The $N-k$ Problem in Power Grids: New Models, Formulations and Computation[1]

Daniel Bienstock and Abhinav Verma
Columbia University
New York

May 2008

Abstract

Given a power grid modeled by a network together with equations describing the power flows, power generation and consumption, and the laws of physics, the so-called $N-k$ problem asks whether there exists a set of k or fewer arcs whose removal will cause the system to fail. We present theoretical results and computation involving two optimization algorithms for this problem.

1 Introduction

Recent large-scale power grid failures have highlighted the need for effective computational tools for analyzing vulnerabilities of electrical transmission networks. Blackouts are extremely rare, but their consequences can be severe. Recent blackouts had, as their root cause, an exogenous damaging event (such as a storm) which developed into a system collapse even though the initial quantity of disabled power lines was small.

As a result, a problem that has gathered increasing importance is what might be termed the *vulnerability evaluation* problem: given a power grid, is there a small set of power lines whose removal will lead to system failure? Here, "smallness" is parameterized by an integer k, and indeed experts have called for small values of k (such as $k = 3$ or 4) in the analysis. Additionally, an explicit goal in the formulation of the problem is that the analysis should be agnostic: we are interested in rooting out small, "hidden" vulnerabilities of a complex system which is otherwise quite robust; as much as possible the search for such vulnerabilities should be devoid of assumptions regarding their structure.

This problem is not new, and researchers have used a variety of names for it: network *interdiction*, network *inhibition* and so on, although the "N - k problem" terminology is common in the industry (where "N" is the number of arcs). We will provide a more complete review of the (rather extensive) literature later on; the core central theme is that the $N-k$ problem is very highly intractable, even for small values of k – the pure enumeration approach is simply impractical. In addition to the combinatorial explosion, another significant difficulty is the need to model the laws of physics governing power flows in a sufficiently accurate and yet computationally tractable manner: power flows are much more complex than "flows" in traditional applications.

A critique that has been leveled against optimization-based approaches to the $N-k$ problem is that they tend to focus on large values of k, say $k = 8$. When k is large the problem tends to become easier, but on the other hand the argument can be made that the cardinality of the attack is unrealistically large. At the other end of the spectrum lies the case $k = 1$, which can be addressed by enumeration but may not yield useful information. The middle range, $2 \leq k \leq 5$, is both relevant and difficult, and is our primary focus.

In this paper we take an approach based on strict optimization. We present results using two models. The first (Section 2.1) is a new linear mixed-integer programming formulation that explicitly models a "game" between a fictional attacker seeking to disable the network, and a controller who tries to prevent a collapse by selecting which generators to operate and adjusting generator outputs and demand levels. As far as we can tell, the problem we consider here is

[1] Partially funded by NSF award 0521741

more general than has been previously studied in the literature; nevertheless our approach yields practicable solution times for larger instances than previously studied.

The second model (Section 3) is given by a new, continuous nonlinear programming formulation whose goal is to capture, in a compact way, the interaction between the underlying physics and the network structure. While both formulations provide substantial savings over the pure enumerational approach, the second formulation appears particularly effective and scalable; enabling us to handle in an optimization framework models an order of magnitude larger than those we have seen in the literature.

1.0.1 Previous work on vulnerability problems

There is a large amount of prior work on optimization methods applied to blackout-related problems. [20] includes a fairly comprehensive survey of recent work.

Typically work has focused on identifying a small set of arcs whose removal (to model complete failure) will result in a network unable to deliver a minimum amount of demand. A problem of this type can be solved using mixed-integer programming techniques techniques, see [2], [21], [3]. We will review this work in more detail below (Section 2.0.6). Generally speaking, the mixed-integer programs to be solved can prove quite challenging.

A different line of research on vulnerability problems focuses on attacks with certain structural properties, see [6], [20]. An example of this approach is used in [20]. Here, as an approximation to the $N - k$ problem with AC power flows, the authors formulate a linear mixed-integer program to solve the following combinatorial problem: remove a minimum number of arcs, such that in the resulting network there is a partition of the nodes into two sets, N_1 and N_2, such that

$$D(N_1) \, + \, G(N_2) \, + \, cap(N_1, N_2) \, \leq \, Q^{min}.$$

Here $D(N_1)$ is the total demand in N_1, $G(N_2)$ is the total generation capacity in N_2, $cap(N_1, N_2)$ is the total capacity in the (non-removed) arcs between N_1 and N_2, and Q^{min} is a minimum amount of demand that needs to be satisfied. The quantity in the left-hand side in the above expression is an upper-bound on the total amount of demand that can be satisfied – the upper-bound can be strict because under power flaw laws it may not be attained.

Thus this is an approximate model that could underestimate the effect of an attack (i.e. the algorithm may produce attacks larger than strictly necessary). On the other hand, methods of this type bring to bear powerful mathematical tools, and thus can handle larger problems than algorithms that rely on generic mixed-integer programming techniques. Our method in Section 3 can also be viewed as an example of this approach.

Finally, we mention that the most sophisticated models for the behavior of a grid under stress attempt to capture the multistage nature of blackouts, and are thus more comprehensive than the static models considered above and in this paper. See, for example, [9]-[12], and [5].

1.0.2 Power Flows

Here we provide a brief introduction to the so-called *linearized*, or *DC* power flow model. For the purposes of our problem, a grid is represented by a directed network $\mathcal{N}$, where:

- Each node corresponds to a "generator" (i.e., a supply node), or to a "load" (i.e., a demand node), or to a node that neither generates nor consumes power. We denote by $\mathcal{G}$ the set of generator nodes.

- If node i corresponds to a generator, then there are values $0 \leq P_i^{min} \leq P_i^{max}$. If the generator is operated, then its output must be in the range $[P_i^{min}, P_i^{max}]$; if the generator is not operated, then its output is zero. In general, we expect $P_i^{min} > 0$.

- If node i corresponds to a demand, then there is a value D_i^{nom} (the "nominal" demand value at node i). We will denote the set of demands by $\mathcal{D}$.

- The arcs of $\mathcal{N}$ represent power lines. For each arc (i,j), we are given a parameter $x_{ij} > 0$ (the resistance) and a parameter u_{ij} (the capacity).

Given a set $\mathcal{C}$ of operating generators, a *power flow* is a solution to the system of constraints given next. In this system, for each arc (i,j), we use a variable f_{ij} to represent the (power) flow on (i,j) – possibly $f_{ij} < 0$, in which case power is effectively flowing from j to i. In addition, for each node i we will have a variable θ_i (the "phase angle" at i). Finally, if i is a generator node, then we will have a variable P_i, while if i represents a demand node, we will have a variable D_i. The constraints are:

$$\sum_{(i,j)\in\delta^+(i)} f_{ij} - \sum_{(j,i)\in\delta^-(i)} f_{ji} = \begin{cases} P_i & i \in \mathcal{C} \\ -D_i & i \in \mathcal{D} \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

$$\theta_i - \theta_j + x_{ij} f_{ij} = 0 \quad \forall (i,j) \tag{2}$$

$$|f_{ij}| \quad \leq \quad u_{ij}, \quad \forall (i,j) \tag{3}$$

$$P_i^{min} \leq P_i \leq P_i^{max} \quad \forall i \in \mathcal{C} \tag{4}$$

$$0 \leq D_j \leq D_j^{nom} \quad \forall j \in \mathcal{D} \tag{5}$$

We will denote this system by $\boldsymbol{P(\mathcal{N}, \mathcal{C})}$. Constraint (1) models flow conservation, while (4) and (5) describe generator and demand node bounds. Optionally, one may impose additional constraints, in particular bounds on the θ_i or on the quantities $|\theta_i - \theta_j|$ (over the arcs (i,j)).

1.0.3 Basic results

A useful property satisfied by the linearized model is summarized by the following result which is not difficult to prove.

Lemma 1.1 *Let $\mathcal{C}$ be given, and suppose $\mathcal{N}$ is connected. Then for any choice of nonnegative values P_i (for $i \in \mathcal{C}$) and D_i (for $i \in \mathcal{D}$) such that*

$$\sum_{i\in\mathcal{C}} P_i \quad = \quad \sum_{i\in\mathcal{D}} D_i, \tag{6}$$

system (1)-(2) has a unique solution in the f_{ij}; the solution is also unique in the $\theta_i - \theta_j$ (over the arcs (i,j)).

Remark 1.2 *We stress that Lemma 1.1 concerns the subsystem of $\boldsymbol{P(\mathcal{N}, \mathcal{C})}$ consisting of (1) and (2). In particular, the "capacities" u_{ij} play no role in the determination of solutions.*

When the network is not connected Lemma 1.1 can be extended by requiring that (6) hold for each component.

Definition 1.3 *Let (f, θ, P, D) be feasible a solution to $P(\mathcal{N}, \mathcal{C})$. The* **throughput** *of (f, θ, P, D) is defined as*

$$\frac{\sum_{i\in\mathcal{D}} D_i}{\sum_{i\in\mathcal{D}} D_i^{nom}}. \tag{7}$$

The throughput of $\mathcal{N}$ is the maximum throughput of any feasible solution to $P(\mathcal{N}, \mathcal{C})$.

1.0.4 DC and AC power flows

Constraint (2) is reminiscent of Ohm's equation – in a direct current (DC) network (2) precisely represents Ohm's equation. In the case of an AC network (the relevant case when dealing with power grids) (2) only *approximates* a complex system of nonlinear equations (see [1]). The issue of whether to use an the more exact nonlinear formulation, or the approximate DC formulation, is rather thorny. On the one hand, the linearized formulation certainly is an approximation only. On the other hand, the AC formulation can prove intractable or otherwise inappropriate (e.g. the formulation may have multiple solutions), and, we stress, is *itself* in any case an approximate model of the underlying physics.

For these reasons, studies that require multiple power flow computations tend to rely on the linearized formulation. This will be the approach we take in this paper, though some of our techniques extend directly to the AC model and this will remain a topic for future research. An approach such as ours can therefore be criticized because it relies on an ostensibly approximate model; on the other hand we are able to focus more explicitly on the basic combinatorial complexity that underlies the $N - k$ problem. In contrast, an approach that uses the AC model would have a better representation of the physics, but at the cost of not being able to tackle the combinatorial complexity quite as effectively, for the simple reason that the theory and computational machinery for linear programming are far more mature, effective *and* scalable than those for nonlinear, nonconvex optimization. In summary, both approaches present limitations and benefits. In this paper, our bias is toward explicitly handling the combinatorial nature of the problem.

A final point that we would like to stress is that whether we use the AC or DC power flow model, the resulting problems have a far more complex structure than (say) traditional single- or multi-commodity flow models because of side side-constraints such as (2). Constraints of this type give rise to counter-intuitive behavior reminiscent of Braess's Paradox [8].

2 The "N - k" problem

Let $\mathcal{N}$ be a network with n nodes and m arcs representing a power grid. We denote the set of arcs by E and the set of nodes by V. A fictional *attacker* wants to remove a small number of arcs from $\mathcal{N}$ in order to maximize damage. Somewhat informally (and, as it turns out, incompletely), the goal of the attacker is that in the resulting network all feasible flows should have low throughput. At the same time, a *controller* is operating the network; the controller responds to an attack by appropriately choosing the set $\mathcal{C}$ of operating generators, their output levels, and the demands D_i, so as to feasibly obtain high throughput.

Thus, the attacker seeks to defeat *all possible courses of action* by the controller, in other words, we are modeling the problem as a Stackelberg game between the attacker and the controller, where the attacker moves first. To cast this in a precise way we will use the following definition. We let $0 \leq T^{min} \leq 1$ be a given value.

Definition 2.1 *Given a network $\mathcal{N}$,*

- *An* **attack** *$\mathcal{A}$ is a set of arcs removed by the attacker.*
- *Given an attack $\mathcal{A}$, the* **surviving network** *$\mathcal{N} - \mathcal{A}$ is the subnetwork of $\mathcal{N}$ consisting of the arcs not removed by the attacker.*
- *A* **configuration** *is a set $\mathcal{C}$ of generators.*
- *We say that an attack $\mathcal{A}$* **defeats** *a configuration $\mathcal{C}$, if either (a) the maximum throughput of any feasible solution to $\boldsymbol{P(\mathcal{N} - \mathcal{A}, \mathcal{C})}$ is strictly less than T^{min}, or (b) no feasible solution to $\boldsymbol{P(\mathcal{N} - \mathcal{A}, \mathcal{C})}$ exists. Otherwise we say that $\mathcal{C}$ defeats $\mathcal{A}$.*
- *We say that an attack is* **successful***, if it defeats* **every** *configuration.*

- *The* **min-cardinality attack problem** *consists in finding a successful attack $\mathcal{A}$ with $|\mathcal{A}|$ minimum.*

Our primary focus will be on the min-cardinality attack problem. Before proceeding further we would like to comment on our model, specifically on the parameter T^{min}. In a practical use of the model, one would wish to experiment with different values for T^{min} – for the simple reason that an attack $\mathcal{A}$ which is not successful for a given choice for T^{min} could well be successful for a slightly larger value; e.g. no attack or cardinality 3 or less exists that reduces demand by 31%, and yet there exists an attack of cardinality 3 that reduces satisfied demand by 30%. In other words, the minimum cardinality of a successful attack could vary substantially as a function of T^{min}.

Given this fact, *it might appear* that a better approach to the power grid vulnerability problem would be to leave out the parameter T^{min} entirely, and instead redefine the problem to that of finding a set of k or fewer arcs to remove, so that the resulting network has minimum throughput (here, k is given). We will refer to this as the *budget-k min-throughput problem.* However, there are reasons why this latter problem is less attractive than the min-cardinality problem.

(a) Clearly, in a sense, the min-cardinality and min-throughput problems are duals of each other. A modeler considering the min-throughput problem would want to run that model multiple times, because given k, the value of the budget-k min-throughput problem could be much smaller than the value of the budget-$(k+1)$ min-throughput problem. For example, it could be the case that using a budget of $k = 2$, the attacker can reduce throughput by no more than 5%; but nevertheless with a budget of $k = 3$, throughput can be reduced by e.g. 50%. In other words, even if a network is “resilient” against attacks of size ≤ 2, it might nevertheless prove very vulnerable to attacks of size 3. For this reason, and given that the models of grids, power flows, etc., are rather approximate, a practitioner would want to test various values of k – this issue is obviously related to what percentage of demand loss would be considered tolerable, in other words, the parameter T^{min}.

(b) From an operational perspective it should be straightforward to identify reasonable values for the quantity T^{min}; whereas the value k is more obscure and bound to models of how much power the adversary can wield.

(c) Because of a subtlety that arises from having positive quantities P_i^{min}, explained next, it turns out that the min-throughput problem is significantly more complex and is difficult to even formulate in a compact manner.

