Power Systems Background and an Introduction to my Ph.D. Research

Keith Moffat

This document is intended to be an introduction to power systems and (some of) my research for a reader who is not already familiar with power systems. Thus, it provides some background on power systems and analogies for the electric grid. It is intended to not assume familiarity with power systems jargon and be accessible to a scientifically-inclined audience. An appropriate treatment would require much more than a few pages but perhaps this document will provide some insight.

Power systems, or "the grid," use electricity[1] to transfer energy from generators to loads so the loads can work. "Generators" are anything that injects power into the grid.[2] "Loads" are anything that intentionally takes power from the grid. Examples of loads include a light that is on, a toaster that is toasting, or an electric vehicle that is charging. "Work" has a technical meaning in physics ("the application of force along a displacement") that is different fom its meanings in everyday language. The technical distinctions of "work" are not important for power systems analysis, however; what is important is that loads take power from the grid. "Power" is the rate at which energy is transferred

---

[1] Water and mountains are an imperfect-but-helpful analogy for electricity. The grid can be thought of as a hypothetical mountain that is used to transmit work from one location to another location. Generators use an external source of power to push water up the mountain to a higher elevation. The water then flows laterally across the mountain through a river to the location where the load is. When it gets to the location where the load is, the water flows down the mountain through a wheel/turbine, powering the load. The power delivered to the load is given by multiplying the river's elevation on the mountain where the load is by the amount of water flowing. In this explanation, the riverbed is analogous to the wires on the grid, the elevation of the riverbed is analogous to electric voltage, and the river is analogous to electric current.

This analogy is intended to help develop intuition for why electric power is the product of voltage and current. In the water-and-mountain analogy, if the load wanted more power, the generator could push more water up the mountain to flow through the riverbed to the load (i.e., send more current at the same voltage), or it could push the same amount of water to a riverbed at a higher elevation (i.e., send the same amount of current at a higher voltage).

[2] Traditionally, generators have been large machines that produce electric power for distribution by spinning a magnet inside coils of wires that are attached to a grid. The magnet spins because it is attached to turbine blades that are pushed by steam or water. Most of the electric power on the grid today is provided by power plants that burn natural gas or coal to produce steam that pushes turbine blades. But there are other types of generators that do not rotate magnets inside of coils of wire, such as solar panels and batteries. Solar panels and batteries are also "generators."

from generators to loads. Thus, power is a measurement that is made at each instant, whereas energy is the cumulative sum of power over a period of time.[3] All the power that loads take from the grid comes from generators. The total generation and total load[4] must be approximately equal at all moments in time. If this balance is disrupted for a significant amount of time, the grid will have a "blackout" and not be able to transfer power from the generators to the loads until the grid is restarted.

Power systems deliver power from generators to loads using power lines.[5] The electric power flowing into, out of, or through any location on the grid is given by multiplying the voltage at the location by the current flowing into, out of, or through the location at each moment in time.[6] Thus, the amount of electric power flowing into, out of, or through a given location on the grid can be increased in two ways: by increasing the voltage at the location, or by increasing the current flowing into, out of, or through the location.

The grid's physical infrastructure constrains the amount of current that can flow on a line as well as the maximum voltage at which it is safe to operate the grid.[7] If there is too much current flowing in the lines, the lines become too hot. These are called "thermal," or "flow" constraints. If the line voltage is too high, then the voltage difference between the power lines can damage equipment. If the grid voltage is too low, then the voltage on the grid can suddenly collapse to zero, causing a blackout. Thus, electric grids are operated in such a way that both the line flows and the voltages stay within upper and lower bounds.

Electric power cannot be told where to flow—once power is on the grid, it flows from the generators to the loads according to the laws of physics. If the loads on the grid are not managed by grid operators, as is traditionally assumed, the power/current flows on the grid can only be altered by changing which generators supply power to the grid at each moment in time.

