Controlling Cascading Failures with Cooperative Autonomous Agents

Paul Hines

5 January 2005

Abstract

Cascading failures in electricity networks cause blackouts, which often lead to severe economic and social consequences. Cascading failures are typically initiated by a set of equipment outages that cause operating constraint violations. When violations persist in a network they can trigger additional outages which in turn may cause further violations. This paper proposes a method for limiting the social costs of cascading failures by eliminating violations before a dependent outage occurs. This global problem is solved using a new application of distributed model predictive control. Specifically, our method is to create a network of autonomous agents, one at each bus of a power network. The task assigned to each agent is to solve the global control problem with limited communication abilities. Each agent builds a simplified model of the network based on locally available data and solves its local problem using model predictive control and cooperation. Through extensive simulations with IEEE test networks, we find that the autonomous agent design meets its goals with limited communication. Experiments also demonstrate that cooperation among software agents can vastly improve system performance.

While the principle contribution of this paper is the development of a new method for controlling cascading failures, several aspects of the included results are also relevant to contemporary policy problems. Firstly, this paper demonstrates that it is possible to perform some network control tasks without large-scale centralization. This property could be valuable in the US where centralization of control and regulatory functions has proved politically difficult. Secondly, this paper presents preliminary estimates of the benefits, costs, and risks associated with this technology. With some additional development, the methods will be useful for evaluating and comparing grid control technologies.

1 Introduction

In 1895 the Niagara Falls Power Company energized the first high capacity 3-phase transmission line, connecting hydroelectric generators at Niagara Falls and consumers 22 miles away in Buffalo, NY. The line operated at 11kV and carried power to customers including the Pittsburgh Reduction Company (now Alcoa) and the Buffalo street-car system. While the new system succeeded in carrying power from Niagara to Buffalo, it proved to be unreliable. Lightning frequently caused faults that damaged equipment and interrupted service (Neil, 1942). Numerous approaches were tried to combat this problem. High powered fuses and eventually circuit breaker/relay systems were installed to interrupt excessive line currents. Parallel transmission lines were added creating redundancy. Eventually, distant portions of the network were interconnected, synchronizing hundreds of large generators.

The emergent system has several important properties. It is able to transmit power over relatively long distances. It can suffer minor disturbances (such as lightning strikes) without sustaining large amounts of equipment damage. If operated correctly, it can endure small outages without significantly disrupting service. And finally, it is susceptible to cascading failures[1] that can result in large blackouts.[2]

1.1 Cascading failures

On November 9, 1965 the Northeastern United States suffered a cascading failure that interrupted service to 30 million customers. A faulty relay setting on a line between Niagara and Toronto tripped. The power was shifted to three parallel lines, which quickly became overloaded, triggering subsequent relay actions. Excess Niagara generation was instantaneously sent south into New York state, overloading additional lines, and eventually resulting in a cascading failure that affected customers in seven states and much of Ontario (Vassell, 1991). If the initial overload on the three remaining Toronto-Niagara lines had been quickly eliminated, the consequences would have been greatly reduced.

It is often difficult to understand the root causes of a cascading failure, but some general properties are known. (Talukdar, 2003) shows that the probability of large blackouts has a power-law

---

[1] A cascading failure is a series of equipment outages, such that an initial disturbance causes one or more dependent equipment outages. Cascading failures can be thought of as state transitions in a hybrid system (Hines, 2004, Antsaklis, 2000).

[2] A blackout is the interruption of electricity service to customers in the network.

tail. Systems that have power-law probability distributions can have very high or even infinite expected consequences.

Others have noted that the probability of a cascading failure increases as transmission system loading increases, and that this probability goes through a sharp phase transition (Dobson, 2004; Liao, 2004). It is also appears that cascading failures are propagated by relays acting in response to operating constraint violations, which often persist for some time before triggering a relay response. While the 1996 western US blackout progressed fairly quickly,[3] the system endured overloads on the western transmission corridor for 22 seconds after the initial disturbance, before a rapid sequence of relay actions commenced (WSCC, 1996).

The consequences of blackouts can be quite severe (see Table 1.1). Because many services, such as stairwell lighting and traffic lights, frequently do not have a source of backup energy, blackouts can have both economic and human health consequences.

Table 1.1—Several large cascading failures (NERC, 2005)

Date	Location	Notable consequences
9 Nov. 1965	Northeastern US, Ontario	30,000,000 customers (20,000 MW) interrupted.
13 July 1977	New York City	9,000,000 customers (6,000 MW) interrupted. Widespread looting and chaos. Police made about 4000 arrests. (Wikipedia, 2005)
2 July 1996	Western US	2,000,000 customers (11,850 MW) interrupted.
3 July 1996	Western US	The disturbance from July 2 reoccurred. Operators interrupted load to most of Boise Idaho, vastly reducing the extent of the cascading failure. (WSCC, 1996)
10 Aug. 1996	Western US	7,500,000 customers (28,000 MW) interrupted. Economic damage estimates range from $1-$3 billion.
25 June 1998	Midwestern US, central Canada	152,000 customers (950 MW) interrupted.
Nov. 1988 to June 2003	Western India	29 large cascading failures over 15 years—1.9 per year. Millions of customers interrupted in most cases. (Roy, 2004)
14 Aug. 2003	Midwestern and Northeastern US, Southeastern Canada	50,000,000 customers interrupted. Estimates of the social costs range from $6 billion (Graves, 2003) to $10 billion (ICF, 2003). Massive traffic jams in New York City.
27 Sept. 2003	Italy	57,000,000 customers interrupted. At least 5 deaths resulted. 30,000 passengers stranded in trains for hours. (BBC, 2003, CNN, 2003)

Historically, cascading failures have opened windows for significant changes in power system regulation and technology. The 1965 blackout led to the creation of the North American Electric

---

[3] Within 1 minute the western grid had separated into 5 islands (WSCC, 1996).

Reliability Council (NERC), the industry's means of self regulating for reliability. As a result of the 1977 event, engineers developed, and NERC adopted, a set of operating states and objectives that remain the primary standard for power system operation.[4] This led to widespread adoption of the "N-1" reliability criteria[5] that most North American system operators (SOs) use to manage the cascading failure risk under normal operating conditions. In the wake of the 2003 blackouts many in industry, government, and academia are advocating that the current practice of self-regulation be replaced with a set of binding, enforceable reliability rules.[6]

1.2 Operating power networks

Power systems are operated with many objectives, including:

- *Economics*–maximize the net economic benefit of service.
- *Reliability*–minimize the risk of service interruption.
- *Protection*–minimize the risk of infrastructure damage.