We will now expand on (c). One would expect that when a configuration $\mathcal{C}$ is defeated by an attack $\mathcal{A}$, it is because only small throughput solutions are feasible in $\mathcal{N} - \mathcal{A}$. However, because the lower bounds P_i^{min} are in general strictly possible, it may also be the case that *no feasible solution to* $\boldsymbol{P(\mathcal{N} - \mathcal{A}, \mathcal{C})}$ *exists.*

Example 2.2 *Consider the following example on a network $\mathcal{N}$ with three nodes, where*

1. *Nodes* 1 *and* 2 *represent generators; $P_1^{min} = 2$, $P_1^{max} = 4$, $P_2^{min} = 0$, and $P_2^{max} = 4$,*
2. *Node* 3 *is a demand node with $D_3^{nom} = 6$. Furthermore, $T^{min} = 1/2$.*
3. *There are three arcs; arc $(1,2)$ with $x_{12} = 1$ and $u_{12} = 3$, arc $(2,3)$ with $x_{23} = 1$ and $u_{23} = 5$, and arc $(1,3)$ with $x_{13} = 1$ and $u_{13} = 1$.*

When the network is not attacked, the following solution is feasible: $P_1 = P_2 = 3$, $D_3 = 6$, $f_{12} = 0$, $f_{13} = f_{23} = 3$, $\theta_1 = \theta_2 = 0$, $\theta_3 = -3$. This solution has throughput 100%. *On the other hand, consider the attack $\mathcal{A}$ consisting of the single arc $(1,3)$, and suppose we choose the configuration $\mathcal{C} = \{1,2\}$ (i.e. we operate both generators). Since $P_1^{min} > u_{12}$,* $\boldsymbol{P(\mathcal{N} - \mathcal{A}, \mathcal{C})}$ *has no feasible solution, and $\mathcal{A}$ defeats $\mathcal{C}$ (in spite of the fact that we can still meet* 100% *of the demand).*

Likewise, $\mathcal{A}$ defeats the configuration where we only operate generator 1. *Thus, $\mathcal{A}$ is successful if and only if it also defeats the configuration where we only operate generator* 2, *which it does not since in that configuration we can feasibly send up to four units of flow on $(2,3)$ and $T^{min} = 1/2 < 4/6$.*

As the example highlights, it is important to understand how an attack $\mathcal{A}$ can defeat a particular configuration $\mathcal{C}$. It turns out that there are *three* different ways for this to happen.

(i) Consider a partition of the nodes of $\mathcal{N}$ into two classes, N^1 and N^2. Write

$$D^k \;=\; \sum_{i \in \mathcal{D} \cap N^k} D_i^{nom}, \quad k = 1, 2, \quad \text{and} \tag{8}$$

$$P^k \;=\; \sum_{i \in \mathcal{C} \cap N^k} P_i^{max}, \quad k = 1, 2, \tag{9}$$

e.g. the total (nominal) demand in N_1 and N_2, and the maximum power generation in N_1 and N_2, respectively. The following condition, should it hold, would guarantee that $\mathcal{A}$ defeats $\mathcal{C}$:

$$T^{min} \sum_{j \in \mathcal{D}} D_j^{nom} - \min\{D^1, P^1\} - \min\{D^2, P^2\} \;>\; \sum_{(i,j) \notin \mathcal{A} \,:\, i \in N^1,\, j \in N^2} u_{ij} \;+ \sum_{(i,j) \notin \mathcal{A} \,:\, i \in N^2,\, j \in N^j} u_{ij}. \tag{10}$$

To understand this condition, note that for $k = 1, 2$, $\min\{D^k, P^k\}$ is the maximum demand within N^k that could possibly be met using power flows that do not leave N^k. Consequently the left-hand side of (10) is a lower bound on the amount of flow that must travel between N^1 and N^2, whereas the right-hand side of (10) is the total capacity of arcs between N^1 and N^2 under attack $\mathcal{A}$. In other words, condition (10) amounts to a mismatch between demand and supply. A special case of (10) is that where in $\mathcal{N} - \mathcal{A}$ there are no arcs between N^1 and N^2, i.e. the right-hand side of (10) is zero.

(ii) Consider a partition of the nodes of $\mathcal{N}$ into two classes, N^1 and N^2, such that in $\mathcal{N} - \mathcal{A}$ there are *no* arcs between N^1 and N^2. Then attack $\mathcal{A}$ defeats $\mathcal{C}$ if

$$\sum_{i \mathcal{D} \cap \in N^1} D_i^{nom} \;<\; \sum_{i \in \mathcal{C} \cap N^1} P_i^{min}, \tag{11}$$

i.e., the minimum power output within N^1 exceeds the maximum demand within N^1. Note that (ii) may apply even if (i) does not.

(iii) Even if (i) and (ii) do not hold, it may still be the case that the system (1)-(5) does not admit a feasible solution. To put it differently, suppose that for every choice of demand values $0 \le D_i \le D_i^{nom}$ (for $i \in \mathcal{D}$) and supply values $P_i^{min} \le P_i \le P_i^{max}$ (for $i \in \mathcal{C}$) such that $\sum_{i \in \mathcal{C}} P_i = \sum_{i \in \mathcal{D}} D_i$ the unique solution to system (1)-(2) in network $\mathcal{N} - \mathcal{A}$ (as per Lemma 1.1) does *not* satisfy the "capacity" inequalities $|f_{ij}| \le u_{ij}$ for all arcs $(i, j) \in \mathcal{N} - \mathcal{A}$. Then $\mathcal{A}$ defeats $\mathcal{C}$. This is the most subtle case of all – it involves the interplay of flow conservation and Ohm's law.

Note that in particular in case (ii), the defeat condition is unrelated to throughput. Nevertheless, should case (ii) arise, it is clear that the attack has succeeded (against configuration $\mathcal{C}$) – this makes the min-throughput problem difficult to model; our formulation for the min-cardinality problem, given in Section 2.1, does capture the three defeat criteria above.

From a practical perspective, it is important to handle models where the values P_i^{min} are positive. It is also important to model *standby* generators that are turned on when needed, and to model the turning off of generators that are unable to dispose of their minimum power output, post-attack. All these features arise in practice. Example 2.2 above shows that models where generators cannot be turned off can exhibit unreasonable behavior. Of course, the ability to select the operating generators comes at a cost, in that in order to certify that an attack is successful we need to evaluate, at least implicitly, a possibly exponential number of control possibilities.

As far as we can tell, most (or all) prior work in the literature **does** require that the controller must always use the configuration $\bar{\mathcal{G}}$ consisting of all generators. As the example shows, however, if the quantities P_i^{min} are positive there may be attacks $\mathcal{A}$ such that $\boldsymbol{P(\mathcal{N} - \mathcal{A}, \bar{\mathcal{G}})}$ is infeasible. Because of this fact, algorithms that rely on direct application of Benders' decomposition or bilevel programming are problematic, and *invalid* formulations can be found in the literature.

Our approach works with general $P^{min} \geq 0$ quantities; thus, it also applies to the case where we always have $P_i^{min} = 0$. In this case our formulation is simple enough that a commercial integer program solver can directly handle instances larger than previously reported in the literature.

2.0.5 Non-monotonicity

Consider the example in Figure 1, where we assume $T^{min} = 0.3$. Notice that there are two parallel copies of arcs $(2,4)$ and $(3,5)$, each with capacity 10 and impedance 1. It is easy to see that the network with no attack is feasible: we operate generator 1 and not operate generators 2 and 3, and send 3 units of flow along the paths $1-6-2-4$ and $1-6-3-5$ (the flow on e.g. the two parallel $(2,4)$ arcs is evenly split).

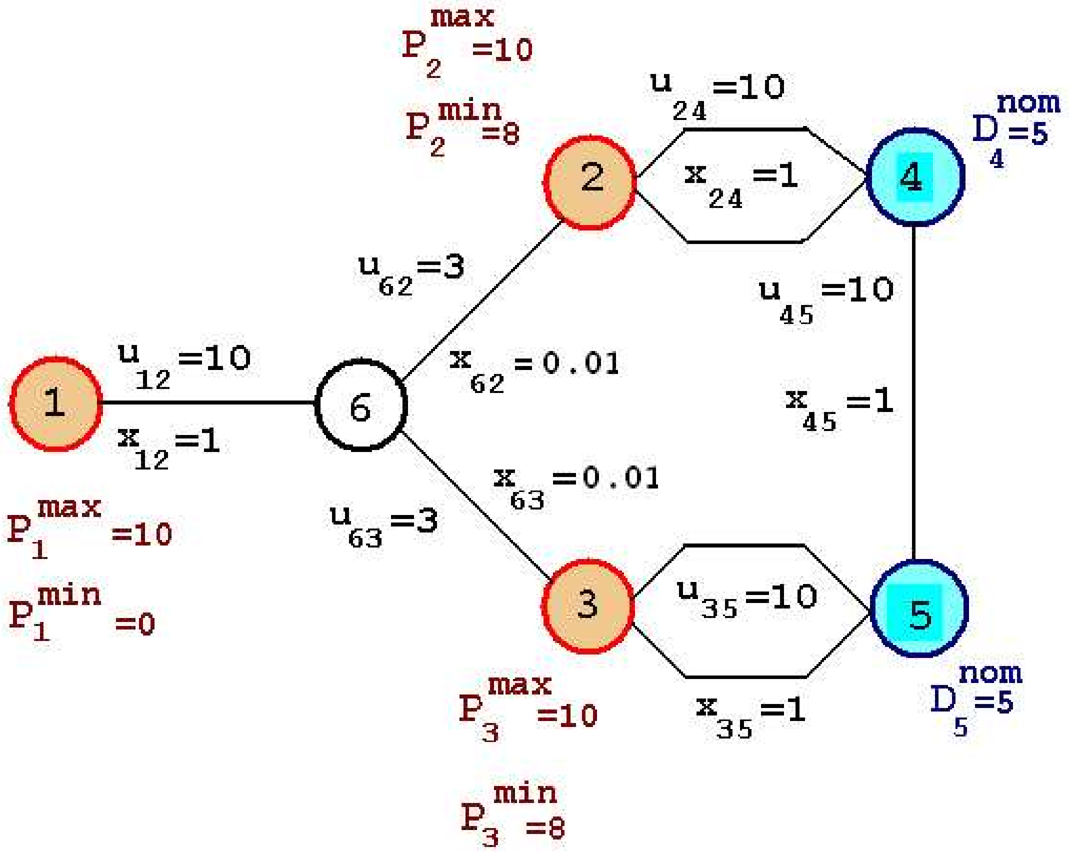

Figure 1: Non-monotone example.

On the other hand, consider the attack consisting of arc $(1,6)$ – we will show this attack is successful. To see this, note that under this attack, the controller cannot operate both generators 2 and 3, since their combined minimum output exceeds the total demand. Suppose, for example, that only generator 3 is operated, and assume by contradiction that a feasible solution exists – then this solution must route *at most* 3 units of flow along $3-6-2-4$, and (since $P_3^{min} = 8$) *at least* 5 units of flow on $(3,5)$ (both copies altogether). In such a case, the voltage drop from 3 to 5 is at least 2.5, whereas the voltage drop from 3 to 4 is at most 1.56. In other words, $\theta_4 - \theta_5 \geq 0.94$, and so we will have $f_{45} \geq 0.94$ – thus, the net inflow at node 5 is at least 5.94. Hence the attack is indeed successful.

However, there is *no* successful attack consisting of arc $(1,6)$ and another arc. To see this, note that if one of $(2,6)$, $(3,6)$ or $(4,5)$ are also removed then the controller can *just* operate one of the two generators 2, 3 and meet eight units of demand. Suppose that (say) one of the two copies of $(3,5)$ is removed (again, in addition to $(1,6)$). Then the controller operates generator 2, sending 2.5 units of flow on each of the two parallel $(2,4)$ arcs; thus $\theta_2 - \theta_4 = 2.5$. The controller also routes 3 units of flow along $2-6-3-5$, and therefore $\theta_2 - \theta_5 = 3.06$. Consequently $\theta_4 - \theta_5 = .56$, and $f_{45} = .56$, resulting in a feasible flow which satisfies 4.44 units of demand at 4 and 3.56 units of demand at 5.

In fact, it is straightforward to show that *no* successful attack of of cardinality 2 exists – hence we observe non-monotonicity.

By elaborating on the above, one can create examples with arbitrary patterns in the cardinality of successful attacks. One can also generate examples that exhibit non-monotone behavior in response to controller actions. In both cases, the non-monotonicity can be viewed as a manifestation of the so-called "Braess's Paradox" [8]. In the above example we can observe combinatorial subtleties that arise from the ability of the controller to choose which generators to operate, and from the lower bounds on output in operating generators. Nevertheless, it is clear that the critical core reason for the complexity is the interaction between voltages and flows, i.e. between "Ohm's law" (2) and flow conservation (1) – the combinatorial attributes of the problem exercise this interaction. Thus, we view it as crucial that an optimization model address the interaction in an explicit manner.

2.0.6 Brief review of previous work

The min-cardinality problem, as defined above, can be viewed as a bilevel program where both the master problem and the subproblem are mixed-integer programs – the master problem corresponds to the attacker (who chooses the arcs to remove) and the subproblem to the controller (who chooses the generators to operate). In general, such problems are extremely challenging. A recent general-purpose algorithm for such integer programs is given in [14].

Alternatively, each configuration of generators can be viewed as a "scenario". In this sense our problem resembles a stochastic program, although without a probability distribution. Recent work [17] considers a single commodity max-flow problem under attack by an interdictor with a limited attack budget; where an attacked arc is removed probabilistically, leading to a stochastic program (to minimize the expected max flow). A deterministic, multi-commodity version of the same problem is given in [18].

Previous work on the power grid vulnerability models proper has focused on cases where either the generator lower bounds P_i^{min} are all zero, or all generators must be operated (the single configuration case). Algorithms for these problems have either relied on heuristics, or on mixed-integer programming techniques, usually a direct use of Benders' decomposition or bilevel programming. [2] considers a version of the min-throughput problem with $P_i^{min} = 0$ for all generators i, and presents an algorithm using Benders' decomposition (also see references therein). They analyze the so-called IEEE One-Area and IEEE Two-Area test cases, with, respectively, 24 nodes and 38 arcs, and 48 nodes and 79 arcs. Also see [21].

[3] studies the IEEE One-Area test case, and allows $P_i^{min} > 0$, but does not allow generators to be turned off; the authors present a bilevel programming formulation which, unfortunately, is incorrect, due to reasons outlined above.

2.1 An algorithm for the min-cardinality problem

In this section we will describe an iterative algorithm for the min-cardinality attack problem. The algorithm iterates in Benders-like fashion, solving at each iteration two mixed-integer programs. Before describing the algorithm we need to introduce some notation and concepts.

