, there are significant differences that bear upon the organizational structures that the committee recommends for the future electric grid. In the atmospheric sciences, government agencies have long gathered weather data and amalgamated these data into publicly available databases. Large-scale modeling efforts for climate models and numerical weather prediction thus have a base of publicly available data resources that serve to coordinate and ground atmospheric modeling research. However, an analogous step for grid modeling research is challenging because a lot of the data are proprietary or are protected for homeland security purposes and designated as Critical Infrastructure Protection (CIP) data.

The field of genomics provides an excellent illustration of how databases can be central to a scientific research area. When the visionary Human Genome Project was initiated to discover the complete DNA sequence of humans and other organisms, its leaders made several astute decisions that have been important factors in the scientific success and impact of the field. Here we mention two. First, the leaders mandated that a comprehensive database would be created and that researchers would be required to add their results to this database in a timely manner. Second, they invested in the development of software tools that would enable all biologists to utilize the information in this database. The National Library of Medicine, part of the National Institutes of Health, was given responsibility for both. It created research groups and tasked software development staff to deal with the data. The impact of these decisions has been phenomenal. All of biology has benefited from the advances in DNA sequencing technology and algorithmic methods for sequence analysis. One of the key provisions was the availability of open and comprehensive data sets that modelers and algorithm and software developers could use for their research.

RECOMMENDATION FOR SYNTHETIC DATA LIBRARIES

Climate research and genomics illustrate the central role that data play in 21st-century science. The electric power industry is poised to make this transition in data intensity, but it has not yet settled on an organizational structure that makes effective use of its data. The sensitive nature of the data (including CIP and proprietary concerns) and the complex manner in which the industry operates call for careful and secure data management before any real data can be released in general to researchers. No entity has yet assumed responsibility for this task on a national level, but creating a center to undertake it would have huge potential benefits to the economics and reliability of the next-generation grid. Below, the committee proposes the establishment of such a multifaceted center.

Any researcher intending to work on a new problem area finds that familiarization with the state of the art in the new area is a difficult task. In the area of power engineering, an additional hurdle, as pointed out in Chapter 6, is the scarcity of data that are representative of the real power grid, including for example non-CIP data. This is true of both the data that describe the power grid as well as the measurement data under various operational conditions. In addition, in developing novel ideas for next-generation algorithms one always goes through a sequence lasting years, where at the beginning the algorithms fail most of the time before finally evolving into ones that are robust. Data that reflect real conditions are necessary for moving this process forward. In this context, it is critical that accessibility to real data not be the limiting factor in developing new algorithms. There may be many reasons why real data are not used in research, and where this is the case we should strive to develop and make available high-quality and accurate synthetic data.

One solution to this problem is to develop synthetic data that exhibit the same behavioral characteristics as real power systems of realistic size, perhaps using some of the tools described in Chapter 6. However it is done, developing synthetic data will require substantial effort so that they can be used effectively as a surrogate for real data—for example, by capturing real behaviors and responding as a real grid would to what-if scenarios.

A recent Funding Opportunity Announcement from ARPA-E asks for the development of synthetic data to be used to test optimal power flow methods. This is a good start, but it is necessary to develop several libraries because not all power systems exhibit all the possible behaviors. There is diversity in both the kinds of data needed to investigate different types of problems and the data based on network architecture. For example, a highly stable system like the Eastern Interconnection is tightly coupled, while a system with longer transmission lines and looser coupling (e.g., the Western Interconnection) can exhibit more dynamic problems. Synthetic data sets should exhibit all the different characteristics that researchers might want to study, without duplicating existing sets.

Recommendation 9: The Department of Energy should sponsor additional efforts to create synthetic data libraries to facilitate studies of, and tool building for, the reliability and control of the future electric grid.