Sometimes these objectives are commensurate, but often they conflict. For example, in a lightning initiated fault on a transmission line, a relay that trips to clear the fault and quickly restores the line to service effectively manages both its reliability and protection objectives. Because the objectives are commensurate it is trivial to manage both simultaneously. During a cascading failure reliability and protection are brought into conflict. Violations such as a transmission line overload cause relays designed for protection to trip, thereby propagating the cascade through the network. During a cascading failure, power systems generally do a poor job of balancing conflicting objectives. This paper proposes to solve this problem by improving the network's ability to react to violations.[7]

System protection measures or special protection schemes (SPS) are control methods designed to preserve the integrity of the network as a whole during an emergency operating condition. According to Anderson (1996) an SPS is a method "that is designed to detect a particular system

---

[4] This method classifies any state as normal, alert, emergency, *in extremis,* or restorative and recommends actions that are appropriate to take in each condition (Fink, 1978). In the normal state a system is to be considered "secure" if no single contingency can cause a cascading failure. A single contingency is the outage of a single element of the network such as a generator or transmission line. A double contingency is the removal of two elements.

[5] The "N-1" reliability criteria, in short, requires that a system be operated such that no single contingency will affect a cascading failure.

[6] Several bills currently in congress (H.R. 6, H.R. 3004, and S. 2236) would facilitate the creation of mandatory reliability rules (see http://www.nerc.com/about/faq.html).

[7] This coincides with recommendation 21 from the August 14, 2003 blackout report, in which the authors recommend that US system operators, "Make more effective and wider use of system protection measures." (US-CA, 2004).

condition that is known to cause unusual stress to the power system, and to take some type of predetermined action to counteract the observed condition in a controlled manner." SPS come in many varieties, but as almost all are preprogrammed to react to very specific circumstances with predetermined control actions. Typically SPS are designed by performing off-line network studies and pre-determining control rules that tend to alleviate a set of potential problems. Newer designs are able to adapt control actions to changing network conditions, but still rely on pre-determined rules (Rehtanz, 2001 ;Novosel, 2004; Madani, 2004). While almost all SPS designs currently in operation are operated out of a centrally located control center (Anderson, 1996), a few SPS design concepts use a more distributed architecture, though agents are generally organized hierarchically and are dependent on central facilities for planning activities (Jung, 2001, 2002; Kamwa, 2001). No existing SPS designs operate solely using distributed autonomous agents.

1.3   Distributed control and multi-agent systems

Power networks are operated by thousands of agents. In the US eastern interconnect there are approximately 100 control areas and about 50,000 buses controlled by hundreds of human, and thousands of mechanical agents. Due to the complexity of power networks, real time control of the entire network from a central location is impossible. Even if doing so were computationally feasible, the system would be highly vulnerable to random failures, organized attacks, and communication problems. For this reason, the control of power networks, as with many complex systems, has been distributed to many autonomous controllers. The vast majority of existing mechanical controllers operate with only local information and follow very simple rules. As communication and computation technologies advance, it is increasingly possible to design distributed agents capable of solving complex network problems.

But, agent-based systems are not without disadvantages. Heterogeneous, distributed agents can be uncoordinated and parochial. To the extent that an agent is autonomous, it can act on its own volition and conflict with other agents. Because a distributed agent generally works with incomplete information, it can, at best, make locally correct decisions, which can be globally wrong. This is the general challenge of designing autonomous agent networks: to design the agents such that locally correct decisions are simultaneously globally correct.

Methods for solving complex problems using distributed software agents are increasingly prevalent in the literature. Fisher (1999) describes an autonomous agent problem decomposition as an emergent algorithm and outlines a strategy for developing agent-based solutions. Camponogara (2000) provides a method of decomposing optimization problems for collaborative agent networks,

provides conditions under which optimal performance can be guaranteed, and demonstrates that these conditions can be relaxed for some applications. Others have shown that distributed optimization methods (Cohen, 1984) can be applied to the optimal power flow (OPF) problem and solved by distributed autonomous agents (Kim, 2000). Attempts to reproduce this method for our application indicate that the method is unreliable and approaches an optimum very slowly if at all (Hines, 2004). Another distributed optimization technique (Modi, 2004) organizes agents hierarchically to solve discrete optimization problems. Agent-based technologies have also been applied to the relay protection problem (Yanxia, 2002; Coury, 2002) and proposed as a means of improving distribution systems (Kueck, 2003).

1.4 Cooperation

We define cooperation as the sharing of useful information and the utilization of commensurate goals. In many applications, as long as communication and calculation costs are negligible, skillful cooperative agents will perform at least as well as agents acting independently or competitively. For example in the prisoner's dilemma game, prisoners who decide *ex ante* to cooperate in concert will likely fare better, and certainly no worse, than prisoners acting independently. Recently engineers and computer scientists have found that cooperation can be a useful technology for software-based systems. Jennings (1999, 2003) discusses cooperative designs for an Energy Management System (EMS)[8], a particle accelerator, and cement factory control. These papers advocate that agents having clearly defined and known intentions and responsibilities. Camponogara (2002) demonstrates that cooperative agents working to control the frequency of a power system can outperform agents acting independently. Cooperation can cause problems as agents must process additional information. This can lead to unbounded problem growth when not properly designed (Durfee, 1999).

1.5 Distributed model predictive control (DMPC)

The autonomous agent network that we use in this paper combines distributed control (spatial problem decomposition) with a method for temporal decomposition called model predictive control (MPC). MPC is a repetitive procedure that combines the advantages of long-term planning (feed-forward control based on performance predictions over an extended horizon) with the advantages of reactive control (feedback using measurements of actual performance). At the beginning of each repetition, the state of the system to be controlled is measured. A time-horizon, stretching into the

---

[8] EMS is the term used for the system control and data communication system that operators use in a control room.

future, is divided into intervals. Models are adopted to predict the effects of control actions on system-states in these intervals. The predictions are used to plan optimal actions for each interval, but only the actions for the first interval are implemented. When this interval ends, the procedure is repeated.