Let $\mathcal{A}$ be a given attack. Suppose the controller attempts to defeat the attacker by choosing a certain configuration $\mathcal{C}$ of generators. Denote by $z^{\mathcal{A}}$ the indicator vector for $\mathcal{A}$, i.e. $z^{\mathcal{A}}_{ij} = 1$ iff $(i,j) \in \mathcal{A}$. Then the controller needs to solve the following linear program:

$$\boldsymbol{K_{\mathcal{C}}(\mathcal{A}):} \qquad t_{\mathcal{C}}(z^{\mathcal{A}}) \quad \doteq \quad \min\ t \tag{12}$$

Subject to:

$$\sum_{(i,j)\in\delta^+(i)} f_{ij} - \sum_{(j,i)\in\delta^-(i)} f_{ji} = \begin{cases} P_i & i \in \mathcal{G} \\ -D_i & i \in \mathcal{D} \\ 0 & \text{otherwise} \end{cases} \tag{13}$$

$$\theta_i - \theta_j + x_{ij} f_{ij} = 0 \quad \forall\ (i,j) \notin \mathcal{A} \tag{14}$$

$$u_{ij}\ t\ -\ |f_{ij}| \quad \geq \quad 0, \quad \forall\ (i,j) \notin \mathcal{A} \tag{15}$$

$$f_{ij} \quad = \quad 0, \quad \forall (i,j) \in \mathcal{A} \tag{16}$$

$$P_i^{min}\ \leq P_i\ \leq P_i^{max} \quad \forall i \in \mathcal{C} \tag{17}$$

$$P_i \quad = \quad 0, \quad \forall i \in \mathcal{G}\ -\ \mathcal{C} \tag{18}$$

$$\sum_{j\in\mathcal{D}} D_j\ \geq\ T^{min}\left(\sum_{j\in\mathcal{D}} D_j^{nom}\right), \tag{19}$$

$$0 \leq\ D_j\ \leq D_j^{nom} \quad \forall j \in \mathcal{D} \tag{20}$$

Remark 2.3 *Using the convention that the value of an infeasible linear program is infinite, $\mathcal{A}$ defeats $\mathcal{C}$ if and only if $t_{\mathcal{C}}(z^{\mathcal{A}}) > 1$.*

Thus, an attack $\mathcal{A}$ is *not* successful if and only if we can find $\mathcal{C} \subseteq \mathcal{G}$ with $t_{\mathcal{C}}(z^{\mathcal{A}}) \leq 1$; we test for this conditions by solving the problem:

$$\min_{\mathcal{C}\subseteq\mathcal{G}}\ t_{\mathcal{C}}(z^{\mathcal{A}}).$$

This is done by replacing, in the above formulation, equations (17), (18) with

$$P_i^{min} y_i\ \leq P_i\ \leq P_i^{max} y_i, \quad \forall i \in \mathcal{G}, \tag{21}$$

$$y_i\ = 0 \ \text{ or } \ 1, \quad \forall i \in \mathcal{G}. \tag{22}$$

Here, $y_i = 1$ if the controller operates generator i.

The min-cardinality attack problem can now be written as follows:

$$\min \sum_{(i,j)} z_{ij} \tag{23}$$

$$t_{\mathcal{C}}\ (z)\ >\ 1, \quad \forall\ \mathcal{C} \subseteq \mathcal{G}, \tag{24}$$

$$z_{ij}\ = 0 \text{ or } 1,\ \ \forall\ (i,j). \tag{25}$$

This formulation, of course, is impractical, because we do not have a compact way of representing any of the constraints (24), and there are an exponential number of them.

Putting these issues aside, we can outline an algorithm for the min-cardinality attack problem. Our algorithm will be iterative, and will maintain a "master attacker" mixed-integer program which will be a *relaxation* of (23)-(25) – i.e. it will have objective (23) but weaker constraints than (24). Initially, the master attacker MIP will include no variables other than the z variables, and no constraints other than (25). The algorithm proceeds as follows.

Basic algorithm for min-cardinality attack problem

```
Iterate:

1. Attacker: Solve master attacker MIP and let $z^*$ be its solution.

2. Controller: Search for a set $\mathcal{C}$ of generators such that $t_{\mathcal{C}}(z^*) \le 1$.

  (2.a) If no such set $\mathcal{C}$ exists, EXIT:
        $\sum_{ij} z^*_{ij}$ is the minimum cardinality of a successful attack.

  (2.b) Otherwise, suppose such a set $\mathcal{C}$ is found.
        Add to the master attacker MIP a system of valid inequalities that cuts off $z^*$.
        Go to 1.
```

As discussed above, the search in Step 2 can be implemented by solving a mixed integer program. Since in 2.b we add valid inequalities to the master, then inductively we always have a relaxation of (23)-(25) and thus the value of the master at any execution of step 1, i.e. the value $\sum_{ij} z^*_{ij}$, is a lower bound on the cardinality of any successful attack. Thus the exit condition in step 2.a is correct, since it proves that the attack implied by z^* is successful.

The implementation of Case 2.b, on the other hand, requires some care. Assuming we are in case 2.b, we have that $t_{\mathcal{C}}(z^*) \le 1$, and certainly the linear program $\boldsymbol{K_{\mathcal{C}}(\mathcal{A})}$ is feasible. The optimal dual solution would therefore (apparently) furnish a Benders cut that cuts off z^*. However this would be incorrect since the structure of constraints (12)-(20)) depends on z^* itself.

Instead, we need to proceed as in two-stage stochastic programming with recourse, where the z variables play the role as "first-stage" variables *and* also appear in the second-stage problem (the subproblem); solutions to the dual of the second-stage problem can then be used to generate cuts to add to the master problem. Toward this goal, we proceed as follows, given $\mathcal{C}$ and z^*:

B.1 Write the *dual* of $K_{\mathcal{C}}(\emptyset)$.

B.2 As is standard in interdiction-type problems (see [18], [17], [14], [2]), the dual is then "combinatorialized" by adding the z variables and additional constraints. For example, if β_{ij} indicates the dual of constraint (14), then we add, to the dual of $K_{\mathcal{C}}(\emptyset)$, inequalities of the form

$$\beta_{ij} - M^1_{ij} z_{ij} \le 0, \;\; -\beta_{ij} - M^1_{ij} z_{ij} \le 0,$$

for an appropriate constant $M^1_{ij} > 0$. We proceed similarly with constraint (15), obtaining the "combinatorial dual". This combinatorial dual is the functional equivalent of the second-stage problem in stochastic programming.

B.3 Fix the z_{ij} variables in the combinatorial dual to z^*; this yields a problem that is equivalent to $K_{\mathcal{C}}(z^*)$ and has the general structure

$$\begin{aligned} t_{\mathcal{C}}(z^*) \; = \qquad & \max \;\; c^T v \\ & Pv \; \le \; b \; + \; Qz^*. \end{aligned} \tag{26}$$

Here, the v are variables, P and Q are matrices, and b is a vector, of appropriate dimensions; and we have a maximization problem since the $K_{\mathcal{C}}()$ are minimization problems. We obtain a cut of the form

$$\bar{\alpha}^T (b + Qz) \ge 1 + \epsilon$$

where $\epsilon > 0$ is a small constant and $\bar{\alpha}$ is the vector of optimal dual variables to (26). Since by assumption $t_{\mathcal{C}}(z^*) \leq 1$ this inequality cuts off z^*.

Note the use of the tolerance ϵ. The use of this parameter gives *more* power to the controller, i.e. "borderline" attacks are not considered successful. In a strict sense, therefore, we are not solving the optimization problem to exact precision; nevertheless in practice we expect our relaxation to have negligible impact so long as ϵ is small. A deeper issue here is how to interpret truly borderline attacks that are successful according to our strict model (and which our use of ϵ disallows); we expect that in practice such attacks would be ambiguous and that the approximations incurred in modeling power flows, estimating demands levels, and so on, not to mention the numerical sensitivity of the integer and linear solvers being used, would have a far more significant impact on precision.

2.1.1 Discussion

In order to make the outline provided in B.1-B.3 into a formal algorithm, we need to specify the constants M^1_{ij}. As is well-known, the folklore of integer programming dictates that the M^1_{ij} should be chosen small to improve the quality of the linear programming relaxation of the master problem.

While this is certainly true, we have found that popular optimization packages show significant numerical instability when solving power flow *linear* programs. In fact, in our experience it is primarily this behavior that mandates that the M^1_{ij} should be kept as small as possible. In particular we do not want the M^1_{ij} to grow with network since this would lead to an nonscalable approach.

It turns out that our formulation $K_{\mathcal{C}}(\mathcal{A})$ is not ideal toward this goal. A particularly thorny issue is that the attack $\mathcal{A}$ may disconnect the network, and proving "reasonable" upper bounds on the dual variables to (for example) constraint (13), when the network is disconnected, does not seem possible. In the next section we describe a different formulation for the min-cardinality attack problem which is much better in this regard. Our eventual algorithm will apply steps B.1 - B.3 to this improved formulation, while the rest of our basic algorithmic methodology as described above will remain unchanged.

2.2 A better mixed-integer programming formulation

As before, let $\mathcal{A}$ be an attack and $\mathcal{C}$ a (given) configuration of generators. Let $y^{\mathcal{C}} \in R^{\mathcal{G}}$ be the indicator vector for $\mathcal{C}$, i.e. $y^{\mathcal{C}}_i = 1$ if $i \in \mathcal{C}$ and $y^{\mathcal{C}}_i = 0$ otherwise. Consider the following *linear* program:

$$\boldsymbol{K^*_{\mathcal{C}}(\mathcal{A})} : \qquad t^*_{\mathcal{C}}(z^{\mathcal{A}}) \quad \doteq \quad \min t \tag{27}$$

Subject to:

$$(\boldsymbol{\alpha^{\mathcal{C}}_{ij}}) \qquad \sum_{(i,j)\in\delta^+(i)} f_{ij} - \sum_{(j,i)\in\delta^-(i)} f_{ji} = \begin{cases} P_i & i \in \mathcal{G} \\ -D_i & i \in \mathcal{D} \\ 0 & \text{otherwise} \end{cases} \tag{28}$$

$$(\boldsymbol{\beta^{\mathcal{C}}_{ij}}) \qquad \theta_i - \theta_j + x_{ij} f_{ij} = 0 \quad \forall\ (i,j) \notin \mathcal{A} \tag{29}$$

$$(\boldsymbol{p^{\mathcal{C}}_{ij}, q^{\mathcal{C}}_{ij}}) \qquad u_{ij}\ t\ -\ |f_{ij}| \quad \geq \quad 0, \quad \forall\ (i,j) \notin \mathcal{A} \tag{30}$$

$$(\boldsymbol{\omega^{\mathcal{C}+}_{ij}, \omega^{\mathcal{C}-}_{ij}}) \qquad t\ -\ |f_{ij}| \quad \geq \quad 1, \quad \forall\ (i,j) \in \mathcal{A} \tag{31}$$

$$(\boldsymbol{\gamma^{\mathcal{C}+}_{i}, \gamma^{\mathcal{C}+}_{i}}) \qquad P^{min}_i y^{\mathcal{C}}_i \ \leq P_i\ \leq P^{max}_i y^{\mathcal{C}}_i \quad \forall i \in \mathcal{G} \tag{32}$$

$$(\boldsymbol{\mu^{\mathcal{C}}}) \qquad \sum_{j \in \mathcal{D}} D_j \ \geq \ T^{min} \left(\sum_{j \in \mathcal{D}} D_j^{nom} \right), \tag{33}$$

$$(\boldsymbol{\Delta_j^{\mathcal{C}}}) \qquad D_j \ \leq D_j^{nom} \quad \forall j \in \mathcal{D} \tag{34}$$

$$P \geq 0, \quad D \geq 0. \tag{35}$$

To the left of each constraint we have indicated the corresponding dual variable – (30) is really two constraints written as one, and the same with (31).

Note that we do not force $f_{ij} = 0$ for $(i,j) \in \mathcal{A}$. Moreover arcs $(i,j) \in \mathcal{A}$ are also exempted from constraint (29). Thus, the controller has significantly more power than in $K_{\mathcal{C}}(\mathcal{A})$. However, because of constraint (31), we have $t^*_{\mathcal{C}}(z^{\mathcal{A}}) > 1$ as soon as any of the arcs in $\mathcal{A}$ actually carries flow. We can summarize these remarks as follows:

Remark 2.4 *$\mathcal{A}$ defeats $\mathcal{C}$ if and only if $t^*_{\mathcal{C}}(z^{\mathcal{A}}) > 1$.*

Note that the above formulation depends on $\mathcal{C}$ only through constraint (32). Using appropriate matrices $\bar{A}_f, \bar{A}_\theta, \bar{A}_P, \bar{A}_D, \bar{A}_t$, and vector $\hat{b}$, the formulation can be abbreviated as

$$\begin{array}{lll} \boldsymbol{K^*_{\mathcal{C}}(\mathcal{A})} : & t^*_{\mathcal{C}}(z^{\mathcal{A}}) & \doteq \quad \min\ t \\ \text{Subject to:} & & \\ & & \bar{A}_f f \ + \bar{A}_\theta \theta \ + \bar{A}_P P \ + \bar{A}_D D \ + \bar{A}_t t \ \geq \ \bar{b} \\ & & P_i^{min} y_i^{\mathcal{C}} \ \leq P_i \ \leq P_i^{max} y_i^{\mathcal{C}}, \quad \forall i \in \mathcal{G} \end{array}$$

Allowing the y quantities to become 0/1 variables, we obtain the problem

$$t^*(z^{\mathcal{A}}) \quad \doteq \quad \min\ t \tag{36}$$

Subject to:

$$\bar{A}_f f \ + \bar{A}_\theta \theta \ + \bar{A}_P P \ + \bar{A}_D D \ + \bar{A}_t t \ \geq \ \bar{b} \tag{37}$$

$$P_i^{min} y_i \ \leq P_i \ \leq P_i^{max} y_i, \quad \forall i \in \mathcal{G} \tag{38}$$

$$y_i \ = \ 0 \text{ or } 1, \quad \forall i \in \mathcal{G}. \tag{39}$$

This is the *controller's problem*: we have that $t^*(z^{\mathcal{A}}) \leq 1$ if and only if there exists some configuration of the generators that defeats $\mathcal{A}$.

However, for the purposes of this section, we will assume $\mathcal{C}$ is given and that the $y^{\mathcal{C}}$ are constants. We can now write the dual of $K^*_{\mathcal{C}}(\mathcal{A})$, suppressing the index $\mathcal{C}$ from the variables, for clarity.

$$\mathbf{A}_{\mathcal{C}}(\mathcal{A}) : \quad \max \quad \sum_{i \in cG} y_i^{\mathcal{C}} P_i^{min} \gamma_i^- - \sum_{i \in \mathcal{G}} y_i^{\mathcal{C}} P_i^{max} \gamma_i^+ - \sum_{j \in \mathcal{D}} D_j^{nom} \Delta_j + \sum_{j \in \mathcal{D}} D_j^{nom} \mu_j + \sum_{(i,j) \in E} \omega_{ij}$$

Subject to:

$$(f_{ij}) \qquad \alpha_i - \alpha_j + x_{ij}\beta_{ij} - p_{ij} + q_{ij} + \omega_{ij}^+ - \omega_{ij}^- \ = \ 0 \ \ \forall (i,j) \in E \tag{40}$$

$$(\theta_i) \qquad \sum_{(i,j) \in \delta^+(i)} \beta_{ij} - \sum_{(j,i) \in \delta^-(i)} \beta_{ji} \ = \ 0 \ \ \forall i \in V \tag{41}$$

$$(t) \qquad \sum_{(i,j) \in E} u_{ij}(p_{ij} + q_{ij}) + \sum_{(i,j) \in E} (\omega_{ij}^+ + \omega_{ij}^-) \ \leq \ 1 \tag{42}$$

$$(P_i) \qquad -\alpha_i - \gamma_i^- + \gamma_i^+ \ = \ 0 \ \ \forall i \in \mathcal{G} \tag{43}$$

$$(D_j) \qquad \alpha_j + \mu - \Delta_j \ \leq 0 \quad \forall j \in \mathcal{D} \tag{44}$$

$$(\xi_{ij}^+, \xi_{ij}^-) \quad x_{ij}^{1/2} |\beta_{ij}| \ \leq \ \mathbf{M}(1 - z_{ij}^{\mathcal{A}}) \ \ \forall (i,j) \in E \tag{45}$$

$$
\begin{aligned}
(\varrho_{ij}) \qquad & p_{ij} + q_{ij} \ \le \ \frac{1}{u_{ij}}(1 - z^{\mathcal{A}}_{ij}) \ \ \forall (i,j) \in E && (46)\\
(\eta_{ij}) \qquad & \omega^{+}_{ij} + \omega^{-}_{ij} \ \le \ z^{\mathcal{A}}_{ij} \ \ \forall (i,j) \in E && (47)\\
& \omega^{+}_{ij} \ge 0, \ \ \omega^{-}_{ij} \ge 0, \ \ p_{ij} \ \ge \ 0, \ \ q_{ij} \ge 0 \ \ \forall (i,j) \in E \\
& \gamma^{+}_{i}, \gamma^{-}_{i} \ \ \ge \ 0 \ \ \forall i \in \mathcal{G} \\
& \Delta_j \ge 0 \ \ \forall j \in \mathcal{D} \\
& \mu \ge 0 \\
& \delta_{ij}, \ \beta_{ij} \text{ free } \ \forall (i,j) \in E \\
& \alpha_i \text{ free } \ \forall i \in V.
\end{aligned}
$$