RECOMMENDATION FOR SOFTWARE LIBRARIES

In addition to having access to synthetic data, one must be able to simulate portions of the grid so as to study the various behaviors in steady-state or faulted conditions, under heavy or light loads, and so on. As mentioned in Chapter 3, simulation tools of various types are used widely for engineering purposes. However, these tools are all commercial products and the programs are proprietary. Thus it is often not possible for researchers to experiment with or add to the models and algorithms in these programs, although the committee notes that some commercial programs support user-defined models and the creation of external programs that extend their functionality. While many vendors provide licenses to their programs for university research and education usage at greatly reduced cost, these licenses are still an impediment for researchers from disciplines other than power systems engineering who wish to explore the problems. Because these programs are highly complex, incorporating many different models of the thousands of components on the grid, it is difficult for an experienced researcher, let alone a newcomer, to duplicate such tools with a reasonable amount of effort. Moreover, development of new algorithms, and their application to power systems, involves research issues that are unlikely to be addressed adequately by software vendors (see Chapter 6). At present, researchers have to try new algorithms or controls on simple test systems with synthetic data (for which custom programs can be written), or they are restricted to making limited add-ons to the existing commercial software. Moreover, the changes happening to the power grid and the IT infrastructure overlaying it are requiring simulation tools that are not just simple extensions of the existing tools. Rather, these changes call for fundamental adjustments to the underlying assumptions in the models and algorithms.

So, in addition to having a library of synthetic data, it would be very beneficial for the research community to have access to a library of simulation software. Even if one has to pay to use the software (because most are commercial products), it becomes very convenient to be able to compare and contrast existing simulation tools. Of course, it would be even more efficient if some of these software packages were open source and researchers could then modify, add, and test their own algorithms. Just starting such a library will encourage researchers and the industry to add their own open-source software to the collection.

Recommendation 10: The Department of Energy and the National Science Foundation should sponsor the development of new open-source software for the next-generation electric grid research community.

RECOMMENDATION FOR INCREASED R&D COORDINATION

DOE has an ongoing effort to coordinate the power grid research that it funds at the national laboratories. In fact, it has tasked a consortium of its national laboratories with mapping out a multiyear research plan for the power grid. A prime objective is the coordination of all the disparate but grid-related research projects being conducted at the national laboratories today. Such coordination would allow, say, the specification of compatible software interfaces—for example, those utilizing standard database structures—to be incorporated into analytical tools made for different purposes, which would improve processing times. The national labs have an invaluable resource in their multicore parallel processors, and their direct involvement in this effort is much desired.

Recommendation 11: In view of the importance of its efforts to coordinate power grid research at the national laboratories, the Department of Energy should broaden this coordination to include academic and industry researchers.

Recommendations 9 and 10, about creating libraries of synthetic data and providing access to software tools, should be well aligned. In addition, since such libraries are fully digital, they do not have to be in one physical location. Creation of interconnected virtual libraries with strict interoperability standards is an option for software tools. Of course, this does not do away with the need for physical locations for testing of actual hardware, but this report is mainly concerned with the research on and development of the analytical and computational tools needed for the planning and operation of the grid.

Implementing Recommendation 11 on the coordination of research efforts across the national laboratories and well beyond will greatly lower the hurdles facing experts from fields other than power engineering who wish to join the research effort. This would allow the multidisciplinary teams needed to solve these complex problems to be formed with greater ease.

RECOMMENDATION FOR A NATIONAL CENTER

As has been noted throughout this report, the power grid is changing ever faster, not only to adopt new technologies but to adapt to the changing climate. The changing climate is driving the world to reduce its use of carbon-based fuels to slow the warming trend, at the same time that work is done to strengthen the grid against extreme weather events. Thus all power engineering research roadmaps as adopted by DOE, NSF, the Electric Power Research Institute, and other research agencies call for new analytics for the planning and operation of the fast evolving power grid. Analytics usually refers to the suite of computer-based tools that is used by power engineers for designing the transmission and distribution systems and developing real-time monitoring and control systems for them.

The first digital analytic tools appeared in the late 1960s and early 1970s. The first power flow and transient stability algorithms, the first optimization algorithms for power flow and scheduling generation (unit commitment), the first Supervisory Control and Data System, and the first energy management system with state estimation and contingency analysis, were all developed in this short period in a great flurry of creative energy. In the next four decades these tools improved with the faster evolving computer hardware and software, but by and large, the methodology and algorithms did not change very much. Overall, the improvements since the 1970s in power engineering analytical tools have been incremental rather than transformational.