MPC, because it uses optimization for making decisions, readily accommodates large numbers of complex constraints. Many other control techniques do not allow inequality constraints. Instead, they require the designer to approximate the effects of constraints with conservative assumptions. (Rawlings, 2000) provides an overview of MPC theory and practice for centralized applications. (Camponogara, 2002) describes the adaptation of MPC to distributed agent networks.

1.6 Project goals

The high-level goal of this work is to provide means for operating power networks with better tradeoffs between conflicting objectives, specifically focusing on tradeoffs between reliability and protection. The specific goal addressed in this paper is to develop a network of distributed, autonomous, cooperative agents capable of eliminating power system violations before the protection system acts to disconnect equipment. If this method can mitigate the effects of at least one future cascading failure without triggering or increasing the severity of others, holding everything else constant, the method will be capable of increasing reliability without negatively affecting other operating objectives. If reliability can be increased without affecting other objectives, effectively improving the Pareto frontier for the operating objectives, it may also be possible to move along the new Pareto surface to obtain better tradeoffs between conflicting objectives.

2 Eliminating violations as a means to prevent cascading failures

This section gives the global problem formulation, which comes from the control goal in section 1.6 and which we use in section 3 to build our problem decomposition.

Most of the state transitions that make up a cascading failure are caused by transmission line relays reacting to high currents and low voltages. These variables are highly sensitive to changes in load levels and generator outputs. In many cases, the network can tolerate violations for a time without negative consequences. A transmission line overcurrent condition can persist for seconds or minutes before the conductors sag enough to allow a phase to ground fault and trigger a relay action. Even a severe overload that could trigger a backup (zone 3) relay will operate with a 1-2 second time delay (Blackburn, 1998, ch. 12). If voltage and current violations can be eliminated through fast load and generator control, transmission line relays will not act to propagate a cascade.

2.1 Problem formulation

With this in mind we use the following control problem as a means of preventing cascading failures: eliminate voltage and current violations with a minimum cost set of load and generation shedding violations before subsequent failures occur. For the sake of this paper, we consider this to be globally correct behavior. This problem can be formulated as a non-linear programming problem, using the steady state power network equations that would ordinarily be used in an optimal power flow formulation (Wood, 1996, ch. 13). This global problem (P) is given in (1a-1h) below.

$$\underset{G,L}{\text{minimize}} \sum_{n \in N} Cost_n \left(G_n - G_{n0}, L_n - L_{n0} \right) \tag{1a}$$

subject to:

$$I = Y_{NN} V \tag{1b}$$

$$G_n - L_n = V_n conj(I_n), \, n \in N \tag{1c}$$

$$\text{Re}\left(L_n / L_{n0} \right) = \text{Im}\left(L_n / L_{n0} \right), \, n \in N \tag{1d}$$

$$G_n^{\text{min}} \leq G_n \leq G_n^{\text{max}}, \, n \in N \tag{1e}$$

$$0 \leq L_n \leq L_{n0}, \, n \in N \tag{1f}$$

$$|V|^{\text{min}} \leq |V| \leq |V|^{\text{max}} \tag{1g}$$

$$\left| I_{nm} \right| = \left| y_{nm}(V_n - V_m) \right| \leq \left| I_{nm} \right|^{\text{max}}, \, n, m \in N, n \neq m \tag{1h}$$

where:

N is the index set of all the nodes in the network.

n is the index of the agent located at bus n.

Q is the index set of all the branches in the network.

V is a complex vector of node voltages. V_{nk} is the voltage at bus n at time step k.

I is a complex vector of currents. I_n is the injection at bus n. I_{nm} is the current along branches between nodes n and m.

G is a complex vector of generation power injections. For the sake of notational simplicity, we assume no more than one generator is located at each bus. It is fairly easy to incorporate multiple generators, but doing so complicates the notation somewhat. G_{n0} is the measured pre-control generator output at bus n.

L is a complex vector of load powers. As above, we assume one load at each bus. L_{n0} is the measured pre-control demand at bus n.

Y_{NN} is the complex node admittance matrix for all the nodes in the network.

Y_Q is the complex branch admittance matrix for the set of all branches in the network.

y_{nm} is the single element of the node admittance matrix that is the admittance between buses n and m.

The costs associated with shedding load (from 1a) are the social costs that would be incurred from the interruption of electrical service. If SOs deem some loads as more valuable than others, the objective function (1a) can be adjusted accordingly. The costs associated with reducing generation come from either the expected equipment damage resulting from rapid deceleration (using techniques such as fast valving or breaking resistors), or the amount that would have to be paid to an independent power producer for such emergency control. Equality constraint (1b) defines the voltage-current relationships in the network. Equality constraint (1c) expresses conservation of energy at each node. Equality constraint (1d) forces the system to shed real and reactive load in equal proportions. Inequality constraints (1e) and (1f) describe the extent to which loads and generation can be adjusted. The final inequality constraints (1g and 1h) define the measures used to identify violations. This formulation can be extended to include constraints on the dynamic system, such as system frequency or generator "out-of-phase" limits, but such extensions are beyond the scope of this paper.

Simulations on several test networks indicate that power system violations can be eliminated by solving this problem and implementing the resulting control actions. We do not presume to be able to eliminate all cascading failures using this method. This method will not likely do much to control high speed (<1 second) cascading failures that result primarily from machine dynamics. Most cascading failures, however, are not of this type and progress over periods of seconds to minutes. We have found that standard non-linear solvers[9] quickly find optimum solutions to this problem for small networks (<200 buses). Figure 2.1 shows the result of one such calculation using the IEEE 39 bus test case.

---

[9] For all simulations in this paper we use the SNOPT solver via the Tomlab interface: www.tomlab.biz.

Figure 2.1—Optimal load and generation shedding actions resulting from the solution of the global problem *P* on the IEEE 39 bus test case. These are the minimum cost actions that eliminate the violations shown. Branch current violations are marked with circles. These violations occurred after a line outage was applied at the location marked with an X. The arrows indicate power flow magnitude and direction.

3 Problem decomposition: distributed MPC and cooperation

Because cascading failures can spread rapidly through an entire synchronous network, obtaining good solutions to *P* requires that *P* be solved over an entire synchronous network. For large networks this is technically impossible, and in many locations would require a degree of centralization that is institutionally impractical. Thus a decentralized solution is necessary.

