Bilevel Nonlinear Programming for Vulnerability Analysis of Power Grids

Reconfiguring Power Systems to Minimize Cascading Failures: Models and Algorithms

Dan Bienstock, Sean Harnett, Taedong Kim, Stephen Wright
Columbia University and University of Wisconsin-Madison

Abstract

Electrical power grids in steady state can be modeled by nonlinear equations that describe the physical laws governing power flows (the "AC equations"), together with prescribed demands at some nodes and power generation at others. When a fault occurs on a grid (in the form of a severed power line, for example), the power flows in the grid change, possibly in a detrimental way that causes other lines to fail. This process can result in a cascade of failures and may possibly lead to a widespread blackout, in which demand cannot be met at many of the nodes.

One way to assess the vulnerability of a grid to faults — including faults caused by a malicious attacker — is to identify those lines whose removal would cause the maximum disruption. A natural model for vulnerability analysis is *bilevel nonlinear optimization*, in which the lower-level problem captures the disruption to the grid arising from a given set of faults, while the variables in the upper-level problem model the faults. The AC power-flow equations appear as constraints in the lower-level program, possibly together with limits on power flows. Essentially, the bilevel optimization problem is to find the set of faults that maximizes the disruption to the grid, subject to some "budget" constraint. The budget is used to map the sensitivity of the grid to different levels of damage, or, in the case of a malicious attacker, the power of the attacker.

Previous approaches to this problem have replaced the lower-level problem with a set of stationarity conditions, thus reformulating the bilevel problem as a fairly standard nonlinear optimization problem. Such formulations have the disadvantage of allowing spurious solutions of the AC power flow equations into the model, giving a misleading picture of the effects of a prescribed set of faults on the grid. We propose instead to retain the bilevel structure, solving the lower-level problem with "black-box" techniques that more reliably rule out spurious solutions.

We describe two techniques for quantifying the disruption. One measures the deviation of the voltage magnitudes at the network nodes from their reference value of 1. (Greater deviation indicates a less stable network.) The second technique measures the total amount of load that needs to be shed at demand nodes in order to restore feasibility to the grid. This quantity can be obtained by solving a "least ℓ_1" solution of the AC power flow equations.

In the upper-level problem, the variables that quantify the faults can be simply binary variables that indicates whether a line is severed. Instead of these discrete variables, we use continuous variables that represent changes to impedances of each transmission line. The upper-level problem maximizes our chosen metric of disruption subject to total increase in impedance staying within some bound. The impedance increases serve as proxies for the binary variables; those lines on which increased impedance leads to severe disruption are the lines that make the grid most vulnerable.

Many algorithmic techniques from optimization are needed to solve the resulting formulations. We describe the use of Frank-Wolfe methods, gradient projection, partial gradient methods, and finite differences in solving the upper-level problem. For the lower-level problem, we describe the use of trust-region methods for the least-ℓ_1 formulation, along with enhanced Newton methods, approximate Newton methods, homotopy techniques, and the semidefinite programming formulations introduced by Lavaei and Low. Computational results of grids with over 15000 nodes and 23000 lines will be presented.