The main reason for the burst of creativity during the early period of power engineering analytics was the entry into the power engineering community of many new people who brought previously unknown mathematical techniques to bear on the problems. Tinney and Walker (1967) introduced sparse matrix techniques that solved power flows for large systems; Schweppe and Wildes (1970) brought in least-squares state estimation to solve the power network equations in real time; and many others brought in new optimization methods, new numerical solutions for grid dynamic behavior, and so on. The committee believes that the changes taking place in the grid today cannot be handled by incremental improvements in grid analytics. To be able to make the transformational changes needed, a research environment needs to be created that attracts mathematicians and experts in computation into power engineering research.

Today's organization structure for conducting power engineering research is being significantly built up to handle the renewed emphasis on power grid analytics. Many of the DOE national laboratories are playing a role in tackling this problem. Presumably this will be further extended by bringing academic institutions and industrial research groups into the fold. NSF continues to support power engineering research in general and two Engineering Research Centers focused on this topic. The DOE Office of Advanced Scientific Computing Research supports some relevant research in applied mathematics and related power engineering research. Although these efforts are very much needed and these organizations are getting bigger to handle the larger volume of research, they are not fundamentally changing the research environment.

The committee believes that for transformational research to take place, the research environment must proactively attract new mathematicians and computation experts into this research. Moreover, it would not be enough to just add a few mathematicians to the existing research teams at the national laboratories and the power engineering research groups at universities. The problems faced by the power grid will require sustained innovation in grid analytics. This calls for a more permanent organization that specifically nurtures talented researchers who may be new to the power engineering community but who would commit themselves to becoming knowledgeable about the research topic. The committee sees the need for a research center to foster such engagement.

Such a center might include a physical facility and staff to support the management of software and data sets (which themselves may be distributed elsewhere), along with virtual or in-person research collaborations among engineers, mathematicians, and other scientists throughout the country. The facility could be housed at a university, a national laboratory, or elsewhere. It would bring together experts from industry, national laboratories, and academia. The center would not necessarily need its own facilities for data storage or for testing the new algorithms and tools; in fact, the high-performance computer capabilities of DOE national laboratories or of cloud-based computing could be utilized.

Recommendation 12: The Department of Energy should establish a National Electric Power Systems Research Center to address fundamental research challenges associated with analysis for the future electric system. The center would act as an interface between the power industry, government, and universities in developing new computational and mathematical solutions for data and modeling challenges and in sharing valuable data.

The recommended center would include on its staff mathematicians, computational scientists and engineers, and power engineers. Their research would focus on the mathematical foundations described earlier in this report, such as fast power-flow programs, optimization programs for planning and operations, and new stochastic tools and methods. Since issues of social science might be of importance for management of the future grid, social scientists might be included in this multidisciplinary team, although social issues have been mentioned only in passing in this report. Some of the social science research needed was presented at the committee's 2015 workshop[2] by Miriam Goldberg. Her remarks are summarized in NASEM (2015). For scientists in these disciplines, one of the obvious benefits in working with the center would be access to otherwise unavailable data, both synthetic and real. Computer models developed at this center using synthetic data could also be tested later on in a secure facility using CIP data.

The center should foster a cross-disciplinary, multi-institutional approach to the analytical problems of the next-generation grid on a scale not normally possible in individual institutions. It should preserve and enhance the natural synergism between research and education while encouraging industry participation in its activities. The center would foster the education of future engineers in both industry and government as well as future faculty members. Through its industry connections, the proposed center could help researchers understand the behavior of the current U.S. power system, contribute to its orderly evolution, and enhance and extend the capabilities of the university and broader industry communities.

REFERENCES

NASEM (National Academies of Sciences, Engineering, and Medicine). 2015. *Mathematical Sciences Research Challenges for the Next-Generation Electric Grid: Summary of a Workshop*. The National Academies Press, Washington, D.C.

Schweppe, F.C., and J. Wildes. 1970. Power system static-state estimation, Part I, II, and III. *IEEE Transactions Power App. Systems* PAS-89(1):120-135.