Given this necessity, our problem is to take *P* and decompose it into tractable sub problems that, when solved and implemented by distributed agents, result in the desired global behavior. Our decomposition method can be described as follows:

- place a software agent at each load and generation bus;
- allow each agent to control only its local control variables;
- allow each agent to gather measurements from a limited portion of the system through communication networks;
- allow each agent to reduce its problem into local, tractable versions of the global problem that can be solved iteratively using MPC and cooperation.

The result is a two-dimensional decomposition of the global control problem. The problem is decomposed in space by assigning the problem to distributed agents and in time by allowing agents to act iteratively using MPC. Figure 3.1 provides a high level view of this method.
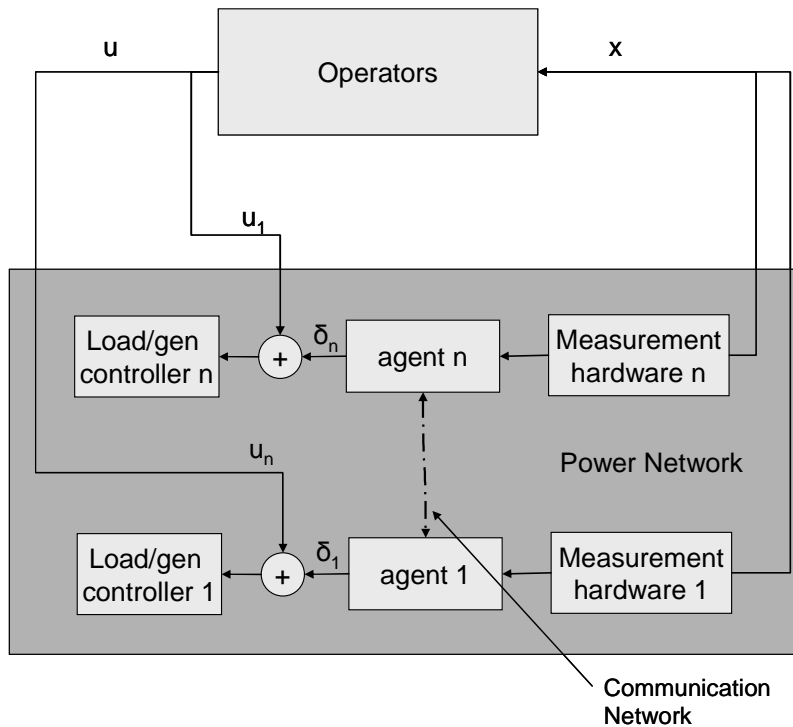
Figure 3.1—Feedback diagram of the system showing how operators and the agent network interact. Under normal conditions operators control load and generator set points (u). When an agent detects a violation, it calculates a plan and implements the local portion of that plan effectively making adjustments (δ_1) to the operator set points.

3.1 Spatial decomposition

In order to balance communication constraints and scalability with the need for agents to make good local decisions, we task each agent with collecting data frequently from its local neighborhood, and rarely or never from more distant locations. Therefore, each agent has a local neighborhood within which the agent collects data frequently (several times per second), and an extended neighborhood within which the agent can take infrequent (daily or weekly) data measurements. An agent cannot take measurements from buses outside of its extended neighborhood. All nodes that can be reached by traveling over no more than r_l branches are inside an agent's local neighborhood. Similarly all nodes within a radius r_e are within an agent's extended neighborhood. Figure 3.2 illustrates this spatial decomposition using the IEEE 118 bus network.
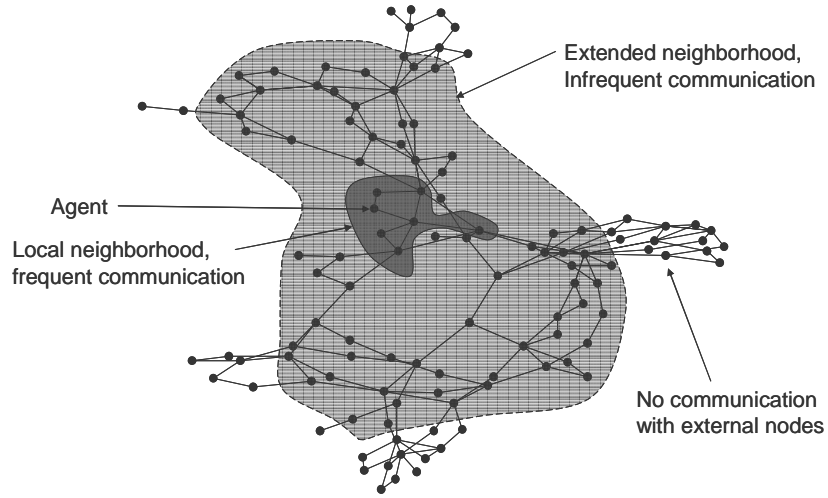
Figure 3.2—Illustration of the spatial decomposition of the global problem (*P*) on the IEEE 118 bus network. An agent gathers measurements and watches for violations frequently within its local neighborhood and occasionally within its extended neighborhood. The radius of an agent's local neighborhood is r_l and the radius of an agent's extended neighborhood is r_e.

Since it is difficult to solve P for large networks, an agent must reduce P into a tractable local version. The agents do this using a linear-difference version of P. This formulation is presented below (2a-2d).

$$\underset{\delta_M}{\text{minimize}} -c_M{}^T \delta_M \tag{2a}$$

subject to:

$$\sum_{g \in G_M} \delta_g = \sum_{l \in L_M} \delta_l \tag{2b}$$

$$\left| I_R \right| = \left| I_{R0} \right| + D_{RM} \delta_M \leq \left| I_R \right|_{\max} \tag{2c}$$

$$\left| V_M \right|_{\min} \leq \left| V_{M0} \right| + E_{MM} \delta_M \leq \left| V_M \right|_{\max} \tag{2d}$$

$$-u_{M0} \leq \delta_M \leq 0 \tag{2e}$$

where:

M is the index set of all nodes (or control variables) within agent *n*'s extended neighborhood. $M \subset N$.

R is the index set of all branches within agent *n*'s extended neighborhood. $R \subset Q$

c is a vector of costs associated with control variable reductions.

u is a vector of control variables: $u = \begin{bmatrix} \text{Re}(G) \\ \text{Re}(L) \end{bmatrix}$.

δ is a vector of control variables changes: $\delta = u - u_0$. δ_M is a sub-vector of δ corresponding to the control variables within an agents external neighborhood.

δ_g, δ_l are elements of the control vector corresponding to generation or load.

