

Controlling Cascading Failures with Cooperative Autonomous Agents

Paul Hines¹ and Sarosh Talukdar^{2,1}

¹Carnegie Mellon U., Department of Engineering and Public Policy

²Carnegie Mellon U., Department of Electrical and Computer Engineering

Abstract

Cascading failures in electricity networks cause blackouts and blackouts often come with severe economic and social consequences. Cascading failures are typically initiated by a set of equipment outages that cause operating constraint violations. These initiating events can be triggered by naturally occurring events, such as a wind storm, or human intervention, such as a terrorist attack. 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. Finally, we discuss the relevance of this work to some current policy issues.

Keywords

Cascading failures, autonomous agents, electrical power networks

Biographies

Paul Hines is currently a doctoral student in the Engineering and Public Policy department at Carnegie Mellon University. Paul earned a Master's degree in Electrical Engineering at the University of Washington in 2001 and a Bachelor's degree in Electrical Engineering at Seattle Pacific University in 1997. Paul previously worked with Alstom ESCA Corp. (now Areva T&D) developing neural network load forecasting tools, and with Black and Veatch Corp. designing electrical substation hardware and controls.

Sarosh Talukdar is Professor of Electrical and Computer Engineering and Engineering and Public Policy at Carnegie Mellon University. Sarosh studied at the Indian Institute of Technology-Madras and Purdue University, West Lafayette, IN. He works in the areas of design, electric power systems, distributed control and complex networks. He was also with McGraw-Edison, Canonsburg, PA, for several years, where he developed simulation and optimization software for electric power circuits.

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 failures¹ 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 tail (see figure 1.1). 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.

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.

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




Figure 1.1—The risk distribution of blackouts in the United States has a fat tail. The dots indicate individual blackouts from NERC data. The dashed line is an exponential distribution (Weibull) fit to the set of blackouts 800MW or less. The solid line fits the same data with a power-law distribution. (source: Talukdar, 2003).

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

<i>Date</i>	<i>Location</i>	<i>Notable consequences</i>
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 technology and regulation. The 1965 blackout led to the creation of the North American Electric 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 criteria⁵ 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 mitigate 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 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.

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).

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 agent networks 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) gives general strategies for developing distributed problem decompositions that achieve a set of desirable properties such as scalability. 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

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

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 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.

This paper is structured as follows. Section 1 is this introduction. Section 2 provides a general description of our problem and solution approach. Section 3 presents the global problem formulation in more detail, and section 4 describes our solution method in some detail. Section 5 describes our verification method and results on the IEEE 118 bus network. Finally, section 6 provides a discussion of the benefits and costs of distributed transmission network control and gives some conclusions.

2 General approach

This section describes the method that we use to develop the multi-agent system presented in this paper. The method takes the general formulation/decomposition procedure from Camponogara (2000) and adapts it to our problem, but in general terms. Sections 3 and 4 present the specific details of the application.

2.1 Global network control problem

In order to design the agent network, the global control problem must be written as an optimization problem. We are particularly interested in those problems where a disturbance causes one or more network excesses, which are detected at time t_0 and must be eliminated before time t_K . If we let X_k be a vector of continuous state variables at time t_k , U_k be a vector of continuous control variables at time t_k , and Ω_k be a vector defining the network configuration the following is the global control problem we get (OP) a formulation of the overall control problem, and therefore the problem that must be solved at each stage of an MPC algorithm.

(OP): Find and implement U_I such that:

$$\underset{U_1, \dots, U_K}{\text{minimize}} \quad C(U_0, U_1, \dots, U_K) \quad (2.1a)$$

$$\text{subject to: } X_{k+1} = \mu(U_k, X_k, \Omega_k); \quad k = 1 \dots K \quad (2.1b)$$

$$G(U_K, X_K, \Omega_K) \leq 0 \quad (2.1c)$$

where:

μ is a function that calculates the next state from the current state and control actions;

G is a function that defines the violations that must be eliminated before the last period of the control horizon.

2.2 Problem decomposition

The second step is to divide the global problem among agents. Suppose there are N autonomous agents distributed over the network. The goals of this section are to decompose the overall problem, (OP), into sub-problems, (SPO_{*i*}), such that: a) each sub-problem can be assigned to an autonomous agent, b) each sub-problem is easier to solve than the overall problem, and c) the optimal solutions of the sub-problems constitute an optimal or at least near optimal solution of the overall problem. We seek to achieve these goals by:

- i) breaking the control vector into disjoint parts,
- ii) making agent- i responsible for calculating only the i -th portion of the control vector,
- iii) allowing agent- i to assume that the other agents will calculate their parts optimally, and
- iv) using predictive models that are simplified so as to de-emphasize those distant parts of the network which are relatively insensitive to agent- i 's decisions.