Tinney, W.F., and J.W. Walker. 1967. Direct solutions of sparse network equations by optimally ordered triangular factorization. *Proceedings of the IEEE - PIEEE* 55(11):1801-1809.

[2] The agenda for the February 11-12, 2015, workshop is reprinted in Appendix A.

Appendixes

A

Workshop Agenda

DAY 1, FEBRUARY 11, 2015

8:00 a.m. **Welcome, Introductions, and Overview**

Opening Remarks and Meeting Overview
Robert J. Thomas, Cornell University, Workshop Planning Committee, Chair
Welcome and Study Objectives
John Guckenheimer, Cornell University, Study Committee, Co-Chair
Thomas Overbye, University of Illinois Champagne-Urbana, Study Committee, Co-Chair

8:15 **Keynote: Setting the Stage**

Low-Cost Pathways to Grid Integration of Renewable Energy: Skating to Where the Puck Is Going to Be
Steven Chu, Stanford University

9:15 **Data and Data Analytics**

Session Co-Chairs:
Cynthia Rudin, Massachusetts Institute of Technology
Marija Ilic, Carnegie Mellon University

Prosumer-centric Power Industry Transformation
David Sun, Alstom
How to Combine Observational Data Sources with First Principles of Physics to Build Stable and Transportable Models for Power System Design and Control
Louis Wehenkel, University of Liège, Belgium
Grid-Scale Data Fusion: Obstacles and Opportunities
Matthew Gardner, Dominion Virginia Power

12:00 p.m. **Breakout Session**

2:00 **Optimization and Control Methods for a Robust and Resilient Power Grid**

Session Chair:
Jeff Dagle, Pacific Northwest National Laboratory

Duration-Differentiated Electric Service for Integrating Renewable Power
Pravin Varaiya, University of California, Berkeley
Demand-Side Flexibility for Reliable Ancillary Services in a Smart Grid: Eliminating Risk to Consumers and the Grid
Sean Meyn, University of Florida
Advances in Mixed-Integer Programming and the Impact on Managing Electrical Power Grids
Robert Bixby, Gurobi

4:45 **Breakout Session**

DAY 2, FEBRUARY 12, 2015

8:30 a.m. **Uncertainty Quantification and Validation**

Session Chair:
Juan Meza, University of California, Merced

How Well Can We Measure What Didn't Happen and Predict What Won't?
Miriam Goldberg, DNV GL
Mathematical Models in Power Markets
Alexander Eydeland, Morgan Stanley

10:30 **Breakout Session**

11:45 **Wrap-up Session**

Session Chair:
Robert Thomas, Cornell University

Presentations from the audience

12:30 p.m. **Final Remarks by the Organizers**

1:00 **Workshop Adjourns**

B

Committee Biographies

JOHN GUCKENHEIMER, *Co-Chair,* holds the A.R. Bullis Chair of Mathematics in the Department of Mathematics at Cornell University. Earlier in his career (1973-1985), he was at the University of California, Santa Cruz. He was a Guggenheim fellow in 1984 and was elected president of SIAM in 1996. Dr. Guckenheimer received a B.A. from Harvard University in 1966 and a Ph.D. from the University of California, Berkeley, in 1970. His book *Nonlinear Oscillations, Dynamical Systems and Bifurcation of Vector Fields* (with Philip Holmes) is an extensively cited work on dynamical systems that was awarded the 2013 Steele Prize for Mathematical Exposition by the American Mathematical Society. He has made contributions in several other disciplines, ranging from neuroscience to fluid dynamics to numerical analysis of dynamical systems with multiple time scales.

THOMAS J. OVERBYE, *Co-Chair,* a member of the National Academy of Engineering (NAE), is the Fox Family Professor of Electrical and Computer Engineering at the University of Illinois, Urbana-Champaign, where he has taught since 1991. He received B.S., M.S., and Ph.D. degrees in electrical engineering from the University of Wisconsin, Madison. His current research interests include electric power system analysis, visualization, dynamics, cybersecurity, and modeling of power system geomagnetic disturbances. Professor Overbye is the original developer of the PowerWorld Simulator, a cofounder of PowerWorld Corporation, and an author of *Power System Analysis and Design*. He was the recipient of the IEEE/PES Walter Fee Outstanding Young Engineer Award in 1993 and its Outstanding Power Engineering Educator Award in 2011; he also participated in the 2003 DOE/NERC blackout investigation.