As before, for each constraint we indicate the corresponding dual variable. In (45), $\mathbf{M}$ is an appropriately chosen constant (we will provide a precise value for it below). Note that we are scaling β_{ij} by $x^{1/2}_{ij}$ – this is allowable since $x^{1/2}_{ij} > 0$; the reason for this scaling will become clear below.
Abbreviating

$$(\alpha^{\mathcal{C}}, \beta^{\mathcal{C}}, p^{\mathcal{C}}, q^{\mathcal{C}}, \omega^{\mathcal{C}+}, \omega^{\mathcal{C}-}, \gamma^{\mathcal{C}-}, \gamma^{\mathcal{C}+}, \mu^{\mathcal{C}}, \Delta^{\mathcal{C}}) = \psi^{\mathcal{C}},$$

we have that $A_{\mathcal{C}}(\mathcal{A})$ can be rewritten as:

$$\max \left\{ w^T_{\mathcal{C}} \, \psi^{\mathcal{C}} \ : \ A\psi^{\mathcal{C}} \ \le b \ + \ B \left(1 - z^{\mathcal{A}}\right) \right\} \tag{48}$$

where A, B, $w_{\mathcal{C}}$ and b are appropriate matrices and vectors. Consequently, we can now rewrite the min-cardinality attack problem:

$$\min \sum_{(i,j)} z_{ij} \tag{49}$$

$$
\begin{aligned}
\text{Subject to:} \qquad\qquad t^{\mathcal{C}} \ & \ge \ 1 + \epsilon, \quad \forall \ \mathcal{C} \subseteq \mathcal{G} && (50)\\
w^T_{\mathcal{C}} \, \psi^{\mathcal{C}} \ - \ t^{\mathcal{C}} \ & \ge \ 0, \quad \forall \ \mathcal{C} \subseteq \mathcal{G}, && (51)\\
A\psi^{\mathcal{C}} \ + \ Bz \ & \le \ b \ + \ B \quad \forall \ \mathcal{C} \subseteq \mathcal{G}, && (52)\\
z_{ij} \ & = \ 0 \text{ or } 1, \ \ \forall \ (i,j). && (53)
\end{aligned}
$$

This formulation, of course, is exponentially large. An alternative is to use Benders cuts – having solved the linear program $A_{\mathcal{C}}(\mathcal{A})$, let $(\bar{f}, \bar{\theta}, \bar{t}, \bar{P}, \bar{D}, \bar{\xi}^{+}, \bar{\xi}^{-}, \bar{\varrho}, \bar{\eta})$ be optimal dual variables. Then the resulting Benders cut is

$$t^{\mathcal{C}} + \sum_{(i,j) \in E} ((\bar{\xi}^{+}_{ij} + \bar{\xi}^{-}_{ij})\mathbf{M}(1 - z_{ij})) + \sum_{(i,j) \in E} (\frac{1}{u_{ij}} \bar{\varrho}_{ij}(1 - z_{ij})) + \sum_{(i,j) \in E} \bar{\eta}_{ij} z_{ij} \ge 1 + \epsilon, \tag{54}$$

We can now update our algorithmic template for the min-cardinality problem.

Updated algorithm for min-cardinality attack problem

Iterate:

1. Attacker: Solve master attacker MIP, obtaining attack $\mathcal{A}$.

2. Controller: Solve the controller's problem (36)-(39) to search for a set $\mathcal{C}$ of generators such that $t^{*}_{\mathcal{C}}(z^{\mathcal{A}}) \ \le \ 1$.

(2.a) If no such set $\mathcal{C}$ exists, **EXIT**:
$\mathcal{A}$ is a minimum cardinality successful attack.

(2.b) Otherwise, suppose such a set $\mathcal{C}$ is found. Then
(2.b.1) Add to the master the Benders' cut (54), and, optionally
(2.b.2) Add to the master the entire system (50)-(52),
Go to 1.

Clearly, option (2.b.2) can only be exercised sparingly (if ever). Below we will discuss how we choose, in our implementation, between (2.b.1) and (2.b.2). We will also describe how to (significantly) strengthen the straightforward Benders cut (54). One point to note is that the cuts (or systems) arising from different configurations $\mathcal{C}$ reinforce one another.

At each iteration of the algorithm, the master attacker MIP becomes a stronger relaxation for the min-cardinality problem, and thus its solution in step 1 provides a lower bound for the problem. Thus, if in a certain execution of step 2 we certify that $t^*_{\mathcal{C}}(z^{\mathcal{A}}) > 1$ for every configuration $\mathcal{C}$, we have solved the min-cardinality problem to optimality.

What we have above is a complete outline of our algorithm. In order to make the algorithm effective we need to further sharpen the approach. In particular, we need set the constant $\mathbf{M}$ to as small a value as possible, and we need to develop stronger inequalities than the basic Benders' cuts.

2.2.1 Setting M

In this section we show how to set for $\mathbf{M}$ a value that does not grow with network size.

Lemma 2.5 *We can set*

$$\mathbf{M} = \max_{(i,j)\in E} \left\{ \frac{1}{\sqrt{x_{ij}}\, u_{ij}}. \right\} \tag{55}$$

Proof. Given an attack $\mathcal{A}$, consider a connected component K of $\mathcal{N} - \mathcal{A}$. For any arc (i,j) with both ends in K, $\omega^+_{ij} + \omega^-_{ij} = 0$ by (47). Hence we can rewrite constraints (40)-(41) over all arcs with both ends in K as follows:

$$\begin{aligned} N_K^T \alpha_K + X_K \beta_K &= p_K - q_K, & (56)\\ N_K \beta_K &= 0. & (57) \end{aligned}$$

Here, N_K is the node arc incidence matrix of K, $\alpha_K, \beta_K, p_K, q_K$ are the restrictions of α, β, p, q to K, and X_K is the diagonal matrix $diag\{x_{ij} \, : \, (i,j) \in K\}$. From this system we obtain

$$N_K X_K^{-1} N_K \alpha_K = N_K X_K^{-1}(p_K - q_K). \tag{58}$$

The matrix $N_K X_K^{-1} N_K$ has one-dimensional null space and thus we have one degree of freedom in choosing α_K. Thus, to solve (58), we can remove from N_K an arbitrary row, obtaining $\tilde{N}_K$, and remove the same row from α_K, obtaining $\tilde{\alpha}_K$. Thus, (58) is equivalent to:

$$\tilde{N}_K X_K^{-1} \tilde{N}_K \tilde{\alpha}_K = \tilde{N}_K X_K^{-1}(p_K - q_K), \tag{59}$$

The matrix $\tilde{N}_K X_K^{-1} \tilde{N}_K$ and thus (59) has a unique solution (given $p_K - q_K$); we complete this to a solution to (58) by setting to zero the entry of α_K that was removed. Moreover,

$$X_K^{-1/2} N_K^T \alpha_K = X_K^{-1/2} \tilde{N}_K^T \tilde{\alpha}_K = X_K^{-1/2} \tilde{N}_K^T (\tilde{N}_K X_K^{-1} \tilde{N}_K^T)^{-1} \tilde{N}_K X_K^{-1} (p_K - q_K). \tag{60}$$

The matrix

$$M = X_K^{-1/2} \tilde{N}_K^T \, (\tilde{N}_K X_K^{-1} \tilde{N}_K^T)^{-1} \, \tilde{N}_K X_K^{-1/2}$$

is symmetric and idempotent, e.g. $MM^T = I$. Thus, from (60) we get

$$\|X_K^{-1/2} N_K^T \alpha_K\|_2 \le \|M\|_2 \, \|X_K^{-1/2}(p_K - q_K)\|_2 \le \|X_K^{-1/2}(p_K - q_K)\|_2, \tag{61}$$

where the last inequality follows from the idempotent attribute. Because of constraints (42), (46) and (47), we can see that the right-hand side of (61) is upper-bounded by the value of the convex maximization problem,

$$\begin{aligned} \max \quad & \sum_{(i,j)\in E} x_{ij}^{-1}(p_{ij} - q_{ij})^2 && (62)\\ \text{s.t.} \quad & \sum_{(i,j)\in E} u_{ij}(p_{ij} + q_{ij}) \leq 1 && (63)\\ & p_{ij} \geq 0,\ q_{ij} \geq 0, && (64) \end{aligned}$$

which, as is easily seen, equals

$$\max_{(i,j)\in E} \left\{ \frac{1}{x_{ij}u_{ij}^2} \right\}.$$

■

2.2.2 Tightening the formulation

In this Section we describe a family of inequalities that are valid for the attacker problem. These cuts seek to capture the interplay between the flow conservation equations and Ohm's law. First we present a technical result.

Lemma 2.6 *Let Q be matrix with r rows with rank r, and let $A = Q^T(QQ^T)^{-1}Q \in \mathcal{R}^{r\times r}$. Let $B := I - A$. Then for any $p \in \mathcal{R}^r$ we have*

$$\begin{aligned} \|p\|_2^2 \quad & = \|Ap\|_2^2 + \|Bp\|_2^2 && (65)\\ \|p\|_1 \quad & \geq |(Ap)_j| + |(Bp)_j| \quad \forall j = 1 \ldots r && (66) \end{aligned}$$

Proof. A and B are symmetric and idempotent, i.e., $A^2 = A$, $B^2 = B$, and any $p \in \mathcal{R}^r$ can be written as $p = Ap + Bp$. Multiplying equation this by p and using the fact that A and B are symmetric and idempotent we get (65):

$$\begin{aligned} p^Tp \quad & = \quad p^TAp + p^TBp && (67)\\ & = \quad p^TA^2p + p^TB^2p && (68)\\ \|p\|_2^2 \quad & = \quad \|Ap\|_2^2 + \|Bp\|_2^2 && (69) \end{aligned}$$

We also have $A^TB = A(I - A) = A - A^2 = 0$, so $y^TA^TBy = 0$ for any $y \in \mathcal{R}^r$. Thus, if we rename $Ap = x$ and $Bp = y$, then the following holds: $p = x + y,\ x^Ty = 0,\ \|p\|_2^2 = \|x\|_2^2 + \|y\|_2^2$.

Let $1 \leq j \leq r$. We have

$$\|p\|_2^2 - (|x_j| + |y_j|)^2 = \|x\|_2^2 + \|y\|_2^2 - (|x_j| + |y_j|)^2 = \sum_{i,i\neq j} x_i^2 + \sum_{i,i\neq j} y_i^2 - 2|x_jy_j|$$

where the first equality follows from (65). Since $x^Ty = 0$, we have $|x_jy_j| = |\sum_{i,i\neq j} x_iy_i|$. Hence,

$$\begin{aligned} \sum_{i,i\neq j} x_i^2 + \sum_{i,i\neq j} y_i^2 - 2|x_jy_j| \quad & = \quad \sum_{i,i\neq j} x_i^2 + \sum_{i,i\neq j} y_i^2 - 2\left|\sum_{i,i\neq j} x_iy_i\right| && (70)\\ & \geq \quad \sum_{i,i\neq j} x_i^2 + \sum_{i,i\neq j} y_i^2 - 2\sum_{i,i\neq j} |x_iy_i| && (71)\\ & = \quad \sum_{i,i\neq j} (|x_i| - |y_i|)^2 && (72)\\ & \geq \quad 0 && (73) \end{aligned}$$

So we have $\|p\|_2^2 - (|x_j| + |y_j|)^2 \geq 0$, which implies $\|p\|_1 \geq \|p\|_2 \geq (|x_j| + |y_j|)\ \ \forall j = 1 \ldots r$. ■

As a consequence of this result we now have:

Lemma 2.7 *Given configuration $\mathcal{C}$, the following inequalities are valid for system (52)-(53) for each $(i,j) \in E$:*

$$x_{ij}^{-\frac{1}{2}}|\alpha_i^{\mathcal{C}} - \alpha_j^{\mathcal{C}}| + x_{ij}^{\frac{1}{2}}|\beta_{ij}^{\mathcal{C}}| \le x_{ij}^{-\frac{1}{2}} w_{ij}^{\mathcal{C}} + \mathbf{M}(1 - z_{ij}) \tag{74}$$

$$x_{ij}^{-\frac{1}{2}}|\alpha_i^{\mathcal{C}} - \alpha_j^{\mathcal{C}}| + x_{ij}^{\frac{1}{2}}|\beta_{ij}^{\mathcal{C}}| \le \sum_{(k,l)} x_{kl}^{-\frac{1}{2}}(p_{kl}^{\mathcal{C}} + q_{kl}^{\mathcal{C}}) + w_{ij}^{\mathcal{C}} \tag{75}$$

where $\mathbf{M} := max_{(k,l)\in E}\{\frac{1}{\sqrt{x_{kl}}u_{kl}}\}$ *as before.*

Proof. Suppose first that $z_{ij} = 0$. Let K be the component containing (i,j) after the attack. Then by (60) and (56),

$$\begin{aligned} X^{-1/2} N_K^T \alpha^{\mathcal{C}} &= A X^{-1/2}(p^{\mathcal{C}} - q^{\mathcal{C}}), && (76)\\ X^{1/2}\beta^{\mathcal{C}} &= (I - A) X^{-1/2}(p^{\mathcal{C}} - q^{\mathcal{C}}), && (77) \end{aligned}$$

where $A = X^{-1/2}\tilde{N_K}^T(\tilde{N_K} X^{-1} \tilde{N_K}^T)^{-1}\tilde{N_K} X^{-1/2}$. Thus, we have

$$x_{ij}^{-1/2}|\alpha_i^{\mathcal{C}} - \alpha_j^{\mathcal{C}}| + x_{ij}^{1/2}|\beta_{ij}^{\mathcal{C}}| \;\; \le \;\; \sum_{(k,l)} x_{kl}^{-1/2}(p_{kl}^{\mathcal{C}} + q_{kl}^{\mathcal{C}}) \;\; \le \;\; \mathbf{M} \tag{78}$$

where the first inequality follows from (66) proved in Lemma 2.6, and the second bound is obtained as in the proof of Lemma 2.5.

Suppose now that $z_{ij} = 1$. Here we have $|\alpha_i^{\mathcal{C}} - \alpha_j^{\mathcal{C}}| \le \omega_{ij}^{\mathcal{C}}$, by (40), (46), (45). Using these (74)-(75) can be easily shown. ■

Inequalities (65)-(66) strengthen system (52)-(53); when case step (2.b.2) of the min-cardinality algorithm is applied then (65), (66) will become part of the master problem. If case (2.b.1) is applied, then the vector $\psi^{\mathcal{C}} = (\alpha^{\mathcal{C}}, \beta^{\mathcal{C}}, p^{\mathcal{C}}, q^{\mathcal{C}}, \omega^{\mathcal{C}+}, \omega^{\mathcal{C}-}, \gamma^{\mathcal{C}-}, \gamma^{\mathcal{C}+}, \mu^{\mathcal{C}}, \Delta^{\mathcal{C}})$ is expanded by adding two new dual variables per arc (i,j).

