Optimal Control of Cascading Power Grid Failures

Reconfiguring Power Systems to Minimize Cascading Failures: Models and Algorithms

Daniel Bienstock
Columbia University

Abstract

In a cascading failure of a power transmission system, an initial event that disables a possibly small subset of the grid conspires with the laws of physics to set off a sequence of additional outages that, in the worst case, accelerates until a large subset of the network is inoperative, resulting in a significant loss of served power.

The mechanics of the process can be summarized as follows: each time a component of the system becomes outaged, a new set of power flows takes hold in the remaining network, following the laws of physics and automatic control actions. Should the new flows, for example, exceed the rating of a given line, then that line will likely become outaged in the near future. In an adverse scenario this gives rise to a vicious cycle which constitutes the cascade.

The complexity of the process is due to interaction between the physics of power flows, as described by an appropriate system of equations, and discrete (i.e., discontinuous) changes in the topology of the network as components are removed from operation. This interaction gives rise to non-monotonic behavior, both from the perspective of control actions (more of a good thing is not necessarily better) and with respect to the severity of the initial event. Non-monotonicity, from a mathematical perspective, is equivalent to non-convexity, and thus, complexity. Additionally, there is poorly understood *noise* inherent in the operation of the system; for example a thermally stressed line is more likely to 'trip' due to many factors that are difficult to model in a precise manner.

We consider algorithms that shed load (demand) as a function of noisy observations taken in real time, with the goal of arresting the cascade with a minimum of demand lost. In principle, one could think of algorithms in this mold that are available *prior to* the cascade; i.e. algorithms independent of the nature of the cascade. However, there is an exponential amount of variability in the structure of the initial event that sets-off the cascade, and in our simulations it is clear that the nature of the cascade can depend strongly on the initial event. Instead, we focus on control algorithms that are computed soon after the onset of the cascade, and which therefore benefit from the knowledge available at that point.

In one version of our work, we assume an initially *slow-moving* cascade so that at the start of the process there is sufficient time (e.g., minutes) to compute an appropriate control algorithm; once computed, the control will be applied as the cascade unfolds. In devising a load-shedding schedule to respond to a potential cascade, one must decide when, where and by how much demand is to be shed. Our methods can be viewed as a data-driven approach for computing such actions – they is data-driven because it relies on the knowledge of the initial event. At the same time, the efficiency of the algorithms makes it potentially feasible to investigate common patterns of desirable control actions that could be directly applied in the event of a contingency.

In another version of our work, we assume a potentially *fast moving* cascade where in addition the initiating event is not observed. Here, offer an "online-learning" variant of classical exponential backoff algorithms.

We present computational results with multi-core algorithms, applied to a large-scale grid, the Eastern Interconnect, with over 15000 buses and 23000 lines.