Therefore, if we consider agent- i and let:

Z_{ik} be the subset of U_k , such that $U_k = [Z_{1k}, Z_{2k}, \dots, Z_{Nk}]$, and Z_{nk} is assigned to agent- i , and

Y_{ik} be the part of U_k that is not assigned to agent- i . In other words: $U_k = [Z_{ik}, Y_{ik}]$,

the following (2.2) is the subproblem that the agent must solve if it acts independently, without communicating with other agents:

(SP_{*i*}): Predict what the other agents will do and how the network will respond; that is, predict Y_{i1}, \dots, Y_{iK} and X_1, X_2, \dots, X_K . Simultaneously, solve the optimization problem:

$$\underset{Z_{i1}, \dots, Z_{iK}}{\text{minimize}} \quad C(U_0, U_1, \dots, U_K) \quad (2.2a)$$

$$G(U_K, X_K, \Omega_K) \leq 0 \quad (2.2b)$$

This problem can be rewritten exclusively in optimization terms as follows:

$$\text{(SPO}_{ik}\text{): } \underset{\{Z_{ik}, Y_{ik}\}_{1 \leq k \leq K}}{\text{minimize}} \quad C(U_0, U_1, \dots, U_K) \quad (2.3a)$$

$$\text{subject to: } x_{ijk+1} = M_{ijk}(U_k, X_k, \Omega_k); \quad j = 1 \dots J, \quad k = 0 \dots K - 1 \quad (2.3b)$$

$$G(U_K, X_K, \Omega_K) \leq 0 \quad (2.3c)$$

where:

x_{ijk} is the state of the network at time, t_k , in region, R_{ij} . In other words, the state of the entire network at time, t_k , is given by: $X_k = \{x_{ijk}\}_{0 \leq j \leq J}$

$R_{i0}, R_{i1}, \dots, R_{iJ}$ are concentric and disjoint regions of the network, such that R_{i0} is centered on agent- i , and R_{ij} is closer to agent- i than R_{ij+1}

M_{ijk} is a network-model such that $x_{ijk+1} = M_{ijk}(U_k, X_k, \Omega_k)$

In words, agent- i predicts the actions of the other agents by assuming their actions will be optimal, and agent- i predicts future states of the network with the aid of models, M_{ijk} , that are centered at its location. The models are specific to the agent. They decrease in fidelity with both distance and time. In other words, each agent has its own suite of models; distant parts of the network are less accurately represented, as are time intervals towards the end of the time-horizon.

Much of agent- i 's efforts are spent in making predictions of the actions of other agents and the response of the network to these actions. There is, of course, a tradeoff between the amount of effort and the quality of the decisions made by agent- i . The more accurate the predictions, the closer the optimal solution of (SPO $_i$) will be to the optimal solution of (OP). The cruder the predictions, the less the effort needed to make them.

Thus, agent- i must solve the overall problem, but conditioned on its unique and simplified view of the network, reflected through its use of the suite of models, $\{M_{ijk}\}$. Of course, even though agent- i predicts the entire control vector, it implements only the fraction assigned to it, and only for the first time-interval. The sub-problem is simpler than the overall problem because of the predictive models used— $\{M_{ijk}\}$ is simpler than μ especially for parts of the network that are far from agent- i .

2.3 Cooperation

We will say that two agents cooperate if they share goals (objectives or constraints) and exchange information to better meet these goals.

Two obvious forms of cooperation are: a) for agents to tell their neighbors what they intend to do, and b) to pass along state-measurements that other agents may not be able to otherwise obtain. (Each agent can sense only a small part of the network; without help from its neighbors, it cannot be expected to obtain an accurate picture of what is happening in the network.)

These two forms of cooperation simplify agent- i 's task as follows:

$$\text{(SPC}_{ik}\text{): } \underset{\{Z_{ik}\}_{x_{i \leq k \leq K}}}{\text{minimize}} \quad C(\hat{U}_0, \hat{U}_1, \dots, \hat{U}_K) \quad (2.4a)$$

$$\text{subject to: } x_{ij1} = M_{ijk}(\hat{U}_0, \hat{X}_0, \Omega_0); \quad j = 1 \dots J \quad (2.4b)$$

$$x_{ijk+1} = M_{ijk}(\hat{U}_k, \hat{X}_k, \Omega_k); \quad j = 1 \dots J, \quad k = 1 \dots K - 1 \quad (2.4c)$$

$$G(\hat{U}_K, \hat{X}_K, \Omega_K) \leq 0 \quad (2.4d)$$

where \hat{U}_k and \hat{X}_k are synthesized from agent- i 's own predictions and measurements as well as those supplied to it by its neighbors.