2.2.3 Strengthening the Benders cuts

Typically, the standard Benders cuts (54) prove weak. One manifestation of this fact is that in early iterations of our algorithm for the min-cardinality attack problem, the attacks produced in Step 1 will tend to be "weak" and, in particular, of very small cardinality. Here we discuss two routines that yield substantially stronger inequalities, still in the Benders mode.

In Step 2 of the algorithm, given an attack $\mathcal{A}$, we discover a generator configuration $\mathcal{C}$ that defeats $\mathcal{A}$, and from this configuration a cut is obtained. However, it is not simply the configuration that defeats $\mathcal{A}$, but, rather, a vector of power flows. If we could somehow obtain a "stronger" vector of power flows, the resulting cut should prove tighter. To put it differently, a vector of power flows that are in some sense "minimal" might also defeat other attacks $\mathcal{A}'$ that are "stronger“ than $\mathcal{A}$; in other words, they should produce stronger inequalities. One way of thinking about this is in analogy with the classical max-flow min-cut paradigm for single commodity flows.

We implement this rough idea in two different ways. Consider Step 2 of the min-cardinality attack algorithm, and suppose case (2.b) takes place. We execute steps I and II below, where in each case $\mathcal{A}^*$ is initialized as $E - \mathcal{A}$, and f^* is initialized as the power flow that defeated $\mathcal{A}$:

(I) First, we add the Benders' cut (54).
Also, initializing $\mathcal{B} = \mathcal{A}$, we run the following step, for $k = 1, 2, \ldots, |E - \mathcal{A}|$:

(I.0) Let $(i_k, j_k) \; = \; \text{argmin}\Big\{|f_{ij}^*| \; : \; (i,j) \; \in \; \mathcal{A}^*\Big\}$.

(I.1) If the attack $\mathcal{B} \cup (i_k, j_k)$ is *not* successful, then reset $\mathcal{B} \leftarrow \mathcal{B} \cup (i_k, j_k)$, and update f^* to the power flow that defeats the (new) attack $\mathcal{B}$.

(I.2) Reset $\mathcal{A}^* \leftarrow \mathcal{A}^* - (i_k, j_k)$.

At the end of the loop, we have an attack $\mathcal{B}$ which is not successful, i.e. $\mathcal{B}$ is defeated by some configuration $\mathcal{C}'$. If $\mathcal{B} = \mathcal{A}$ we do nothing. Otherwise, we add to the master problem the Benders cut arising from $\mathcal{B}$ and $\mathcal{C}'$.

(II) Set $\mathcal{F} = \emptyset$ and $\mathcal{C}' = \mathcal{C}$. We run the following step, for $k = 1, 2, \ldots, |E - \mathcal{A}|$:

(II.0) Let $(i_k, j_k) \in \mathcal{A}^*$ be such that its flow has minimum absolute value.

(II.1) Test whether $\mathcal{A}$ is successful against a controller which is forced to satisfy the condition

$$f_{ij} \,=\, 0, \;\; \forall \;(i,j) \,\in\, \mathcal{F} \cup (i_k, j_k). \tag{79}$$

(II.2) If *not* successful, let $\mathcal{C}'$ be the configuration that defeats the attack, and reset f^* to the corresponding power flow that satisfies (79). Reset $\mathcal{F} \leftarrow \mathcal{F} \cup (i_k, j_k)$,

(II.3) Reset $\mathcal{A}^* \leftarrow \mathcal{A}^* - (i_k, j_k)$.

Comment. Procedure (I) produces attacks of increasing cardinality. At termination, if $\mathcal{A} \neq \mathcal{B}$, then and $\mathcal{C} \neq \mathcal{C}'$, and yet $\mathcal{B}$ is still not successful. In some sense in this case $\mathcal{C}'$ is a 'stronger' configuration than $\mathcal{C}$ and the resulting Benders' cut 'should' be tighter than the one arising from $\mathcal{C}$ and $\mathcal{A}$. We say 'should' because the previously discussed non-monotonicity property of power flow problems could mean that $\mathcal{C}'$ does not defeat $\mathcal{A}$. Nevertheless, *in general*, the new cut is indeed stronger.

In contrast with (I), procedure (II) considers a progressively weaker controller. In fact, because we are forcing flows to zero, but we are not voiding Ohm's equation (2), the power flow that defeats $\mathcal{A}$ while satisfying (79) is a feasible power flow for the original network. Thus, at termination of the loop,

$$\mathcal{C}' \text{ defeats every attack } \mathcal{A}' \text{ of the form } \mathcal{A}' \,=\, \mathcal{A} \cup \mathcal{E} \text{ for each } \mathcal{E} \subseteq \mathcal{F}.$$

Thus, if $\mathcal{F} \neq \emptyset$ the cut obtained in (II) should be particularly strong.

One final comment on procedures (I) and (II) is that each "test" requires the solution of the controller's problem (36)-(39), a mixed-integer program. In our testing, such problems can be solved *extremely* fast using a commercial solver.

2.3 Implementation details

Our implementation is based on the updated algorithmic outline given in Section 2.2. In step (2.b.1) we add the Benders' cut with strengthening as in section 2.2.3, so we may add two cuts. We execute Step (2.b.2) so that the relaxation includes up to two full systems (50)-(52) at any time: when a system is added at iteration k, say, it is replaced at iteration $k+4$ by the system corresponding to the configuration $\mathcal{C}$ discovered in Step 2 of that iteration. Because at each iteration the cut(s) added in step (2.b.1) cut-off the current vector $z^{\mathcal{A}}$, the procedure is guaranteed to converge.

2.4 Computational results the with min-cardinality model

. We tested our algorithm on a number of problems based on networks derived from the IEEE test cases [16]: (a) a 49-node, 84-arc network with 14 demand nodes, and (b) a 98- node, 204-arc network, with 28 demand nodes

Tables 1 presents experiments with our algorithm on the 49-node, 84-arc network, first using a set with 4 and then using 8. The sum of maximum generator outputs, $\sum_{i \in \mathcal{G}} P^{max})i$, is the same for both cases; the demand nodes and their nominal demand values are identical.

Each table shows in summarized form the progress of the algorithm. Each row corresponds to a value of the minimum throughput T^{min}, while each column corresponds to an attack cardinality. For each (row, column) combination, the corresponding cell is labeled "Not Enough" when using any attack of the corresponding cardinality (or smaller) the attacker will not be able to reduce demand below the stated throughput, while "Success" means that some attack of the given cardinality (or smaller) does succeed. Further, we also indicate the number of iterations that the algorithm took in order to prove the given outcome (shown in parentheses) as well as the corresponding CPU time in seconds.

49 nodes, 84 arcs
Entries show: (iteration count), time,
Attack status (**F** = cardinality too small, **S** = attack success)

4 generators

	Attack cardinality			
Min. throughput	**2**	**3**	**4**	**5**
0.84	(4), 129, **F**	(4), 129, **S**		
0.82	(4), 364, **F**	(35), 1478, **F**	(36), 1484, **S**	
0.78	(4), 442, **F**	(4), 442, **F**	(26), 746, **S**	
0.74	(4), 31, **F**	(11), 242, **F**	(168), 4923, **F**	(168), 4923, **S**
0.70	(3), 31, **F**	(4), 198, **F**	(10), 1360, **F**	(203), 3067, **S**
0.62	(4), 86, **F**	(4), 86, **F**	(131), 2571, **F**	(450), 34298, **F**

8 generators

	Attack cardinality			
Min. throughput	**2**	**3**	**4**	**5**
0.90	(1), 13, **F**	(3), 133, **S**		
0.86	(1), 59, **F**	(5), 357, **F**	(13), 1291, **S**	
0.84	(1), 48, **F**	(4), 227, **F**	(41), 2532, **F**	(43), 2535, **S**
0.80	(1), 14, **F**	(4), 210, **F**	(8), 1689, **F**	(50), 2926, **S**
0.74	(1), 8, **F**	(3), 101, **F**	(10), 1658, **F**	(68), 23433, **F**

Table 1: ***Min-cardinality problem, small network***

Thus, for example, in Table 1, the algorithm *proved* that using an attack of size 4 or smaller we cannot reduce total demand below 70% of the nominal value; this required 10 iterations which overall took 1360 seconds. At the same time, in 203 iterations (3067 seconds) the algorithm found a successful attack of cardinality 5.

Not surprisingly, the network with 8 generators proves more resilient – for example, an attack of cardinality 5 is needed to reduce throughput below 84%, whereas the same can be achieved with an attack of size 3 in the case of the 4-generator network. Also note that the running-time performance does not significantly degrade as we move to the 8-generator case, even though the number of generator configurations is 511.

Table 2 describe similar tests, but now on the 98-node, 204-arc network. Note that in the 15 generator case there are over 30000 generator configurations that must be examined, at least implicitly, in order to certify that a given attack is successful.

2.4.1 Comparison with pure enumeration

Here we compare our algorithm with the pure enumeration approach. As noted before, even though the controller's problem (36)-(39) is a mixed-integer program, modern commercial solvers handle

98 nodes, 204 arcs Entries show: (iteration count), time, Attack status (**F** = cardinality too small, **S** = attack success)			
10 generators			
	Attack cardinality		
Min. throughput	**2**	**3**	**4**
0.89	2, 177, **F**	30, 555, **S**	
0.86	(2), 195, **F**	(12), 5150, **F**	(14), 5184, **S**
0.84	(2), 152, **F**	(11), 7204, **F**	(35), 223224, **F**
0.82	(2), 214, **F**	(9), 11458, **F**	(16), 225335, **F**
0.75	(2), 255, **F**	(9), 5921, **F**	(17), 151658, **F**
0.60		(1), 4226, **F**	N/R
12 generators			
	Attack cardinality		
Min. throughput	**2**	**3**	**4**
0.92	(2), 318, **F**	(11), 7470, **F**	(14), 11819, **S**
0.90	(2), 161, **F**	(11), 14220, **F**	(18), 16926, **S**
0.88	(2), 165, **F**	(10), 11178, **F**	(15), 284318, **S**
0.84	(2), 150, **F**	(9), 4564, **F**	(16), 162645, **F**
0.75	(2), 130, **F**	(9), 7095, **F**	(15), 93049, **F**
15 generators			
	Attack cardinality		
Min. throughput	**2**	**3**	**4**
0.94	(2), 223, **F**	(11), 654, **S**	
0.92	(2), 201, **F**	(11), 10895, **F**	(18), 11223, **S**
0.90	(2), 193, **F**	(11), 6598, **F**	(16), 206350, **S**
0.88	(2), 256, **F**	(9), 15445, **F**	(18), 984743, **F**
0.84	(2), 133, **F**	(9), 5565, **F**	(15), 232525, **F**
0.75	(2), 213, **F**	(9), 7550, **F**	(11), 100583, **F**

Table 2: ***Min-cardinality problem, larger network***

it with ease. Thus the enumeration approach, where we enumerate all possible attacks of a given cardinality, should be applicable at least in case of small problems. When a successful attack of the cardinality under consideration exists, the enumeration approach might "get lucky" and find it quickly; on the other hand when the given cardinality is insufficient to defeat the controller *all* attacks will need to be enumerated.

In order to effect a comparison, we first estimated, for each network, the time needed to solve one controller's problem by choosing 1000 random attacks and averaging their solution time. We then multiplied this estimated average time by the number of cases that need to be enumerated.

In the following tables we tabulate the projected time(in seconds) it would take if a pure numeration approach was used. The column 'time per MIP' indicates the average time (in seconds) taken by $CPLEX$ to solve one instance of controller MIP. The following table summarizes our results; the numbers in parentheses indicate the total number of enumerations required, while each cell entry indicates the projected total CPU time.

2.4.2 One configuration problems

For completeness, in Table 4 we present results where we study *one-configuration* problems where the set of generators that the controller operates are *fixed*. Problems of this type correspond most closely to those previously studied in the literature. Here we applied the mixed-integer programming

		Attack cardinality		
		2	**3**	**4**
		(20706)	(1394204)	(7005871)
10 generators				
Min. throughput	**Time per MIP**			
0.89	0.051550	1067	71870	
0.86	0.052284	1083	72894	3662973
0.84	0.052853	1094	73687	3702811
0.82	0.055451	1148	77310	3884826
0.75	0.077676	1608	108296	5441916
0.60	0.110078	2279	153471	7711957
12 generators				
Min. throughput	**Time per MIP**			
0.94	0.0546667	1132	76216	
0.92	0.056725	1174	79086	3974116
0.90	0.052853	1685	113518	5704293
0.88	0.063490	1314	88518	4448030
0.84	0.090882	1881	126708	6367104
0.75	0.113589	2351	158365	7957849
15 generators				
Min. throughput	**Time per MIP**			
0.92	0.066127	1369	92195	4632806
0.90	0.052853	1685	113518	5704293
0.88	0.097627	2024	136290	6848586
0.84	0.116882	2420	162957	8188631
0.75	0.124245	2576	173496	7711927

Table 3: ***Pure enumeration, 98 nodes 204 arcs***

formulation (49)-(53) restricted to the single configuration $\mathcal{C} = \mathcal{G}$. Rather than use our algorithm, we simply solved these problems using a commercial solver, Cplex [13]. The table shows the CPU time needed to solve the minimum-cardinality problem corresponding to the minimum throughput shown in the first column.

As expected, the problem becomes *easier* as the required minimum attack size increases – more candidates (for optimal attack) exist.

3 A continuous, nonlinear attack problem

In this section we study a new attack model. Our goals are twofold:

- First, we want to more explicitly capture how the flow conservation equations (1) interact with the power-flow law (2) in order to produce flows in excess of capacities. More generally, we are interested in directly incorporating the interaction of the laws of physics with the graph-theoretic structure of the network into an algorithmic procedure. It is quite clear that the complexity of combinatorial problems on power flows, such as the min-cardinality attack problem, is primarily due to this interaction.

- Second, there are ways other than the outright disabling of a power line, in which the functioning of the line could be hampered. There is a sense (see e.g. [23]) that recent real-world blackouts were not simply the result of discrete line failures; rather the system as a whole was already under "stress" when the failures took place. In fact, the operation of a power grid can be viewed as a noisy process, this in addition to the fact that even the AC power flow model

Min. Throughput	Min. Attack Size	Time (sec.)
0.95	2	2
0.90	3	20
0.85	4	246
0.80	5	463
0.75	6	2158
0.70	6	1757
0.65	7	3736
0.60	7	1345
0.55	8	2343
0.50	8	1328

Table 4: **49 nodes, 84 arcs, one configuration**

is an approximation. Rather than attempting to model the noise and complexity in detail, we seek a generic modeling methodology that can serve to expose system vulnerabilities.

The approach we take relies on the fact that one can approximate a variety of complex physical phenomena that (negatively) affect the performance of a line by simply perturbing that line's resistance (or, for AC models, the conductance, susceptance, etc.). In particular, by significantly increasing the resistance of an arc we will, in general, force the power flow on that line to zero. This modeling approach becomes particularly effective, from a system perspective, when the resistances of many arcs are simultaneously altered in an *adversarial* fashion.

Accordingly, our second model works as follows:

(I) The attacker *sets* the resistance x_{ij} of any arc (i, j).

(II) The attacker is constrained: we must have $x \in F$ for a certain known set F.