DANIEL BIENSTOCK is a professor in Columbia University's Industrial Engineering and Operations Research Department, where he has been since 1989. He received his Ph.D. in operations research from the Massachusetts Institute of Technology (MIT). His research focuses on optimization and high-performance computing, with a second focus on the use of computational mathematics in the analysis and control of power grids, especially the study of vulnerabilities and of cascading blackouts. Prior to joining Columbia University, Dr. Bienstock was in the combinatorics and optimization research group at Bellcore. He received the 2013 INFORMS fellow award, a Presidential Young Investigator award, and an IBM Faculty award, and he gave a plenary address at the 2005 Optimization Conference of the Society for Industrial and Applied Mathematics (SIAM) and a semiplenary address at the 2006 International Symposium on Mathematical Programming.

ANJAN BOSE, a member of the NAE, has over 45 years of experience in industry and academia as an engineer, educator, and administrator. He is now a Regents professor in the School of Electrical Engineering and Computer Science at Washington State University and holds the endowed Distinguished Professorship in Power Engineering. From 1998 until 2005 he served as dean of that University's College of Engineering and Architecture. From 1993 to 1998, he was the director of the School of Electrical Engineering and Computer Science. A fellow of the IEEE, he was the recipient of the Outstanding Power Engineering Educator Award (1994), the Third Millennium Medal (2000), and the Herman Halperin Electric Transmission and Distribution Award (2006. He has been recognized as a distinguished alumnus of the Indian Institute of Technology, Kharagpur (2005), and the College of Engineering at Iowa State University (1993). In 2012 and 2013, Professor Bose served as a senior advisor at DOE, where he led the Grid Tech team, which identified DOE priorities in the context of the next-generation grid.

W. TERRY BOSTON, a member of the NAE, retired recently after 6 years as CEO of PJM Interconnection, the largest power grid in North America and the largest electricity market in the world. Mr. Boston is president of the Association of Edison Illuminating Companies and past president of GO 15, the association of the world's largest power grid operators. He also served as a U.S. vice president of the International Council of Large Electric Systems and is a past chair of the North American Transmission Forum. He also was one of the eight industry experts selected to direct the North American Electric Reliability Corporation (NERC) investigation of the August 2003 Northeast/Midwest blackout. In 2011, Mr. Boston was honored with the Leadership in Power award from the IEEE Power and Energy Society. He also was chosen by *Intelligent Utilities Magazine* as one of the Top 11 Industry Movers and Shakers and led PJM to win Platts Global Energy awards in Industry Leadership 2010 and Excellence in Electricity in 2012. Mr. Boston received a B.S. in engineering from the Tennessee Technological University and an M.S. in engineering administration from the University of Tennessee.

JEFFERY DAGLE has worked at the Pacific Northwest National Laboratory, operated by Battelle for DOE, since 1989 and currently manages several projects in transmission reliability and security, including the North American SynchroPhasor Initiative (NASPI) and cybersecurity reviews for the DOE Smart Grid Investment Grants and Smart Grid Demonstration Projects. He is a senior member of the IEEE and a member of the International Society of Automation and the National Society of Professional Engineers. He received the 2001 Tri-City Engineer of the Year award from the Washington Society of Professional Engineers; led the data requests and management task for the U.S.-Canada Power System Outage Task Force investigation of the August 14, 2003, blackout; supported the DOE Infrastructure Security and Energy Restoration Division with on-site assessments in New Orleans following Hurricane Katrina in the fall of 2005; and is the recipient of two patents, a Federal Laboratory Consortium (FLC) Award in 2007, and an R&D 100 Award in 2008 for grid friendly appliance controller technology. Mr. Dagle was a member of a National Infrastructure Advisory Council study group formed in 2010 to establish critical infrastructure resilience goals. He received B.S. and M.S. degrees in electrical engineering from Washington State University in 1989 and 1994, respectively.