D is a matrix of branch current sensitivity factors such that $|I_{Qk}| \cong |I_{Q,k-1}| + D\delta_k$.

E is a matrix of bus voltage sensitivity factors such that $|V_{N,k}| \cong |V_{N,k-1}| + E\delta_k$.

In this simplification of P, we reduce all of the non-linear equality constraints into a simple balancing constraint (2b) that forces the system to choose to shed load and generation in equal quantities. While this does not guarantee that the system will maintain a perfect energy balance, the small errors should be easily picked up by existing frequency regulation mechanisms. If a significant frequency imbalance were detected, (2b) could be weighted to compensate for the imbalance.

This formulation has several distinct advantages over the full non-linear problem (P). Constraints that are distant or unimportant and remote variables that will not affect the final solution can be easily removed from the local problem. The difference nature of (2) allows for state measurements to be incorporated into the problem, making iterative solutions feasible. Finally, compared to (1), (2) is noticeably less sensitive to errors in remote control variable estimates. This reduced sensitivity stems from two aspects of the formulation. Firstly, it uses the DC load flow approximations to calculate the load distribution factor matrix (D) as follows:

$$D = \mathrm{Im}(Y_{QN})\,\mathrm{Im}(Y_{NN}^{-1})\Lambda \tag{3}$$

where Λ is a matrix that translates the control variables into bus power injections:

$$\Lambda: \ G - L = \Lambda \begin{bmatrix} G \\ L \end{bmatrix} \tag{4}$$

This allows the agent to calculate D without any knowledge of the network state (voltage, current) and control (generation, load) variables. An agent must only know the status of the branches in the network to calculate D correctly. Even if some status errors exist in the agent's model of distant parts of the network, the important elements of D will be nearly correct. Secondly, the agent only needs to know the amount of load and generation at a particular location if it decides that the entire quantity should be eliminated from the network. Since this case occurs infrequently, large errors are infrequent. The disadvantage is that solutions to (2) will be sub-optimal solutions to P, even given perfect information. This sub-optimality is acceptable since we are primarily concerned with eliminating violations. Doing so at minimum cost is secondary.

3.2 Temporal decomposition

It is not difficult to adapt the linear difference formulation (2) to make use of MPC. We add a time dimension to the control vector (δ), giving us a decision matrix (Δ)—a two dimensional control plan. The cost function therefore is a summation of control costs over the time horizon. The costs are discounted so that the least expensive actions will be chosen first, and more expensive actions later. The solution appears to be independent of the discount rate chosen for discount rates such that $0<\rho<1$. Additionally, we add some slack to the constraints so that the violations need not be entirely eliminated during the first period, but can be gradually eliminated over time. Finally, we add ramp rate constraints on the generators since there are natural limits to how fast a generator can decelerate. Thus, the DMPC problem for agent n at time k_0 (problem P_{n,k_0}) can be written:

$$\underset{\Delta_{MK}}{\text{minimize}} \sum_{k=k_0}^{K} e^{-\rho k} c_M^T \delta_{Mk} \tag{5a}$$

subject to (for $k=1\ldots K$):

$$\sum_{g \in G_M} \delta_{gk} = \sum_{l \in L_M} \delta_{lk} \tag{5b}$$

$$\left| I_{Rk} \right| \cong \left| I_{Rk-1} \right| + D_{RM} \delta_{Mk} \leq f_1(I_{R,0\ldots k-1})\left| I_R \right|^{\max} \tag{5c}$$

$$f_2(V_{M,0\ldots k})\left| V_M \right|_{\min} \leq \left| V_{Mk} \right| \cong \left| V_{M,k-1} \right| + E_{MM} \delta_M \leq f_3(V_{M,0\ldots k})\left| V_M \right|_{\max} \tag{5d}$$

$$RR_g \leq \delta_{gk} \leq 0, \; g \in G_M \tag{5e}$$

$$-u_{l0} \leq \delta_{lk} \leq 0, \; l \in L_M \tag{5f}$$

$$-u_{M0} \leq \sum_{k=k_0}^{K} \Delta_{Mk} \leq 0 \tag{5g}$$

where:

k is the current time step,

K is final time step for the solution horizon, and

ρ is a discount factor.

System operators should be able to estimate how quickly a violation must be eliminated to prevent relay operation. This will depend on both the magnitude of the violation and the time that the violation has persisted on the system. To prevent a zone three or a time over-current relay operation, a violation will need to be eliminated fairly quickly (1-2 seconds). In order to minimize the risk of a line sagging and causing a fault, longer time delays (seconds to minutes) will likely be acceptable. This relationship between time and violation magnitude is encoded into (5) in the

constraint multiplier functions f_1, f_2, and f_3 in (5c, 5d). In this paper we use four period simulations and define f such that agents will seek to reduce current violations linearly from 130% of the limit in the first period to 100% in the final period (figure 4.1 shows this control goal). Voltage violations (5d) are currently neglected, but will be included in future revisions. If a violation persists past the original planning horizon the agent continues to act to reduce the violation below the threshold. The result of each calculation is a control plan $\Delta_{WK}{}^*$. This plan is illustrated in figure 3.3.

$$\Delta_{MK} = \begin{bmatrix} \delta_{1k_0} & \cdots & \delta_{1K} \\ & \cdots & \\ \delta_{Mk_0} & \cdots & \delta_{MK} \end{bmatrix}$$

Immediate action to be taken locally

Space

Estimate of what a remote agent will do some time in the future.
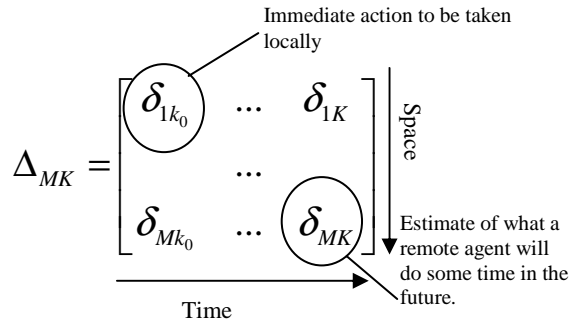
Time

Fig. 3.3—Diagram of the two dimensional decision plan calculated by agent n. The agent obtains this control plan by solving (5) given the data that it was able to collect from other agents. The columns of the matrix represent a set of calculated control actions for a single time period and for every control variable in set M. The rows represent a set of calculated control actions for a single location over the entire solution horizon.