(III) The output of each generator i is fixed at a given value P_i, and similarly each demand value D_i is also fixed at a given value.

(IV) The objective of the attacker is to maximize the overload of any arc, that is to say, the attacker wants to solve

$$\max_{x \in F} \max_{ij} \left\{ \frac{|f_{ij}|}{u_{ij}} \right\}, \tag{80}$$

where the f_{ij} are the resulting power flows.

In view of Lemma 1.1, (III) implies that in (d) the vector f is unique for each choice of x; thus the problem is well-posed.

In future work we plan to relax (III). But (I), (II), (IV) already capture a great deal of the inherent complexity of power flows. Moreover, suppose that e.g. the value of (80) equals 1.25. Then even if we allow demands to be reduced, but insist that this be done under a *fair* demand-reduction discipline (one that decreases all demands by the same factor) the system will lose 25% of the total demand if overloads are to be avoided (and it is not surprising that the same qualitative conclusion holds even if demands are "unfairly" reduced to minimize maximum overload; see Table 11). Thus we expect that the impact of (III), under this model, may not be severe.

For technical reasons, it will become more convenient to deal with the inverses of resistances, the so-called "conductances." For each $(i, j) \in E$, write $y_{ij} = 1/x_{ij}$, and let y be the vector of y_{ij}. Then we are interested in a problem of the form

$$\max_{y \in \Gamma} \max_{ij} \left\{ \frac{|f_{ij}(y)|}{u_{ij}} \right\}, \tag{81}$$

where Γ is an appropriate set, and as just discussed the notation $f_{ij}(y)$ is justified.

A relevant example of a set Γ is that given by:

$$\sum_{ij} \frac{1}{y_{ij}} \le B, \qquad \frac{1}{x_{ij}^U} \le y_{ij} \le \frac{1}{x_{ij}^L} \quad \forall\,(i,j), \tag{82}$$

where B is a given 'budget', and, for any arc (i,j), x_{ij}^L and x_{ij}^U and indicates a minimum and maximum value for the resistance at (i,j). Suppose the initial resistances x_{ij} are all equal to some common value $\bar{x}$, and we set $x_{ij}^L = \bar{x}$ for every (i,j), and $B = k\,\theta\,\bar{x} + (|E| - k)\bar{x}$, where $k > 0$ is an integer and $\theta > 1$ is large. Then, roughly speaking, we are approximately allowing the adversary to make the resistance of (up to) k arcs "very large", while not decreasing any resistance, a problem closely reminiscent of the classical $N - K$ problem. We will make this statement more precise later.

If the objective in (81) is convex then the optimum will take place at some extreme point. In general, the objective is not convex; but computational experience shows that we tend to converge to points that are either extreme points, or very close to extreme points (see the computational section).

3.1 Solution methodology

Problem (81) is not smooth. However, it is equivalent to:

$$\max_{y,p} \quad \sum_{ij} \frac{f_{ij}(y)}{u_{ij}}(p_{ij} - q_{ij}) \tag{83}$$

$$\text{s.t.} \quad \sum_{ij}(p_{ij} + q_{ij}) \;=\; 1, \tag{84}$$

$$y \in \Gamma, \quad p, q \ge 0. \tag{85}$$

In order to work with this formulation we need to develop a more explicit representation of the functions $f_{ij}(y)$. This will require a sequence of technical results given in the following section; however a brief discussion of our approach follows.

Problem (83), although smooth, is not concave. A relatively recent research thrust has focused on adapting techniques of (convex) nonlinear programming to nonconvex problems. This work has resulted in a very large literature with interesting and useful results; see [15], [4]. Since one is attempting to solve non-convex minimization (and thus, NP-hard) problems, there is no guarantee that a global optimum will be found by these techniques. One can sometimes assume that a global optimum is approximately known; and the techniques then are likely to converge to the optimum from an appropriate guess.

In any case, (a) the use of nonlinear models allows for much richer representation of problems, (b) the very successful numerical methodology backing convex optimization is brought to bear, and (c) even though only a local optimum may be found, at least one is relying on an agnostic, "honest" optimization technique as opposed to a pure heuristic or a method that makes structural assumptions about the nature of the optimum in order to simplify the problem.

In our approach we will indeed rely on this methodology – items (a)-(c) precisely capture the reasons for our choice. Points (a) and (c) are particularly important in our blackout context: we are very keen on modeling the nonlinearities, and on using a truly agnostic algorithm to root out hidden weaknesses in a network. And from a computational perspective, the approach does pay off, because we are able to comfortably handle problems with on the order of 1000 arcs.

As a final point, note that in principle one could rely on a branch-and-bound procedure to actually find the global optimum. This will be a subject for future research.

3.1.1 Laplacians

In this section we present some background material on linear algebra and Laplacians of graphs – the results are standard but we include a proof for completeness and continuity. See [7] for relevant material.

As before we have a directed network G with n nodes and m arcs and with node-arc incidence matrix N. As before we assume G is connected. For a positive diagonal matrix $Y \in \mathcal{R}^{m \times m}$ we will write

$$L \,=\, NYN^T, \qquad J \,=\, L + \frac{1}{n}\,\mathbf{1}\mathbf{1}^T. \tag{86}$$

where $\mathbf{1} \in \mathcal{R}^n$ is the vector $(1, 1, \ldots, 1)^T$. L is called a *generalized Laplacian.* We have that L is symmetric positive-semidefinite. If $\lambda_1 \le \lambda_2 \le \ \ldots\ \le \lambda_n$ are the eigenvalues of L, and $v^1, v^2, \ldots, v^n$ are the corresponding unit-norm eigenvectors, then

$$\lambda_1 = 0, \quad \text{but} \quad \lambda_i > 0 \quad \text{for } i > 1, \tag{87}$$

because G is connected, and thus L has rank $n-1$. The same argument shows that since $N\mathbf{1} = 0$, we can assume $v^1 \,=\, n^{-1/2}\ \mathbf{1}$. Finally, since different eigenvectors are are orthogonal, we have $\mathbf{1}^T v^i \,=\, 0$ for $2 \le i \le n$.

Lemma 3.1 *L and J have the same eigenvectors, and all but one of their eigenvalues coincide. Further, J is invertible.*

Proof. By (87),

$$Lv^1 \,=\, 0, \qquad Jv^1 \,=\, \frac{1}{n}\,\mathbf{1}\mathbf{1}^T v^1 \,=\, v^1, \tag{88}$$

and further

$$Jv^i \,=\, Lv^i \,=\, \lambda_i v^i.\blacksquare \tag{89}$$

Lemma 3.2 *Let $b \in \mathcal{R}^n$. Any solution to the system of equations $L\alpha = b$ is of the form*

$$\alpha = J^{-1}b + \delta\mathbf{1},$$

for some $\delta \in \mathcal{R}$.

Proof. We have that $L \,=\, \sum_{i=2}^n \lambda_i v_i v_i^T$, and, by Lemma 3.1, $J^{-1} \,=\, \sum_{i=2}^n \frac{1}{\lambda_i} v_i v_i^T \,+\, \frac{1}{n}\mathbf{1}\mathbf{1}^T$. Now, the system of equations $L\alpha = b$ is feasible if and only if b lies in the column space of matrix L and when it is so we can write $b = \sum_{i=2}^n v_i(v_i^T b)$. Assuming that this is the case, defining

$$\hat{\alpha} \,\doteq\, J^{-1}b \;\;=\;\; \sum_{i=2}^n \frac{1}{\lambda_i} v_i(v_i^T b) \tag{90}$$

we will have $L\hat{\alpha} \,=\, b$. Suppose that $\bar{\alpha}$ is another vector satisfying $L\bar{\alpha} = b$. Then $L(\bar{\alpha} - \hat{\alpha}) = 0$, and consequently $\bar{\alpha} \,=\, \hat{\alpha} + \delta\mathbf{1}$, for some δ. $\blacksquare$

Define

$$P = I - J.$$

Note that the eigenvalues of P are 0 and $1 - \lambda_i$, $2 \le i \le n$; thus if we have

$$\sum_{(u,v)} y_{uv} \,<\, 1/2, \quad \text{for all } u, \tag{91}$$

then it is not difficult to show that

$$0 \; < \; 1 - \lambda_i < 1, \qquad \text{for all } i \geq 2. \tag{92}$$

(See [19] for related background). In such a case we can write

$$J^{-1} \; = \; (I - P)^{-1} \; = \; I + P + P^2 + P^3 + \ldots, \tag{93}$$

in other words, the series in (93) converges to J^{-1}.

Lemma 3.3 *For any integer $k > 0$, $P^k \; = \; (I - NYN^T)^k - \frac{1}{n}\mathbf{1}\mathbf{1}^T$.*

Proof. We will prove the statement by induction on k, while also proving that $(I - NYN^T)^k\mathbf{1}\mathbf{1}^T = \mathbf{1}\mathbf{1}^T$. The case $k = 1$ holds by definition. For the general inductive step, we have

$$\begin{aligned} P^{k+1} \; &= \; \left[(I - NYN^T)^k - \frac{1}{n}\mathbf{1}\mathbf{1}^T\right] P & (94) \\ &= \; (I - NYN^T)^{k+1} \; - \; \frac{1}{n}(I - NYN^T)^k\mathbf{1}\mathbf{1}^T \; - \; \frac{1}{n}\mathbf{1}\mathbf{1}^T\left[(I - NYN^T) - \frac{1}{n}\mathbf{1}\mathbf{1}^T\right] & (95) \\ &= \; (I - NYN^T)^{k+1} - \frac{1}{n}\mathbf{1}\mathbf{1}^T, & (96) \end{aligned}$$

because by induction

$$(I - NYN^T)^k\mathbf{1}\mathbf{1}^T \;\; = \;\; (I - NYN^T)^{k-1}(I - NYN^T)\mathbf{1}\mathbf{1}^T \; = \; \mathbf{1}\mathbf{1}^T, \tag{97}$$

and

$$\mathbf{1}\mathbf{1}^T\left[(I - NYN^T) - \frac{1}{n}\mathbf{1}\mathbf{1}^T\right] \;\; = \;\; \mathbf{1}\mathbf{1}^T - \frac{1}{n}\mathbf{1}\mathbf{1}^T\mathbf{1}\mathbf{1}^T \; = \; 0. \tag{98}$$

The second inductive statement is similarly proved. ■

3.2 Model details

We will now apply the above techniques to our problem (83)-(85), where, as per our modeling assumption (III), b denote the (fixed) net supply vector, i.e. $b_i = P_i$ for a generator i, $b_i = -D_i$ for a demand node i, and $b_i = 0$ otherwise. Denoting by Y the diagonal matrix with entries $1/y_{ij}$, we have that given Y the unique power flows f and voltages θ are obtained by solving the system

$$\begin{aligned} N^T\theta - Y^{-1}f \;\; &= \;\; 0 \\ Nf \;\; &= \;\; b. \end{aligned}$$

Note that if we scale Y and b by a multiplicative factor $\mu > 0$ then we obtain an equivalent system, e.g. the power flows f increase by a factor of μ and the angles θ do not change. Thus, assuming $\Gamma \subseteq R^n_+$ is bounded (as is the case if we use (82)) then as a first step to solving (83)-(85) we can scale Γ so that condition (91) holds for every $y \in \Gamma$. Consequently, by (92), we can assume that there is a constant $r < 1$ such that $1 - \lambda_i < r$ for $2 \leq i \leq n$. In what follows we will always make this assumption.

By Lemma 3.2 each solution to (99)-(99) is of the form

$$\begin{aligned} &\theta = J^{-1}b + \delta\mathbf{1} \quad \text{ for some } \delta \in \mathcal{R}, & (99) \\ &f = YN^TJ^{-1}b. & (100) \end{aligned}$$

For each arc (i, j) denote by n_{ij} the column of N corresponding to (i, j), i.e., $n_{ij} := Ne_{ij}$, where $e_{ij} \in \mathcal{R}^m$ is the vector with a 1 at entry (i, j) and zero otherwise. Using (93) we therefore have

$$\begin{aligned} f_{ij} \;\; &= \;\; y_{ij}n^T_{ij}\left[I + P + P^2 + P^3 + \ldots\right]b, \quad \forall(i, j), \quad \text{and} & (101) \\ \theta_i - \theta_j \; = \; n^T_{ij}\theta \;\; &= \;\; n^T_{ij}\left[I + P + P^2 + P^3 + \ldots\right]b \; = \; n^T_{ij}\sum_{k=0}^{\infty} P^k \; b, & (102) \end{aligned}$$

In the following we will be handling expressions with infinite series such as the above. In order to facilitate the analysis we need a 'uniform convergence' argument, as follows. Given $y \in \Gamma$, note that we can write

$$P \, = \, P(y) \, = \, U(y)\Lambda(y)U(y)^T, \tag{103}$$

where $U(y)$ is a unitary matrix and $\Lambda(y)$ is the diagonal matrix containing the eigenvalues of $P(y)$. Hence, for any $k \geq 1$ and any arc (i,j) (and dropping the dependence on y_{st} for simplicity),

$$|n_{ij}^T P^k b| \, = \, |n_{ij}^T U \Lambda^k U^T b| \, < \, \nu^k, \tag{104}$$

for some $\nu < 1$, by (92). We will rely on this bound below.

As a first consequence of (104) we have the following result, showing that appropriate assumptions the continuous model we consider is related to the network vulnerability models in Section 2.

Lemma 3.4 *Let S be a set of arcs whose removal does not disconnect G. Suppose we fix the values $y_{ij} \, = \, 1/x_{ij}$ for each arc $(i,j) \notin S$, and we likewise set $y_{st} \, = \, \epsilon$ for each arc $(s,t) \in S$. Let $(f(y), \theta(y))$ denote the resulting power flow, and let $(\bar{f}, \bar{\theta})$ the solution to the power flow problem on $G - S$.*

Then

- *(a)* $\lim_{\epsilon \to 0} f_{st}(y) \, = \, 0$, *for all* $(s,t) \in S$,
- *(b) For any* $(u,v) \notin S$, $\lim_{\epsilon \to 0} f_{uv}(y) \, = \, \bar{f}_{uv}$.
- *(c) For any* (u,v), $\lim_{\epsilon \to 0}(\theta_u(y) - \theta_v(y)) \, = \, \bar{\theta}_u - \bar{\theta}_v$.

Proof. (a) Let $\tilde{G} = G - (s,t)$, let $\tilde{N}$ be node-arc incidence matrix of $\tilde{G}$, $\tilde{Y}$ the restriction of Y to $E - (s,t)$, and $\tilde{P} = I - \tilde{N}\tilde{Y}\tilde{N}^T - \frac{1}{n}\mathbf{1}\mathbf{1}^T$.

For any integer $k \geq 1$ we have by Lemma 3.3

$$\lim_{\epsilon \to 0} P^k = \lim_{\epsilon \to 0} (I - NYN^T)^k - \frac{1}{n}\mathbf{1}\mathbf{1}^T = (I - \tilde{N}\tilde{Y}\tilde{N}^T)^k - \frac{1}{n}\mathbf{1}\mathbf{1}^T = \tilde{P}^k. \tag{105}$$

Consequently, by (101), for any $(s,t) \in S$,

$$\lim_{\epsilon \to 0} f_{st} \;\; = \;\; \lim_{\epsilon \to 0} \left[y_{st} n_{st}^T \left(\sum_{k=0}^{\infty} P^k \right) b \right] \; = \; \sum_{k=0}^{\infty} \left[\lim_{\epsilon \to 0} y_{st} \left(n_{st}^T P^k b \right) \right] \; = \; 0, \tag{106}$$

where the exchange between summation and limit is valid because of (104). The proof of (b), (c) are similar.■

Lemma 3.4 can be interpreted as describing a particular type of attack that is feasible for the adversary under our models. Our computational experiments show that the pattern assumed by the Lemma is approximately correct: given an attack budget, the attacker tends to concentrate most of the attack on a small number of arcs (essentially, making their resistance very large), while at the same time attacking a larger number of lines with a small portion of the budget.