MARIJA ILIC is a professor at Carnegie Mellon University (CMU) with joint appointments in electrical and computer engineering and engineering and public policy. She is director of the Electric Energy Systems Group at CMU and also serves as Honorary Chaired Professor for Control of Future Electricity Network Operations at the Delft University of Technology (the Netherlands). An IEEE fellow, Professor Ilic has over 30 years of experience in teaching and research in the area of electrical power system modeling and control. Her main interest is the systems aspects of operations, planning, and economics of the electric power industry. She has coauthored several books in her field of interest, most recently coediting *Engineering IT-Enabled Sustainable Electricity Services: The Case of Low-Cost Green Azores Island*.

CHRISTOPHER K.R.T. JONES is the Bill Guthridge Distinguished Professor of Mathematics at the University of North Carolina (UNC) at Chapel Hill. He received his Ph.D. in mathematics from the University of Wisconsin, Madison and, before joining UNC, was a professor of applied mathematics at Brown University for 13 years. The main thrust of Dr. Jones's research is the use of dynamical systems as a tool for solving problems that originate in applications—in particular, the use of dynamical systems methods in the study of nonlinear wave motion in neuroscience and optics, ocean dynamics, and, more recently, climate. He is director of the Mathematics and Climate Research Network, a broadly based NSF-funded effort to engage the mathematical sciences community in climate science and to define the problems that will form an emerging area of "climate mathematics."

FRANK KELLY is a foreign member of NAE and a professor of the mathematics of systems at the University of Cambridge. His main research interests are in random processes, networks, and optimization, and he is especially interested in applications to the design and control of networks and to the understanding of self-regulation in large-scale systems. From 2003 to 2006 he served as chief scientific adviser to the U.K. Department for Transport. He was chair of the U.K. Council for the Mathematical Sciences from 2010 to 2013 and is a member of RAND Europe's Council of Advisors as well. Professor Kelly is a fellow of the Royal Society.

YANNIS KEVREKIDIS is the Pomeroy and Betty Perry Smith Professor in Engineering at Princeton University. He received a B.S in chemical engineering from the National Technical University in Athens and M.S. and Ph.D. degrees in mathematics from the University of Minnesota. After a year at the Center for Nonlinear Studies in Los Alamos in 1985 and 1986, he moved to Princeton, New Jersey, where he teaches courses in chemical engineering and applied and computational mathematics. His research interests center on the dynamics of physical and chemical processes, types of instabilities, pattern formation, and the ways to study and understand such phenomena computationally. Recently he has also developed an interest in multiscale computations. Professor Kevrekidis has been a Packard fellow, a Presidential Young Investigator, a Guggenheim fellow, and a Ulam Scholar at Los Alamos. He has won the Colburn and Wilhelm awards of the American Institute of Chemical Engineers, SIAM's Crawford prize, and a Humboldt Prize.

RALPH D. MASIELLO, a member of the NAE, is an Industry Advisor at Quanta Technology. A recognized leader in next-generation electric grid systems, his focus in recent years has included energy storage applications and system integration, renewables integration in markets and operations, and development and integration of distributed energy resources. Dr. Masiello is a life fellow of the IEEE and has served as chairman of the IEEE section on power system engineering, chairman of the IEEE section on power industry computing applications, on the editorial board of the *IEEE Proceedings*, and on the advisory board for *IEEE Spectrum* magazine. He is the recipient of the 2009 IEEE Power Engineering Concordia award for Power System Engineering. Dr. Masiello received his B.S., M.S., and Ph.D. degrees from MIT in electrical engineering.