3 Global control problem definition

This section adapts the general global problem formulation (OP) to the specific problem of controlling cascading failures. This problem formulation lays the foundation for the problem decomposition that we define in section 4.

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.

3.1 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 (OP) is given in (3.1a-3.1h) below.

$$\text{minimize}_{G,L} \sum_{k=1}^K \sum_{n \in N} \text{Cost}_n (G_{nk-1} - G_{nk}, L_{nk-1} - L_{nk}) \quad (3.1a)$$

subject to:

$$|V|^{\min} \leq |V_k| \leq |V|^{\max} \quad (3.1b)$$

$$|I_{nm,k}| = |y_{nm}(V_{nk} - V_{mK})| \leq |I_{nm}|^{\max}, \quad n, m \in N, n \neq m \quad (3.1c)$$

and for all $k \in \{1 \dots K\}$:

$$I_k = Y_{NN} V_k \quad (3.1d)$$

$$G_{nk} - L_{nk} = V_{nk} \text{conj}(I_{nk}), \quad n \in N \quad (3.1e)$$

$$\text{Re}(L_{nk}/L_{n0}) = \text{Im}(L_{nk}/L_{n0}), \quad n \in N \quad (3.1f)$$

$$G_n^{\min} \leq G_{nk} \leq G_n^{\max}, \quad n \in N \quad (3.1g)$$

$$0 \leq L_{nk} \leq L_{n0}, \quad n \in N \quad (3.1h)$$

where:

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

n is the index for an individual member of N .

K is the final time step in a given sequence of actions.

k is an index variable for time.

V is a complex vector of node voltages. V_n is the voltage at bus n .

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. See (Wood, 1996) for definition.

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 3.1a) are the social costs that would be incurred from the interruption of electrical supply or demand. If SOs deem some customers as more valuable than

others, the objective function (3.1a) 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. The first two inequality constraints (3.1b) and (3.1c) define the measures used to identify violations, which together make up the inequality constraints in our general formulation (2.1x). This formulation could 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. Equality constraint (3.1d) defines the voltage-current relationships in the network. Equality constraint (3.1e) expresses conservation of energy at each node. Equality constraint (3.1f) forces the system to shed real and reactive load in equal proportions. Inequality constraints (3.1g) and (3.1h) describe the extent to which loads and generation can be adjusted.

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 solvers⁹ quickly find optimum solutions to this problem for small networks (<200 buses). Figure 3.1 shows the result of one such calculation using the IEEE 39 bus test case.




Figure 3.1—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.

4 Problem decomposition method: distributed MPC and cooperation

Because cascading failures can spread rapidly through an entire synchronous network, obtaining good solutions to the global problem (3.1) requires that it 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. For these, among other, reasons a decentralized solution is necessary. With this in mind, we use the method described generally in section 2 to decompose the global problem into agent sub-problems. We take the following steps to decompose the problem:

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

- place a software agent at each load and generation bus;
- task each agent with controlling only its local load and generation;
- allow each agent to gather measurements from a limited portion of the system through communication networks and use that data to populate simplified network models;
- allow each agent to solve its local problem iteratively using MPC;
- allow each agent to improve its local solutions using cooperation.

The result is essentially 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 4.1 provides a high level view of this method.

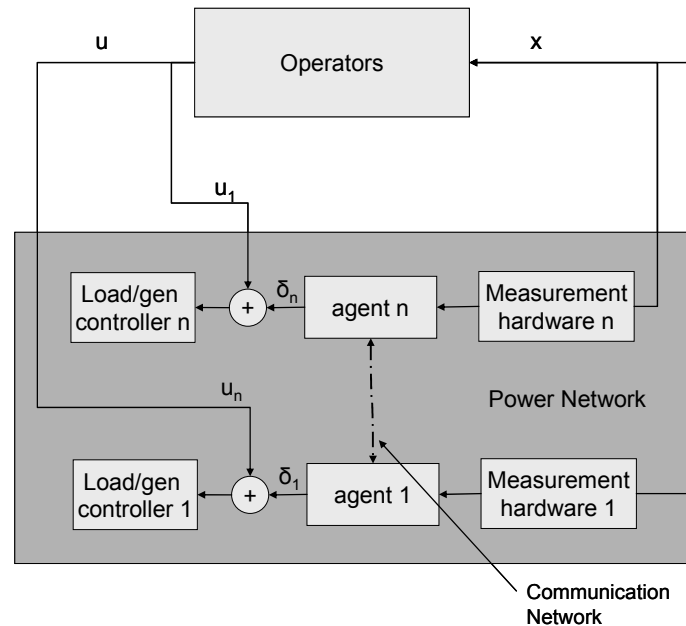


Figure 4.1—Feedback diagram of the system showing how operators and the agent network interact. Under normal conditions operators control load and generator set points (U). When an agent detects a violation, it calculates a plan and implements the local portion of that plan effectively making adjustments (δ_i) to the operator set points.