In the following set of results we determine efficient closed-form expressions for the gradient and Hessian of the objective in (81). As before, we denote by n_{ij} the column of the node-arc incidence matrix of the network corresponding to arc (i,j).

Lemma 3.5 *For any integer $k > 0$, and any arc (i,j)*

$$\begin{aligned} (a) \quad & \mathbf{1}^T P^k \;=\; 0, \\ (b) \quad & \frac{\partial}{\partial y_{ij}}\left[P^k b\right] \;=\; P\frac{\partial}{\partial y_{ij}}\left[P^{k-1}b\right] - n_{ij}n_{ij}^T P^{k-1}b. \end{aligned}$$

Proof. Note that $\mathbf{1}^T P = \mathbf{1}^T(I-J) = \mathbf{1}^T(I - NYN^T - \frac{1}{n}\mathbf{1}\mathbf{1}^T) = 0$. Hence $\mathbf{1}^T P^k = 0$.

$$\begin{aligned}
\frac{\partial}{\partial y_{ij}}\left[P^k b\right] &= \frac{\partial}{\partial y_{ij}}\left[PP^{k-1}b\right] \\
&= \frac{\partial}{\partial y_{ij}}\left[\left(I - \sum_{(u,v)\in E} y_{uv}\, n_{uv}n_{uv}^T - \frac{1}{n}\mathbf{1}\mathbf{1}^T\right)P^{k-1}b\right] \\
&= \frac{\partial}{\partial y_{ij}}\left[P^{k-1}b\right] - \frac{\partial}{\partial y_{ij}}\left[\left(\sum_{(u,v)\in E} y_{uv}\, n_{uv}n_{uv}^T\right)P^{k-1}b\right] - \frac{\partial}{\partial y_{ij}}\left[\frac{1}{n}\mathbf{1}\mathbf{1}^T P^{k-1}b\right] \\
&= \frac{\partial}{\partial y_{ij}}\left[P^{k-1}b\right] - \frac{\partial}{\partial y_{ij}}\left[\left(\sum_{(u,v)\in E} y_{uv}\, n_{uv}n_{uv}^T\right)P^{k-1}b\right] \\
&= \frac{\partial}{\partial y_{ij}}\left[P^{k-1}b\right] - \sum_{(u,v)\in E}\frac{\partial}{\partial y_{ij}}\left[y_{uv}\, n_{uv}n_{uv}^T P^{k-1}b\right] \\
&= \frac{\partial}{\partial y_{ij}}\left[P^{k-1}b\right] - \sum_{(u,v)\in E}\left[\frac{\partial y_{uv}}{\partial y_{ij}}\right]\; n_{uv}n_{uv}^T P^{k-1}b - \sum_{(u,v)\in E} y_{uv}\frac{\partial}{\partial y_{ij}}\left[n_{uv}n_{uv}^T P^{k-1}b\right] \\
&= \frac{\partial}{\partial y_{ij}}\left[P^{k-1}b\right] - n_{ij}n_{ij}^T P^{k-1}b - \sum_{(u,v)\in E} y_{uv}\, n_{uv}n_{uv}^T \frac{\partial}{\partial y_{ij}}\left[P^{k-1}b\right] \\
&= \left[I - \sum_{(u,v)\in E} y_{uv}\, n_{uv}n_{uv}^T\right]\frac{\partial}{\partial y_{ij}}\left[P^{k-1}b\right] - n_{ij}n_{ij}^T P^{k-1}b \\
&= \left[P + \frac{1}{n}\mathbf{1}\mathbf{1}^T\right]\frac{\partial}{\partial y_{ij}}\left[P^{k-1}b\right] - n_{ij}n_{ij}^T P^{k-1}b \\
&= P\frac{\partial}{\partial y_{ij}}\left[P^{k-1}b\right] - n_{ij}n_{ij}^T P^{k-1}b + \frac{\partial}{\partial y_{ij}}\left[\frac{1}{n}\mathbf{1}\mathbf{1}^T P^{k-1}b\right] \\
&= P\frac{\partial}{\partial y_{ij}}\left[P^{k-1}b\right] - n_{ij}n_{ij}^T P^{k-1}b.
\end{aligned}$$

where the third and the last equality follow from (a). ■

Using the above recursive formula we can write the following expressions:

$$\begin{aligned}
\frac{\partial}{\partial y_{ij}}[Pb] &= -n_{ij}n_{ij}^T b \\
\frac{\partial}{\partial y_{ij}}[P^2 b] &= P\frac{\partial}{\partial y_{ij}}[Pb] - n_{ij}n_{ij}^T Pb \\
\frac{\partial}{\partial y_{ij}}[P^3 b] &= P^2\frac{\partial}{\partial y_{ij}}[Pb] - Pn_{ij}n_{ij}^T Pb - n_{ij}n_{ij}^T P^2 b \\
\frac{\partial}{\partial y_{ij}}[P^4 b] &= P^3\frac{\partial}{\partial y_{ij}}[Pb] - P^2 n_{ij}n_{ij}^T Pb - Pn_{ij}n_{ij}^T P^2 b - n_{ij}n_{ij}^T P^3 b \\
&\;\;\vdots \\
\frac{\partial}{\partial y_{ij}}[P^k b] &= P^{k-1}\frac{\partial}{\partial y_{ij}}[Pb] - P^{k-2}n_{ij}n_{ij}^T Pb - P^{k-3}n_{ij}n_{ij}^T P^2 b - \ldots - n_{ij}n_{ij}^T P^{k-1}b
\end{aligned}$$

Consequently, defining

$$\tilde{\nabla}_{ij} \;=\; \frac{\partial}{\partial y_{ij}} \Big[I + P + P^2 + \ldots \Big]\, b, \tag{107}$$

we have

$$\begin{aligned}
\tilde{\nabla}_{ij} &= \Big[I + P + P^2 + \ldots \Big] \frac{\partial}{\partial y_{ij}} [Pb] \;-\; \Big(I + P + P^2 + \ldots \Big)\, n_{ij} n_{ij}^T \,\Big(P + P^2 + P^3 + \ldots \Big)\, b \\
&= -\Big[I + P + P^2 + \ldots \Big]\, n_{ij} n_{ij}^T b \;-\; \Big(I + P + P^2 + \ldots \Big)\, n_{ij} n_{ij}^T \,\Big(I + P + P^2 + \ldots - I \Big)\, b \\
&= -\Big(I + P + P^2 + \ldots \Big)\, n_{ij} n_{ij}^T \,\Big(I + P + P^2 + \ldots \Big)\, b \\
&= -J^{-1}\, n_{ij} n_{ij}^T \,\theta,
\end{aligned} \tag{108}$$

where the last equality follows from (99) and (93), and the fact that $n_{ij}^T \mathbf{1} = 0$.

Using (101), the gradient of function $f_{uv}(y)$ with respect to the variables y_{ij} can be written as:

$$\frac{\partial f_{uv}}{\partial y_{ij}} \;=\; y_{uv}\, n_{uv}^T \frac{\partial}{\partial y_{ij}} \Big[I + P + P^2 + P^3 + \ldots \Big]\, b \;=\; y_{uv}\, n_{uv}^T \,\tilde{\nabla}_{ij}, \qquad (i,j) \neq (u,v) \tag{109}$$

$$\begin{aligned}
\frac{\partial f_{ij}}{\partial y_{ij}} &= n_{ij}^T \Big[I + P + P^2 + P^3 + \ldots \Big]\, b \;+\; y_{ij}\, n_{ij}^T \frac{\partial}{\partial y_{ij}} \Big[I + P + P^2 + P^3 + \ldots \Big]\, b \\
&= n_{ij}^T \tilde{\nabla}_{ij} \;+\; y_{ij} n_{ij}^T \tilde{\nabla}_{ij}.
\end{aligned} \tag{110}$$

We similarly develop close-form expressions for the second order derivatives. For $(u,v) \neq (i,j), (u,v) \neq (h,k)$, we have the following :

$$\begin{aligned}
\frac{\partial^2 f_{uv}}{\partial y_{ij} \partial y_{hk}} &= y_{uv} n_{uv}^T \,[\, (I + P + P^2 + P^3 + \ldots)\, n_{ij} n_{ij}^T \,(I + P + P^2 + P^3 + \ldots) n_{hk} n_{hk}^T \\
&\qquad + (I + P + P^2 + P^3 + \ldots)\, n_{hk} n_{hk}^T \,(I + P + P^2 + P^3 + \ldots)\, n_{ij} n_{ij}^T \,\theta \\
&= -y_{uv} n_{uv}^T J^{-1} \Big[n_{ij} n_{ij}^T \tilde{\nabla}_{hk} + n_{hk} n_{hk}^T \tilde{\nabla}_{ij} \Big].
\end{aligned} \tag{111}$$

Similarly, the remaining terms are:

$$\frac{\partial^2 f_{uv}}{\partial y_{uv}^2} \;=\; 2\, n_{uv}^T \tilde{\nabla}_{uv} \;-\; 2\, y_{uv} n_{uv}^T J^{-1} n_{uv} n_{uv}^T \tilde{\nabla}_{uv}, \tag{112}$$

$$\frac{\partial^2 f_{uv}}{\partial y_{uv} \partial y_{ij}} \;=\; n_{uv}^T \,\tilde{\nabla}_{ij} \;-\; y_{uv}\, n_{uv}^T J^{-1} \Big[n_{ij} n_{ij}^T \,\tilde{\nabla}_{i} \;+\; n_{uv} n_{uv}^T \,\tilde{\nabla}_{ij} \Big] \tag{113}$$

3.3 Implementation details

We use LOQO [22] to solve problem (83)-(85), using $\Gamma \;=\; \Big\{ y \geq 0 \;:\; \sum_{ij} \frac{1}{y_{ij}} \;\leq\; B \Big\}$ with values of B that we selected. LOQO is an infeasible primal-dual, interior-point method applied to a sequence of quadratic approximations to the given problem. The procedure stops if at any iteration the primal and dual problems are feasible and with objective values that are *close* to each other, in which case a local optimal solution is found. For numerical reasons, LOQO additionally uses an upper bound on the overall number of iterations to perform.

At each iteration of the method applied by LOQO, it requires the Hessian and gradient of the objective function and the constraints. The latter are easy to derive. Note that using (109), (110), (111)-(113) one can obtain compact, closed-form expressions for the Hessian and gradient of the objective. This approach requires the computation of quantities $n_{uv}^T J^{-1} n_{ij}$ for each pair of arcs (i,j), (u,v). At any given iteration, we compute and (appropriately) store these quantities (which can be done in $O(n^2 + nm)$ space).

In order to compute $n_{uv}^T J^{-1} n_{ij}$, for given (i,j) and (u,v), we simply solve the sparse linear system on variables κ, λ:

$$\begin{aligned} N^T\kappa - Y^{-1}\lambda &= 0 && (114)\\ N\lambda &= n_{ij}. && (115) \end{aligned}$$

As in (99), we have $\kappa = J^{-1}n_{ij} + \delta\mathbf{1}$ for some real δ. But then $n_{uv}^T\kappa = n_{uv}^T J^{-1} n_{ij}$, the desired quantity. In order to solve (114)-(115) we use Cplex (to solve a nominal linear program).

We point out that, alternatively, LOQO can perform symbolic differentiation in order to directly compute the Hessian and gradient. We could in principle follow this approach in order to solve a problem with objective (83), constraints (84), (85) *and* (1), (2). We prefer our approach because it employs fewer variables (we do not need the flow variables or the angles) and primal feasibility is far simpler.

In our implementation, we fix a value for the iteration limit, but apply additional stopping criteria:

(1) If both primal and dual are feasible, we consider the relative error between the primal and dual values, $\epsilon = \frac{\text{PV - DV}}{\text{DV}}$, where 'PV' and 'DV' refer to primal and dual values respectively. If the relative error ϵ is less than some desired threshold we stop, and report the solution as "ϵ-locally-optimal."

(2) If on the other hand we reach the iteration limit without a stopping as in [(1)], then we consider the last iteration at which we had both primal and dual feasible solutions. If such an iteration exists, then we report the corresponding configuration of resistances along with the associated congestion value. If such an iteration does not exist, then the report the run as unsuccessful.

Finally, we provide to LOQO the starting point $x_{ij} = x_{ij}^L$ for each arc (i,j).

3.4 Computational testing

We applied our algorithm to a number of test cases, using three constraint sets Γ as in (82):

(1) $\mathbf{\Gamma(1)}$, where for all (i,j), $x_{ij}^L = 1$ and $x_{ij}^U = 5$,

(2) $\mathbf{\Gamma(2)}$, where for all (i,j), $x_{ij}^L = 1$ and $x_{ij}^U = 10$,

(3) $\mathbf{\Gamma(3)}$, where for all (i,j), $x_{ij}^L = 1$ and $x_{ij}^U = 20$.

In each case, we set $B = \sum_{(i,j)} x_{ij}^L + \Delta\text{B}$, where ΔB represents an "excess budget".

We used data sets derived from the IEEE test cases [16], and, in addition, we used the following procedure to generate larger, examples. Let $\mathcal{N}^1$ and $\mathcal{N}^2$ be two power networks. We create a new network, $\mathcal{N}^3$, by taking a copy of $\mathcal{N}^1$ and a (disjoint) copy of $\mathcal{N}^2$, and adding a random set of arcs between the two copies. Each potential arc between the two copies is added with a given probability $0 < p < 1$; furthermore, the resistance and capacity of this arc are chosen equal to the (corresponding) average (among all arcs in $\mathcal{N}^1$ and $\mathcal{N}^2$), plus a small random perturbation.

In the tables below, we state the iteration limit and the ϵ parameter used to control termination of LOQO. For each run, we state the objective value (i.e. the maximum arc congestion as in (80)) at termination, the corresponding run time and number of iterations used, and the termination status. This is indicated by "Exit Status", with the following interpretation:

(1) '**ϵ-L-opt.**': the algorithm computed an ϵ-locally-optimal solution.

(2) '**PDfeas, Iter: lastItn**': the algorithm reached the iteration limit without finding an ϵ-locally-optimal solution, but there was an iteration at which both primal and dual problems were feasible. 'lastItn' gives the last iteration at which both primal and dual solutions were feasible.

(3) '**opt.**': the algorithm attained LOQO's internal optimality tolerance.

Tables 5 and 6 summarize the results for the 49-node, 84-arc network, with 14 demand nodes and 4 generators that we considered in section 2.4, using sets $\Gamma(1)$ and $\Gamma(2)$ respectively. Note that for example in the case $\Delta\mathrm{B} = 30$, under $\Gamma(2)$, the attacker can increase (from their minimum value) the resistance of up to 3 arcs by a factor of 10 (with 3 units of budget left over). And under $\Gamma(1)$, up to 6 arcs can have their resistance increased by a factor of 5. In either case we have a situation reminiscent of the $N - k$ problem, with small k.