3.3 Cooperation

According to our earlier definition of cooperation as sharing goals and exchanging useful information, an agent that merely solves (5) and acts is not cooperative. Such an agent uses an overlapping objective function (5a) but does not exchange useful information with its neighbors before taking action. In order to improve performance we have studied a variety of designs for cooperative agents.

The algorithm presented in this paper is based on our finding that agents with only local information can overlook important data located just outside the agent's local area. Consider two agents: A and B. A is near a violation that B should react to, but B is unaware of the problem because it lies just outside of B's neighborhood (but not A's). If A solves its problem and calculates that B should act, and then shares the important violation data with B, B will likely be able to make better decisions about its local control actions. If B replies and shares its local data with A, A may also be able to improve its solution. Table 3.1 provides a more general description of this cooperation algorithm.

Table 3.1—Agent algorithm with cooperation. Unilateral agents skip steps 6-10.

Condition	Step	Action
No violations detected	1	Collect data from neighbors
	2	Update network model
	3	Check for violations. If violations found goto 5.
	4	Repeat from 1.
One or more violations detected	5	Solve (5) to obtain the control vector for the current time period ($\delta_{M,k}$)
	6	Determine a set of agents (Q) that appear to require control actions.
	7	Compare solutions with those agents in set Q.
	8	If a large discrepancy is found, exchange data with the agents with whom there exists a discrepancy.
	9	Re-solve (5) with the updated data.
	10	Iterate from 6 until consensus is reached, or until a maximum number of iterations has occurred.
	11	Implement the local portion of the calculated control actions

This rather simple cooperation algorithm was found to be quite effective. Each agent may begin with severely limited information but through the cooperation process the relevant agents obtain more detailed information about important aspects of the network. In our simulations we found that agents reach consensus within one or two iterations. We limit this process to three iterations.

4   Verification

In this section we describe the results of simulations designed to evaluate this method. The following experiments are specifically designed to determine the relationship between agent performance and communication abilities. The below results apply to simulations on the IEEE 118 bus test case, though similar results obtain using other networks. The 118 bus case was modified slightly from the original to match its properties to those of a typical contemporary power system.

4.1   Simulation model description

For the following simulations we use a standard, non-linear, power flow network model with constant real/reactive power loads and constant power/voltage generators. The network is assumed to perform frequency regulation through a single slack bus. The initial condition of the network is calculated with an optimal power flow algorithm (Zimmerman, 1997). One agent is placed at each bus and has the capabilities specified in section 3. Table 4.1 summarizes the important model input parameters and assumptions.

Table 4.1—Model input parameters

Input	Description
Network data	Modified from the IEEE 118 bus test case (see Appendix A)
Load shedding costs	Randomly assigned between $500/MW and $1500/MW
Generator shedding costs	Assigned uniformly at $30/MW
Solution horizon (K)	4 time steps
External neighborhood radius (r_e)	10 branches
Local neighborhood radius (r_l)	Varies between 1 and 6 branches
External data estimation error	15% coefficient of variation (δ_x/x)
Initiating disturbances	Chosen randomly from a set of 100 violation inducing double branch outages

A simulation is initiated by choosing a disturbance, a local neighborhood radius (r_l), and allowing agents to sample data from the pre-fault condition of their external networks. During each simulation time step the agents solve their local problems, and implement the required local control action. After the agents have finished their calculations, the affect of agent control actions is calculated using a power flow routine. Table 4.2 describes this simulation procedure in more detail. For every disturbance/radius combination steps 5-11 were repeated for both cooperative agents and and unilateral agents.

Table 4.2—Simulation procedure

Step	Action
1	Choose a disturbance randomly from the set of double contingencies.
2	Choose an internal neighborhood radius (r_l)
3	Run an optimal power flow to obtain the pre-disturbance network conditions.
4	Allow the agents to take noisy measurements from buses within their extended neighborhoods.
5	Model the disturbance and run a power flow calculation.
6	Set $k=0$, $K=4$.
7	Allow the agents to take measurements from their local neighborhoods.
8	Allow the agents to calculate control plans for the control horizon ($k+1…K$)
9	Incorporate the agent control actions into the network data by changing load and generation.
10	Calculate the new network conditions using a power flow.
11	Increment k (and K if $k+1>K$).
12	Repeat from 6 until all of the violations are eliminated, or until it is clear that the agents will not be able to eliminate the remaining violations.

4.2 Results

What follows are results from 771 simulations using the above procedure and sampled using the assumptions in table 4.1. Each simulation is repeated for agents with and without cooperation.

Figure 4.1 shows the trajectory of the most severe violation resulting from a typical disturbance for both cooperative and unilateral agents. Figures 4.2 and 4.3 show the relationship between the quantity of communication (internal neighborhood size) and two measures of performance: control error and completion time (see figure 4.1 for definitions).


Figure 4.1—Violation trajectories that result from agents reacting to a disturbance in the IEEE 118 bus test case. The agents have a small local neighborhood ($r_l = 2$). The disturbance consists of outages on branches 8 and 40. The cooperative agents eliminate the violation nearly along the control goal. The control error is the area of the space between the control goal and the actual trajectory. For the cooperative agents this area is quite small, while for the unilateral agents it is rather large. The completion time is the number of time iterations required to reduce the violation to no more than 0.05 (5% above the constraint). For the cooperative case, the completion time is 4. For the unilateral case, the completion time is set to 10 (beyond the solution horizon).


Figure 4.2—The relationship between r_l and control error (see figure 4.1 for definition).

Figure 4.3—The relationship between r_l and completion time (see figure 4.1 for definition). On average the cooperative agents require one additional time step to remove their violations, independent of the local neighborhood size.

4.3 Discussion

The above experiments reveal some important properties of the agent network design. The experiments indicate that it is possible to eliminate power system violations using autonomous agents working with a power flow network model. This conclusion applies to an actual power system so long as the time between MPC iterations is sufficiently large that the network nearly arrives at a steady state before the next control action occurs. As long as the generator actions can be accomplished quickly, this condition should hold. Tests using a dynamic power system simulator may provide additional insight. The experiments also demonstrate the value of even simple cooperation schemes in agent networks. Without cooperation, the communication required to obtain acceptable performance may be beyond what can be expected from existing technology.