JUAN C. MEZA is dean of the School of Natural Sciences at the University of California, Merced. Before that Dr. Meza was for many years the head and senior scientist of the High Performance Computing Research Department at E.O. Lawrence Berkeley National Laboratory, where he oversaw work in computational science and mathematics, computer science and future technologies, scientific data management, visualization, numerical algorithms, and application development. His current research interests include nonlinear optimization, with an emphasis on methods for parallel computing. He has also worked on various scientific and engineering applications, including scalable methods for nanoscience, power grid reliability, molecular conformation problems, optimal design of chemical vapor deposition furnaces, and semiconductor device modeling. Before joining Lawrence Berkeley, Dr. Meza was a distinguished member of the technical staff at Sandia National Laboratories and served as manager of the Computational Sciences and Mathematics Research department. Dr. Meza has been named by *Hispanic Business Magazine* as one of the Top 100 Influentials in the area of science. A fellow of the American Association for the Advancement of Science (AAAS), Dean Meza was the 2008 recipient of the Blackwell-Tapia Prize and the SACNAS Distinguished Scientist Award. He was also a member of the team that won the Association for Computing Machinery's (ACM) 2008 Gordon Bell Award for Algorithm Innovation. Dr. Meza has served on numerous external committees, including DOE's Advanced Scientific Computing Research Advisory Committee, NSF's Mathematical and Physical Sciences Advisory Committee and its Advisory Committee for Cyberinfrastructure, and the SIAM Board of Trustees.

CYNTHIA RUDIN is an associate professor of statistics at MIT and directs the Prediction Analysis Lab. Before joining MIT, Dr. Rudin held positions at the Center for Computational Learning Systems at Columbia University and at New York University (NYU). She holds an undergraduate degree from the University at Buffalo, where she received the College of Arts and Sciences Outstanding Senior Award in Sciences and Mathematics. She received a Ph.D. in applied and computational mathematics from Princeton University. She is the recipient of the 2013 INFORMS Innovative Applications in Analytics Award, an NSF CAREER award, and was named as one of the "Top 40 Under 40" by Poets and Quants in 2015.

ROBERT J. THOMAS is professor emeritus of electrical and computer engineering at Cornell University, where he began teaching in 1973. The author of over 100 technical papers and two book chapters, Professor Thomas has been a member of the Energy Policy Committee of the Institute of Electrical and Electronics Engineers (IEEE) since 1991 and was that committee's chair in 1997 and 1998. He has also been a member of the IEEE Technology Policy Council, has served as the IEEE-USA vice president for technology policy, and has been a member of several university, government, and industry advisory boards. He has published in the areas of transient control and voltage collapse problems as well as on the technical, economic, and institutional impacts of restructuring. He is the founding director of the 13-university member, National Science Foundation (NSF)-sponsored Power Systems Engineering Research Center (PSERC). He was a member of the DOE Secretary's Power Outage Study Team and is a founding member of the Coalition for Electric Reliability Solutions (CERTS). Professor Thomas was on assignment to the DOE in 2003 as a senior advisor to the director of the Office of Electric Transmission and Distribution and a member of the DOE team investigating the August 14, 2003, blackout, and he has also spent time with the DOE Office of Electric Energy Systems and at the NSF as the first program director for the Power Systems Program. He contributed to the 2007 National Interest Electric Transmission Corridor study and was an advisor to three DOE assistant secretaries for electricity delivery and energy reliability from 2002 to 2011. He served as one of 30 inaugural members of the DOE Secretary's Electricity Advisory Committee from 2008 until 2010. Professor Thomas has received five teaching awards and the IEEE Centennial and Millennium medals, and is an IEEE life fellow.

MARGARET H. WRIGHT, a member of both the National Academy of Sciences and the NAE, is Silver Professor of Computer Science and Mathematics in the Courant Institute of Mathematical Sciences at NYU. She received B.S. (mathematics), M.S., and Ph.D. (computer science) degrees from Stanford University. Her research interests include optimization, scientific computing, and real-world applications. Before joining NYU, she worked at Bell Laboratories (Lucent Technologies). Professor Wright has served as president of SIAM and on numerous advisory committees for the DOE and several mathematical sciences institutes.