Table 5: ***49 nodes, 84 arcs,*** $\mathbf{\Gamma(1)}$
Iteration Limit: 800, $\boldsymbol{\epsilon = 0.01}$

	ΔB					
	5	**10**	**15**	**20**	**25**	**30**
Max Cong	0.673054	0.750547	0.815623	0.865806	0.901453	0.951803
Time (sec)	12	15	18	19	28	22
Iterations	258	347	430	461	Limit	492
Exit Status	ϵ-L-opt.	ϵ-L-opt.	ϵ-L-opt.	ϵ-L-opt.	PDfeas Iter: 613	ϵ-L-opt.

Table 6: ***49 nodes, 84 arcs,*** $\mathbf{\Gamma(2)}$
Iteration Limit: 800, $\boldsymbol{\epsilon = 0.01}$

	ΔB					
	5	**10**	**15**	**20**	**25**	**30**
Max Cong	0.67306	0.751673	0.815584	0.8685	0.91523	0.9496
Time (sec)	9	13	34	3	29	30
Iterations	177	295	Limit	Limit	Limit	Limit
Exit Status	ϵ-L-opt.	ϵ-L-opt.	PDfeas Iter: 800	PDfeas Iter: 738	PDfeas Iter: 624	PDfeas Iter: 656

Table 7 presents similar results for a network with 300 nodes, 409 arcs, 42 generators and 172 loads. Note that for the runs $\Delta\mathrm{B} \geq 20$ the maximum load value is identical; the optimal solutions x_{ij} were nearly identical, independent of the initial point given to LOQO.

Table 8 presents similar results for a network with 600 nodes, 990 arcs, 344 demand nodes and 98 generators, under set $\Gamma(2)$. We observed an interesting issue in the case where $\Delta\mathrm{B} = 10$. Here, LOQO terminated with a solution in which for some arc (i, j), both $p_{ij} > 0$ and $q_{ij} > 0$ (refer to formulation (83)-(85). The value in parenthesis indicates the true value of the congestion obtained by solving the network controller's problem if we were to use the resistance values (x_{ij}) given by LOQO.

Finally, Table 9 presents experiments on a network with 649 nodes and 1368 arcs. Here, exit status 'DF' means that dual feasibility was achieved, but not primal feasibility.

Table 7: ***300 nodes, 409 arcs,* $\mathbf{\Gamma(2)}$**
Iteration Limit: 500, $\epsilon = 0.01$

	ΔB			
	9	**18**	**27**	**36**
Max Cong	0.590690	0.694101	0.771165	0.771165
Time (sec)	208	1248	981	825
Iterations	91	Limit	406	320
Exit Status	opt.	PDfeas Iter: 318	opt.	opt

Table 8: ***600 nodes, 990 arcs,* $\mathbf{\Gamma(2)}$**
Iteration Limit: 300, $\epsilon = 0.01$

	ΔB				
	10	**20**	**27**	**36**	**40**
Max Cong	0.082735 (0.571562)	1.076251	1.156187	1.088491	1.161887
Time (sec)	11848	7500	4502	11251	7800
Iterations	Limit	210	114	300	208
Exit Status	PDfeas Iter: 300	ϵ-L-opt.	ϵ-L-opt.	PDfeas Iter: 300	ϵ-L-opt.

3.4.1 Distribution of attack weights

Table 10 describe the distribution of x_{ij} values at termination of the algorithm, for a number of networks and attack budgets. For each test we show first (in parentheses) the number of nodes and arcs, followed by the the attack budget and constraint set. The data for each test shows, for each range of resistance values, the number of arcs whose resistance falls in that range.
Note that in each test case the adversary can increase up to three resistances to their maximum value. In all three cases many resistances take relatively small values and a small number of arcs have high resistance. Recall that for set $\Gamma(2)$ we always have $x_{ij}^{max} = 10$, thus in the case of the $(300, 409)$ network exactly three arcs are in the top range, while for the $(600, 990)$ network two are in the top range and one more has relatively high resistance. In the case of the small network the distribution seems more "continuous", although as we will see in the next section that the three highest resistance arcs play a significant role.

3.4.2 Comparison with the minimum-cardinality attack model

In this section we describe some comparisons with the mixed-integer programming model considered in Section 2.1. A direct comparison on case-by-case basis is not possible, because the nonlinear model assumes that all demands and generator outputs are fixed, whereas the model in Section 2.1 in particular allows load-shedding with a minimum desired throughput of D^{min} – we could set up problem instances where $D^{min} = 1.0$ but in that case an attack that disconnects a demand node, even one with tiny demand, would be considered a success for the attacker.

To deal with these issues and still obtain a meaningful comparison, we set an example with 49

Table 9: ***649 nodes, 1368 arcs,*** $\mathbf{\Gamma(2)}$
Iteration Limit: 500, $\boldsymbol{\epsilon = 0.01}$

	ΔB			
	20	**30**	**40**	**60**
Max Cong	(0.06732) 1.29567	1.942652	(0.049348)1.395398	2.045111
Time (sec)	66420	36274	54070	40262
Iterations	Limit	374	Limit	Limit
Exit Status	DF	ϵ-L-opt.	DF	PDfeas Iter: 491

Table 10: **Solution histogram**

$\mathbf{(49, 90)}$ $\mathbf{\Delta B = 57, \Gamma(3)}$		$\mathbf{(300, 409)}$ $\mathbf{\Delta B = 27, \Gamma(2)}$		$\mathbf{(600, 990)}$ $\mathbf{\Delta B = 36, \Gamma(2)}$	
Range	**Count**	**Range**	**Count**	**Range**	**Count**
[1, 1]	8	[1, 1]	1	[1, 1]	14
(1,2]	72	(1, 2]	405	(1, 2]	970
(2,3]	4	(2, 9]	0	(2, 5]	3
(5,6]	1	(9, 10]	3	(5, 6]	0
(6,7]	1			(6, 7]	1
(7,8]	4			(7, 9]	0
(8,20]	0			(9, 10]	2

nodes and 88, and an example with 49 nodes and 90 arcs, in which no demand or generator node can be disconnected from the rest by removing up to three arcs. In each case there are 4 generators and 14 demand nodes. By scaling up all capacities by a common constant we then obtained a family of problems.

In terms of the mixed-integer programming model, we then set-up a one-configuration problem with $D^{min} = 1$, with the goal of investigating its vulnerability should up to three arcs be removed. Here we remind the reader that the algorithms 2.1 seek a minimum-cardinality attack that defeat the controller, and not the most severe attack of a given cardinality. Once our problem is solved the optimal attack is certified to be successful (and of minimum-cardinality), but not necessarily *the* most severe attack of *that* cardinality. Nevertheless, by adjusting our formulation (49)-(53) we can search for *a* successful attack of any given cardinality, if it exists. The problem we obtain is:

$$t^* \quad = \quad \max t \tag{116}$$

$$\text{Subject to:} \qquad \sum_{(i,j)} z_{ij} \quad \leq \quad k, \tag{117}$$

$$w_{\mathcal{C}}^T \psi^{\mathcal{C}} - t \quad \geq \quad 0, \quad \forall \, \mathcal{C} \subseteq \mathcal{G}, \tag{118}$$

$$A\psi^{\mathcal{C}} + Bz \quad \leq \quad b + B \quad \forall \, \mathcal{C} \subseteq \mathcal{G}, \tag{119}$$

$$z_{ij} \quad = \quad 0 \text{ or } 1, \ \forall \, (i,j). \tag{120}$$

where k $(= 3)$ is a the number of arcs that the attacker can be remove. However, all this formulation guarantees is that $t^* > 1$ if and only if a successful attack of cardinality $\leq k$ exists – because of the nature of our formulation, when $t^* > 1$ then t^* will be an approximation (in general, close) to the highest severity. A final detail is that since 3 lines will not disconnect the demands from the generators, the “severity” of an attack as per formulation (117)-(120) is the maximum arc congestion post-attack; thus putting the problem on a common ground with the nonlinear models we consider.

For our experiments we used $\Gamma(1)$ (which allows resistances to increase by up to a factor of 20) with an excess budget of 60, on the network with 49 nodes, 90 arcs, 4 generators and 14 demand nodes. Note that the parameters allow the attacker to concentrate the budget on three arcs.

Table 11 contains the results. Each row corresponds to a different experiment, where a common multiplier used to scale up all capacities. We used, respectively, $1.0, 1.2, 1.4, 1.6, 1.8, 2.0$, thus obtaining examples that are progressively more difficult to interdict.

In the 'MIP' section, the column headed 'Cong' indicates the congestion (max. arc overload) in the network obtained by removing the arcs produced by the mixed-integer programming model, and the column headed 'ATTACK' indicates which arcs were removed by the MIP.

In the 'NONLINEAR' section, 'Cong' indicates the maximum congestion resulting from the increase in resistances computed by the model. We also list the 6 arcs with highest resistances (and the resistance values).

The column headed 'Impact' indicates the maximum congestion obtained by deleting the three arcs with maximum resistance (as computed by the model), while leaving all other resistances unchanged. We also performed the following test: we removed the top three highest resistance arcs, while keeping all other resistances unchanged, but now we allow the controller to reduce total demand by up to 10% with the objective of minimizing the maximum congestion; the resulting congestion value is shown in the column labeled 'I-10%'.

Finally, using all resistance values as computed by the nonlinear model, and without removing any arcs, we allowed the controller to reduce total demand by up to 10%, again with the objective of minimizing the maximum congestion. The column labeled 'C-10%' shows the resulting congestion value.

Table 11: **Comparison between models**

MIP		**NONLINEAR**					
Cong	**Attack**	**Cong**	**Top 6 Arcs**	**Impact**	**I-10%**	**C-10%**
1.440880	29,32,45	2.149673	29(7.79), 27(7.20), 41(7.03), 67(7.02), 54(6.72), 79(5.71)	1.717584	1.334536	1.671452
1.431320	27,29,41	1.786874	29(8.28), 27(7.72), 41(7.32), 67(7.19), 54(6.92), 79(5.78)	1.431320	1.112113	1.386416
1.226846	27,29,41	1.556341	29(8.31), 27(7.74), 41(7.53), 67(7.48), 54(7.18), 79(6.15)	1.226846	0.953240	1.213288
1.073490	27,29,41	1.359954	29(8.18), 27(7.58), 41(7.53), 67(7.58), 54(7.22), 79(6.25)	1.073490	0.834085	1.054584
0.692488	18,57,60	1.202712	29(8.43), 27(7.90), 41(7.53), 67(7.48), 54(7.18), 79(6.12)	0.954213	0.741409	0.935953
0.686301	20,89,45	1.077328	29(7.87), 27(7.29), 41(7.04), 67(7.01), 54(6.70), 79(5.63)	0.858792	0.667268	0.838777

Note that the results in Table 11 show some significant, overlap between the results from the two models. As before, we see that the solutions to the nonlinear model tend to concentrate the attack on a relatively small number of lines, while at the same time investing small portions of the attack budget on other lines.

Moreover, the two models are consistent: the severity of the attack, for both models, decrease as the scale increases (as one should expect). Finally, the top three highest resistance arcs selected by the nonlinear model have significant impact from the point of view of the min-cardinality model.

References

[1] G. Andersson, *Modelling and Analysis of Electric Power Systems.* Lecture 227-0526-00, Power Systems Laboratory, ETH Zürich, March 2004. Download from http://www.eeh.ee.ethz.ch/downloads/academics/courses/227-0526-00.pdf.

[2] R. Alvarez, Interdicting Electric Power Grids, Masters' Thesis, U.S. Naval Postgraduate School, 2004.

[3] J. Arroyo and F. Galiana, On the Solution of the Bilevel Programming Formulation of the Terrorist Threat Problem, *IEEE Trans. Power Systems*, Vol. 20 (2005), 789–797.

[4] H. Y. Benson, D. F. Shanno and R. J. Vanderbei, Interior-point methods for nonconvex nonlinear programming: jamming and comparative numerical testing, *Math. Programming* **99**, 35 – 38 (2004).

[5] D. Bienstock and S. Mattia, Using mixed-integer programming to solve power grid blackout problems , *Discrete Optimization* **4** (2007), 115–141.

[6] V.M. Bier, E.R. Gratz, N.J. Haphuriwat, W. Magua, K.R. Wierzbickiby, Methodology for identifying near-optimal interdiction strategies for a power transmission system, Reliability Engineering and System Safety **92** (2007), 1155–1161.

[7] S. Boyd, Convex Optimization of Graph Laplacian Eigenvalues, *Proc. International Congress of Mathematicians* **3** (2006), 1311–1319.

[8] D. Braess, Über ein Paradox der Verkerhsplannung, Unternehmenstorchung Vol. 12 (1968) 258–268.

[9] B.A. Carreras, V.E. Lynch, I. Dobson, D.E. Newman, Critical points and transitions in an electric power transmission model for cascading failure blackouts, Chaos, vol. 12, no. 4, 2002, 985-994.

[10] B.A. Carreras, V.E. Lynch, D.E. Newman, I. Dobson, Blackout mitigation assessment in power transmission systems, 36th Hawaii International Conference on System Sciences, Hawaii, 2003.

[11] B.A. Carreras, V.E. Lynch, I. Dobson, D.E. Newman, Complex dynamics of blackouts in power transmission systems, Chaos, vol. 14, no. 3, September 2004, 643-652.

[12] B.A. Carreras, D.E. Newman, I. Dobson, A.B. Poole, Evidence for self organized criticality in electric power system blackouts, IEEE Transactions on Circuits and Systems I, vol. 51, no. 9, Sept. 2004, 1733- 1740.

[13] ILOG CPLEX 11.0. ILOG, Inc., Incline Village, NV.

[14] S.T. DeNegre and T.K Ralphs, A Branch-and-cut Algorithm for Integer Bilevel Linear Programs, COR@L Technical Report, Lehigh University (2008).

[15] R. Fletcher, N. I. M. Gould, S. Leyffer, Ph. L. Toint, and A. Wächter, Global convergence of trust-region SQP-filter algorithms for general nonlinear programming, *SIAM J. Optimization* **13**, 635–659 (2002).

[16] The IEEE reliability test system–1996, IEEE Trans. Power Syst., vol. 14 (1999) 1010 - 1020.

[17] U. Janjarassuk and J. T. Linderoth, Reformulation and Sampling to Solve a Stochastic Network Interdiction Problem, to appear, *Networks* (2008).

[18] C. Lim and J.C. Smith, Algorithms for Discrete and Continuous Multicommodity Flow Network Interdiction Problems, *IIE Transactions* **39**, 15-26, 2007.

[19] B. Mohar, The Laplacian spectrum of graphs, in: Y. Alavi, G. Chartrand, O. Oellermann, A. Schwenk (Eds.), *Graph Theory, Combinatorics, and Applications*, London Math. Soc. Lecture Notes, Wiley-Interscience, 871-898 (1991).

[20] A. Pinar, J. Meza, V. Donde, and B. Lesieutre, Optimization Strategies for the Vulnerability Analysis of the Power Grid, submitted to *SIAM Journal on Optimization* (2007).

[21] J. Salmeron, K. Wood and R. Baldick, Analysis of Electric Grid Security Under Terrorist Threat, *IEEE Trans. Power Systems* **19** (2004), 905–912.

[22] Vanderbei, R. 1997. LOQO User's manual, *Statistics and Operations Research* Technical report No SOR-97-08, Princeton University.

[23] *Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations*, U.S.-Canada Power System Outage Task Force, April 5, 2004. Download from: https://reports.energy.gov.