5 Policy issues: a preliminary analysis

Special protection schemes can have substantial benefits to a system, but these benefits are in terms of services such as reliability of the bulk electric grid and transmission capacity. In the case of reliability, the benefits will accrue to all customers almost uniformly (especially those without backup generation). In the case of transmission capacity it is difficult to determine the distribution of benefits. The costs however fall to those that own and operate the power grid: vertically integrated utilities, transmission owners, and independent system operators. Both reliability and network

capacity have properties of public goods. Additionally, a distributed-agent SPS will only be able to control cascading failures on control area seams if the systems in adjacent control areas are coordinated. Due to the difficult coordination issues involved, the concentrated nature of the costs, and the dispersed nature of the benefits, investment is unlikely to occur without regulatory intervention. Before regulators act to promote the dispersion of any technology, the costs, benefits, and risks of that technology should be carefully weighed. This section gives a preliminary description of the nature of some of these costs, risks, and benefits specifically for the US Eastern Interconnect.

5.1 Benefits

The primary benefit of this technology is to reduce the cascading failure risk. If we know the expected cost of cascading failures to a system, and know that a given technology can reduce this probability by some portion, we can estimate the reliability value of that technology. If a cascading failure on the scale of the 2003 Northeast blackout occurs once every 15 years and one half this size occurs twice as frequently, assuming that blackout costs scale linearly from the $6 billion estimate, the expected cost of large cascading failures is $800 million/year.[10] A technology that could cut both frequencies in half would have a $400 million/year reliability benefit.

A secondary benefit of this technology would be the ability to use existing transmission capacity more efficiently. Because power grids do not currently react to disturbances well, operators must use transmission capacity sparingly to allow for possible disturbances. If the grid had better reflexes it would be possible in many locations to use transmission grids more aggressively, allowing for efficiency gains.

In future work we plan to study both of these benefits for test case networks to estimate their relative magnitudes and significance.

5.2 Costs

Because the proposed method does not require additional high-voltage hardware, the costs per location should be low. The computational requirements for the agents themselves are no more than the abilities of a standard PC. As the costs of wireless, satellite, and power line communication equipment decrease, the communication system costs should be reasonable as well. Table 5.1 shows an order-of-magnitude approximation for the cost of implementing this technology for the US

---

[10] This is close to the $1 billion/year cost used in (Apt, 2004)

eastern interconnect.[11] Relative to the social cost of blackouts and the cost of building new transmission it seems likely that the benefits of this technology could outweigh its costs.

Table 5.1—A preliminary cost estimate for installing a network of autonomous control agents at every node of the eastern interconnect. Assumed discount rate is 7% with a 30 year planning period.

Agent hardware	50,000 buses x $2000 each = $100 million
Installation	50,000 buses x $10,000 each = $500 million
Maintenance	50,000 buses x $1000 every 5 years = NPV: $100 million
Communications	25,000 buses x $2000 each = $50 million
Total	$750 million or $60 million/year

5.3 Risks

One interpretation of the No Free Lunch Theorem of Optimization is that it is impossible to make a complex system resistant to one set of disturbances without also making it more susceptible to others (Ho, 2001). It is possible that wide-spread implementation of SPS like that presented here could result in a system that was more resistant to some conditions and more susceptible to others. There is therefore a non-zero chance that the new technology could increase rather than decrease the cascading failure risk.

For example, it is possible that during a very high speed cascading failure with voltages and currents fluctuating rapidly this method would propagate rather than arrest the problem. An important step in developing this method for practical use would be to develop a method for characterizing the current local condition as one that the agent is capable of reacting to successfully, or not. An agent could monitor the rate of change of voltage phase angles and frequency, and classify the current condition based on this data.

The existence of risks like the above does not mean that SPS should not be used. The important question is whether the benefits justify the risks. Further research is needed to understand the effect that distributed agents would have on cascading failure risk.

6   Conclusions

Cascading failures and blackouts result from violations that have been allowed to persist long enough to trigger the protective system. Experiments performed for this paper demonstrate that it is possible to design a network of autonomous agents with limited communication abilities that can eliminate power network violations before they can trigger the protective system. Experiments also

---

[11] See (Apt, 2004) for a related calculation

demonstrate that cooperative agents can outperform agents acting unilaterally by a large margin. While related to existing work on DMPC, the proposed method is new. With additional development, it seems probable that this technology could improve the ability of power systems to make better tradeoffs between conflicting objectives. In future work we plan to refine this method, perhaps through the use of improved network models, or machine learning, and further study the benefits, costs, and risks of distributed agents for power networks.

While this work is primarily technology-focused, there are important implications for current regulatory issues. The benefits of improved transmission control are diverse, difficult to accurately quantify, and widely distributed. Under current regulation, transmission owners, regional transmission organizations, and state regulated utilities are responsible to fund the transmission investments. Large investments are unlikely to occur unless regulatory bodies ensure that investors can recover their capital. Additionally, the installation of a distributed agent network will require substantial coordination among system operators across state boundaries. This is likely to require at least some regulatory oversight. On the other hand, building an agent control network, does not require large centralized regional transmission operators (RTO). This is an important result, because it implies that it is possible to solve some global power network problems without centralization. This property may ease implementation if it is found that a technology like this one is needed, but it is institutionally infeasible to organize the US grid into large RTOs. Finally, a coordinated method of determining the social cost of load shedding must be adopted before this method can be implemented. If SOs could set costs unilaterally, neighboring systems may be able to prevent local load shedding by setting very high local load values. Uniform cost assignment may be the best method initially, until a more refined method can be implemented.

Acknowledgements

Works Cited

Anderson, P. M. and K. K. LeReverend. "Industry Experience with Special Protection Schemes." *IEEE Transactions on Power Systems*, v. 11, no. 3. 1996.

Apt, J., L. B. Lave, S. Talukdar, M. G. Morgan, and M. Ilic. "Electrical Blackouts: A Systemic Problem." *Issues in Science and Technology*. Summer 2004.

BBC. "Italy launches blackout inquiry." BBC news online. 30 Sept. 2003. Online: news.bbc.co.uk.

Blackburn, J. L. *Protective Relaying: Principles and Applications*. Marcel Dekker, New York, 1997.

Camponogara, E., D. Jia, D., B.H. Krogh and S. Talukdar, "Distributed Model Predictive Control," *IEEE Control Systems Magazine*, v. 22, no. 1, 2002.

Camponogara, E.. *Controlling Networks with Collaborative Nets*. Electrical and Computer Engineering. Pittsburgh, Carnegie Mellon University. 2000.