The Internet and relatively inexpensive, powerful computers have made it

---

[3] To illustrate the difference between power and energy, consider the example of two generators with one attached to a solar panel and the other attached to a battery. The two generators have the same capacity, but the generator attached to the solar panel only generates electricity when the sun is shining on the solar panel, whereas the battery generates power whenever it is told to (if the battery has not been fully discharged). Thus, if the sun shines for the first five minutes of an hour and then is blocked by rain clouds, the solar panel will generate full electricity for only the first five minutes then turn off for the rest of the hour. The battery, however, may generate electricity for the full hour. The solar panel and battery may thus be able to provide the same amount of electric power for the grid, but the solar panel will provide less energy over the course of the hour than the battery.

[4] In addition to the loads, some power is "lost" as heat on the power lines when power flows from the generators to the loads. We define the "total load" on the grid to be the sum of all the loads attached to the network and all of the losses on the network.

[5] Power lines are large metal (aluminum) wires that can transfer large amounts of electricity.

[6] See footnote 1.

[7] Continuing the water analogy to understand grid constraints, the flow/thermal constraints are the maximum amount of current that can flow through a riverbed/wire before bad things happen (i.e., overheating the wire). Voltage constraints are defined as the range of elevations at which it is safe for a riverbed to be. For reasons specific to electric grids, there are both upper and lower limits on voltage.

possible to manage loads in real time, upending the traditional assumption that only generators can be controlled. It is now possible to avoid thermal and voltage constraints by reducing the power drawn by certain loads on the system in real time. For example, it is now possible to tell electric vehicles in each neighborhood that there is too much load in that neighborhood at a given moment and the power lines have hit their thermal limit. The utility company that manages the grid can offer to pay the electric vehicle owners if they reduce their vehicle charging rate until the power/current flows on the lines are below their limits.

In the future there might be thousands of loads on the grid which can be managed to avoid constraint violations. Thus, it will be necessary to have tools that automatically make decisions for each load based on high-level objectives that are set by the utility company/grid operator. It is worth noting that the same tools that automatically make decisions for loads can also be used to automatically make decisions for generators distributed across the grid, such as rooftop solar panels or Tesla Powerwall batteries. By thinking of the distributed generators as "negative loads," the tools used to automatically manage loads can also be used to control the power output of distributed generators without any modification to the tools.

It is also helpful to understand the difference between transmission networks and distribution networks. Transmission networks span large areas, such as the western United States and parts of Canada and Mexico, with thick wires that transfer a lot of power at "very high" voltage (e.g., hundreds of kilovolts (kV)). Distribution networks attach to the transmission grid via substations, which step the voltage down to just "high" voltage (e.g., tens of kV). A single transmission system can provide power to thousands of distribution networks.

While the same physics apply to transmission and distribution networks, the two types of systems are operated in different ways. Typically, transmission systems have been actively managed to avoid constraint violations—that is, care is taken when grid operators decide how much power to generate at each generator so that the power/current flows on the transmission lines do not violate thermal constraints.

Distribution networks, on the other hand, are not actively managed. The power flows from the substation to the loads without any coordination of the loads and generators on the distribution network. Thermal and voltage constraints are avoided on distribution networks by 1) building the distribution infrastructure for the most challenging possible circumstance, and 2) limiting new connections onto the grid.[8] Requiring distribution networks to support the most challenging possible circumstance without any coordination of applicable

---

[8] For example, if an individual household or business wanted to build a powerful electric vehicle charging station, the utility that owns and operates the distribution network would limit if/where the household/business could build the charging station based on where the grid has capacity for large new loads. The capacity for new loads is based on the single moment of the year that is most challenging for the thermal and voltage constraints. For the rest of the year, the line flows and network voltages are below thermal and voltage constraints.

Similarly, not every house is allowed to install new rooftop solar panels that will inject power loads and generators results in unnecessary interconnection restrictions and the need to install larger, more expensive equipment than would be necessary if the loads and generators were coordinated.

My research is focused on the emerging practice of managing/controlling the power consumption or generation of distributed energy resources to avoid violating distribution network thermal and voltage constraints. "Managing distributed energy resources" could mean managing how much power an electric vehicle charging station uses, how much power a Tesla Powerwall battery injects back into the grid, or reducing the power generated by a rooftop solar panel. Managing distributed energy resources to avoid network constraints is not something that has been done historically,[9] but it may become commonplace in the coming years as the amount of electric power the grid delivers from generators to loads increases. Managing distributed energy resources to avoid network constraints will result in more affordable, more accessible, and more efficient electricity service.