C

Acronyms

ACE	area control error
ACOPF	alternating current, optimal power flow
AEMO	Australian Energy Market Operator
AGC	automatic generation control
AMI	advanced metering infrastructure
ARPA-E	Advanced Research Products Agency-Energy
ARRA	(U.S.) American Recovery and Reinvestment Act (2009)
ASO	Ancillary Services Operator
AVR	automatic voltage regulator
BMSA	Board on Mathematical Sciences and Their Applications
BPA	Bonneville Power Administration
CA	contingency analysis
CAISO	California Independent System Operator
CCVT	coupled capacitive voltage transformer
CEII	Critical Energy Infrastructure Information
CIP	critical infrastructure protection
ConEdison	Consolidated Edison
CT	current transformer
Δ-connected	delta-connected
DA	data assimilation
DAE	differential algebraic equation
DER	Direct Energy Resources
DFIG	doubly fed inductions generator
DOE	Department of Energy
dPIN	double Pareto-log normal (distribution)
DS	dynamic stability

DSO	Distribution System Operator
EMP	electromagnetic pulse
EMS	Energy Management System
EP	equilibrium point
E-Pro	electric grid protection
ERCOT	Electric Reliability Council of Texas
EV	electric vehicle
FACTS	Flexible Alternating Current Transmission Systems
FERC	Federal Energy Regulatory Commission
FPC	Federal Power Commission
FTR	financial transmission right
GENI	Global Energy Network Institute
GIC	geomagnetically induced current
GMD	geomagnetic disturbance
GPS	Global Positioning System
GW	gigawatt
HEMP	high-altitude electromagnetic pulse
HILF	high impact, low frequency
HVDC	high-voltage direct current
Hz	hertz
iBA	intelligent Balancing Authority
ICAP	installed capacity
ICT	information and communication technology
IEEE	Institute of Electrical and Electronics Engineers
ISO	Independent System Operator
ISO-NE	ISO New England
IT SCED	intermediate-term security-constrained economic dispatch
KCL	Kirchhoff's current law
kV	kilovolt
LFC	load frequency control
LMP	locational marginal cost
LOLP	loss of load probability
LP	linear programming
LPC	locational pricing calculator
LTC	load tap-changing
MCRN	Mathematics Climate Research Network
MINLP	mixed-integer nonlinear programming
MIP	mixed-integer programming
MISO	Midcontinent ISO
MPC	model predictive control
Mvar	megavar
MVA	megavolt-ampere

MW/MWh	megawatt
NAPSIC	North American Power Systems Interconnection Committee
NDA	nondisclosure agreement
NEM	National Electricity Market
NEPSRC	National Electric Power Systems Research Center
NERC	North American Electric Reliability Corporation
NR	Newton-Raphson (method)
NRC	National Research Council
NSF	National Science Foundation
NYISO	New York Independent Systems Operator
ODE	ordinary differential equation
OPF	optimal power flow
P	power
PC	polynomial chaos (expansion)
PID	proportional integral derivative
PJM	Pennsylvania-New Jersey-Maryland Interconnection, LLC
PMU	phasor measurement unit
PSERC	Power Systems Engineering Research Center
PSS	power system stabilizer
PT	potential transformer
PTDF	power transfer distribution factor
PU	per unit
PV	photovoltaic
Q	Reactive Power
QC	Quadratic Convex
QCQP	Quadratically Constrained Quadratic Program
RAS	Remedial Action Scheme
RTCA	real-time contingency analysis
RTO	Regional Transmission Organizations
RT SCED	real-time security-constrained economic dispatch
RTU	remote terminal unit
SCADA	Supervisory Control and Data Acquisition
SciDAC	Scientific Discovery through Advanced Computing
SCOPF	security-constrained optimal power flow
SCUC	security-constrained unit commitment
SDE	stochastic differential equation
SDP	semidefinite programming
SE	state estimation
SFT	simultaneous feasibility test
SPD	scheduling, pricing, and dispatch
SPP	Southwest Power Pool
SPS	Special Protection System
SQP	sequential quadratic programming
SSCI	subsynchronous control instability

SVC	static var compensator
TCP/IP	Transmission Control Protocol/Internet Protocol
TCSC	thyristor-controlled series capacitor
TEB	transient excitation boost
TS	transient stability
TVA	Tennessee Valley Authority
TWh	terawatt
V	voltage magnitude
VSD	variable speed drive
WAMS	wide-area measurement system
WECC	Western Interconnection
Y-connected	wye-connected