CNN. "Italy recovering from big blackout." CNN. 28 Sept. 2003. Online: www.cnn.com/world.

Cohen, Guy and Dao Li Zhu, "Methods in Large Scale Optimization Problems." *Advances in Large Scale Systems,* v. 1, 1984.

Coury, D.V., J.S. Thorp, K. M. Hopkinson, and K. P. Birman, "An agent-based current differential relay for use with a utility intranet." *IEEE Trans. on Power Delivery*, v. 17, no. 1. Jan. 2002.

Dobson, Ian, B. A. Carreras, D. E. Newman. "A criticality approach to monitoring cascading failure risk and failure propagation in transmission systems." Proc. of the Carnegie Mellon Transmission Conf., Pittsburgh, Dec. 2004. Online: www.ece.cmu.edu/~tanddconf_2004.

Durfee, E. H.. "Practically Coordinating." *AI Magazine*, v.20, no.1. 1999.

Fink, L. H. and K. Carlsen. "Operating under Stress and Strain." *IEEE Spectrum*: Mar. 1978.

Fink, L. H. and P. J. M. van Son. "On system control within a restructured industry." *IEEE Transactions on Power Systems*, v. 13, no. 2. 1998.

Fisher, David A. and Howard F. Lipson, "Emergent Algorithms: A New Method for Enhancing Survivability in Unbounded Systems." *Proc. of the 32nd Hawaii Int. Conf. on System Sciences*. IEEE, 1999.

Graves, Frank. and Lisa Wood. "Economic Cost of the August 14th 2003 Northeast Power Outage." The Brattle Group. August 2003. Online: www.brattle.com.

Hines, Paul and Sarosh Talukdar. "Reducing the Costs of Disturbances to the Electric Power Network." Presented at the Int. Assn. for Energy Economics meeting, Washington D.C., June 2004.

Hines, Paul, Dong Jia, and Sarosh Talukdar. "Distributed Model Predictive Control for Electric Grids." *Proc. of the Carnegie Mellon Transmission Conf.*, Pittsburgh, Dec. 2004. Online: www.ece.cmu.edu/~tanddconf_2004.

Hines, Paul, H. Liao, D. Jia, and S. Talukdar. "Autonomous Agents and Cooperation for the Control of Cascading Failures in Electric Grids." *Proc. of the IEEE Conf. on Networking, Sensing, and Control.* Tuscan, Mar. 2005.

Ho, Yu-Chi and David L. Pepyne. "Simple Explanation of the No Free Lunch Theorem of Optimization." *Proc. of the 40th IEEE Conf. on Decision and Control*, Orlando, Dec. 2001.

ICF Consulting. "The Economic Cost of the Blackout." August 2003. Online: www.icfconsulting.com.

Jung, J. and C. C. Liu, "Multi-agent technology for vulnerability assessment and control," *Proc. Power Engineering Society Summer Meeting*, July 2001.

Jung, J., C. Liu, S. L. Tanimoto, V. Vittal. "Adaptation in Load Shedding Under Vulnerable Operating Conditions." *IEEE Transactions on Power Systems.* v. 17, no. 4, Nov. 2002.

Kamwa, I., R. Grondin, Y. Herbert. "Wide-area measurement based stabilizing control of large power systems-a decentralized/hierarchical approach." *IEEE Transactions on Power Systems.* v. 16, no. 1., Feb. 2001.

Kim, B. H. and R. Baldick, "A Comparison of Distributed Optimal Power Flow Algorithms." *IEEE Trans. on Power Systems.* v. 15, no. 2, May 2000.

Kolluri, S., K. Tinnium, and M. Stephens. "Design and operating experience with fast acting load shedding scheme in the Entergy System to prevent voltage collapse." *Proc. of Power Engineering Society Winter Meeting*, IEEE. Jan. 2000.

Kueck, J. D. and B. J. Kirby. "The Distribution System of the Future." *The Electricity Journal.* 2003

Liao, Huaiwei, Jay Apt, and Sarosh Talukdar. "Phase Transitions in the Probability of Cascading Failures." *Proc. of the Carnegie Mellon Transmission Conf.*, Pittsburgh, Dec. 2004. Online: www.ece.cmu.edu/~tanddconf_2004.

Madani, V., M. Adamiak, and M. Thakur, "Design and Implementation of Wide Area Special Protection Schemes," *Proc. 57th Annual Texas A&M University Conference for Protective Relay Engineers,* 2004.

Modi, P., W. Shen, W. Tambe, and M. Yokoo, "ADOPT: Asynchronous Distributed Constraint Optimization with Quality Guarantees." *Artificial Intelligence Journal*, 2004.

Neil, Charles E. "Entering the Seventh Decade of Electric Power." *Edison Electric Institute Bulletin.* Sept. 1942.

NERC. "Examples of major bulk electric system power outages." NERC. Online: www.nerc.com. "BlackoutTable.pdf" accessed Jan. 2005.

Novosel, D., M. M. Begovic, and V. Madani. *IEEE Power and Energy Magazine*, v. 2 , no. 1 , Jan-Feb 2004.

Rehtanz, C. and J. Bertsch. "A new wide area protection system." *IEEE Power Tech Conf. Proc.*, Sept. 2001.

Roy, Anjan, and P. Pentayya. "Experience of Blackouts and Restoration Practices in the Western Region of India." Slides presented at the IEEE Power Engineering Society General Meeting, Denver, July 2004.

Talukdar, S. and E. Camponogara. "Collaborative Nets." *Proc. of the 33rd Hawaii International Conference on System Sciences*. IEEE, Hawaii, 2000.

Talukdar, S. N., J. Apt, M. Ilic, L. B. Lave, and M. G. Morgan. "Cascading failures: Survival versus Prevention." *The Electricity Journal.* Nov. 2003.

Vassell, Gregory S. "Northeast Blackout of 1965." *IEEE Power Engineering Review.* Jan. 1991.

Wikipedia. "New York City blackout of 1977." Online: www.wikipedia.org. Jan. 2005.

Wood, A. J. and B. F. Wollenberg, *Power Generation Operation and Control.* Wiley & Sons, New York, 1996.

WSCC Operations Committee. *Western Systems Coordinating Council Disturbance Report.* WSCC, Sept. 1996.

Zimmerman, R. D. and D. Gan. "MATPOWER User's Manual, Version 2.0." PSERC, Dec. 1997. Online: www.pserc.cornell.edu/matpower.