To transition to renewable energy we will need to connect many electric vehicles, batteries, rooftop solar panels, and other distributed energy resources to the grid. Managing these resources will be challenging because the decisions for many resources must be made at once, and they must be made in a "fair" manner.[10] Furthermore, distribution network operators do not generally know exactly how a change in the power consumption or generation from one resource will affect a given voltage or line flow that is in danger of violating a constraint. Consider an example distribution network which supplies power to both an electric vehicle charging station in town and homes outside of town. If the voltage is too low at the homes outside of town, it would be helpful for distribution network operators to know exactly how sensitive the voltage at those homes is to changes in the power consumption at the electric vehicle charging station. At present, distribution network operators do not know the sensitivities that describe how a change in the power injections will change the voltages and line flows on a distribution network.[11] This lack of knowledge prevents distribution network operators from actively managing the resources on distribution

---

back into the grid. The decision to install new solar panels that will inject power back into the grid must be cleared with the utility, which is allowed to say that a household/business is not allowed to inject power back into the grid.

[9] There are a number of utility programs and companies that control distributed energy resources, but these programs/companies focus on reducing load or increasing generation so that the total load and the total generation on the grid are equal, not avoiding distribution network constraints. "Virtual Power Plant" is a term commonly used to describe controlling distributed energy resources so that the total generation matches the total load.

[10] The question of what is "fair" allocation of grid resources is critically important, but outside of the scope of my research, which focuses on the engineering challenges. My research seeks to provide tools to operate the grid more efficiently, and to make the trade-offs of policy decisions transparent.

[11] Distribution network operators have an estimate or intuition for the sensitivities of the voltages and line flows, but do not generally have an accurate mathematical expression for them because there are challenging nuances in accurately determining the sensitivities. Also, traditional distribution network operation has not required exact sensitivities. Accurate mathematical expressions for the sensitivities are helpful for automatically managing the power injections of distributed energy resources to avoid distribution network constraint violations.

networks.

Broadly speaking, my research works on methods for addressing the challenges described in this introduction.

One example of my research is the introduction of new equations that describe how changes in the power injections of the resources attached to the grid will affect the voltages and line flows on the grid. These equations that describe the "sensitivity" of the grid could be derived from a model of the grid that includes how the grid wires are connected, what the wires are made of, and how long the wires are.

Alternatively, when grid operators do not have an accurate model of the grid, the sensitivity of the grid—how changes in the power injections of the resources attached to the grid will affect the voltages and line flows on the grid—can be "learned" from measurements without taking the time to learn the physical traits of the grid. These "data-driven" methods are easier to implement and, perhaps, more accurate. The data-driven methods have limitations, though, based on how accurate the sensors are. My work seeks to understand these limitations so that someone implementing the data-driven method has guarantees of how well the method will perform.

Once a model for the grid (sensitivity) is known, the final step is managing the distributed energy resources—e.g. managing electric vehicle charging, managing the charging/discharging of a Tesla Powerwall battery, or managing the power output of a solar panel or wind turbine. My research has introduced/developed two different tools for automatically managing distributed energy resources to avoid thermal and voltage constraint violations on distribution networks.

The two tools, "Voltage Phasor Control" and "Feedback Optimization," make automatic decisions based on an objective function. An objective function is a mathematical expression that codifies the grid-operation goals.[12] Thus, the objective function can be crafted so that the Voltage Phasor Control tool and/or the Feedback Optimization tool control distributed energy resources in a fair manner. The tools themselves do not specify an objective function—that is left up to the policy-makers, grid operators, and people who use the grid.

The research problems above were the focus of my Ph.D. Since moving to ETH Zurich for my postdoc I have been focused on grid stability, which is a different than the challenges described in this document. In a sentence, this research can be described as "how to prevent power outages and other bad things from happening when we replace large rotating machine generators[13] with solar, wind, and battery generators." This work requires a different set of background knowledge which will be the focus of Power Systems Background Chapter 2, coming soon.

---

[12] Examples of goals that can be expressed as objective functions include "deliver as much electric power as possible," or "give all participants the same amount of electric power."

[13] Gas, nuclear, coal, and hydro plants fall into the "large rotating machine" category.