4.1 Spatial decomposition and data collection

In our decomposition we place one agent at each bus in the network, each of which has responsibility to make emergency adjustments to control variables (load and generation) at its location. Each agent obtains local state (current and voltage) measurements through local measurement hardware, and obtains more remote measurements through a communication network.

For the sake of network modeling, an agent (agent- i) divides the network into four regions, R_0 through R_3 . R_0 contains the local node (bus i) where agent- i has direct control and measurement abilities and the state variables accessible from this location. R_1 extends to every bus that can be reached by traveling over no more than r_1 branches (the sub-network of radius r_1 around bus i). We refer to this region as the agent's local neighborhood. Agent- i obtains constant measurements from buses within R_1 and therefore maintains good models of this region. The next region, R_2 , extends to every bus within a radius of r_2 from bus i . In this extended neighborhood the agent obtains infrequent (daily or weekly) measurements such that it can estimate the quantity of load and generation at these locations. This provides agent- i with crude approximations of the control abilities of agents at these remote locations but not state variables since these can change quickly during stressed conditions. The remainder of the network falls into R_3 . The agent does not take any measurements or estimates of the state or control

variables in R_3 . It does estimate the configuration of the R_3 network (Ω_{R_3}) by assuming that the remote branches are in some default state., perhaps given to the agent by a system operator.

Thus agent- i populates four network models with decreasing fidelity as the models become more remote. Table 4.1 describes the agent regions and models in more detail, and figure 4.2 illustrates the four regions on a graph of the IEEE 118 bus network model.

Table 4.1—Summary of agent measurement and prediction for the four agent regions

Region	Measurement			Prediction		
	u_{ijk}	x_{ijk}	Ω_{ijk}	u_{ijk+1}	x_{ijk+1}	Ω_{ijk+1}
R_0	local	local	local	controlled directly	predicted with M_{ijk}	$\Omega_{ijk+1} = \Omega_{ijk}$
R_1	constant measurements	constant measurements obtained	constant measurements obtained	assumed optimal given M_{ijk}	predicted with M_{ijk}	$\Omega_{ijk+1} = \Omega_{ijk}$
R_2	estimated	through cooperation	through cooperation	assumed optimal given M_{ijk}	predicted for non-empty elements of x_{ijk}	$\Omega_{ijk+1} = \Omega_{ijk}$
R_3	none	none	all branches connected	$u_{ijk+1} = u_{ijk}$	no prediction	$\Omega_{ijk+1} = \Omega_{ijk}$

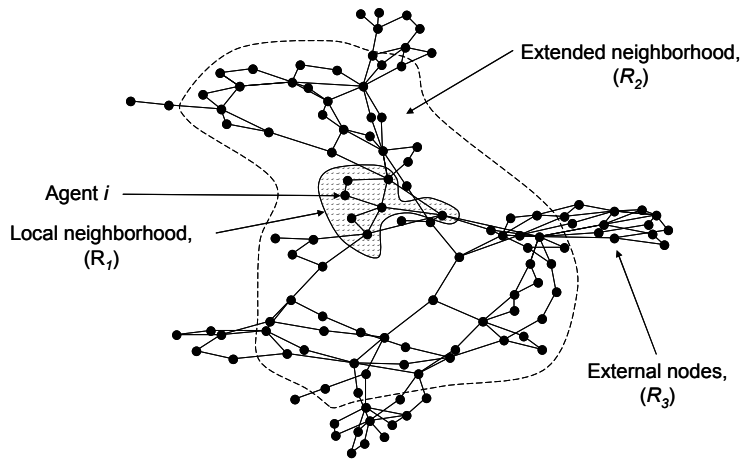


Figure 4.2—Illustration of the spatial decomposition of the overall problem on the IEEE 118 bus network. An agent gathers measurements and watches for violations frequently within its local neighborhood and occasionally within its extended neighborhood. The radius of an agent’s local neighborhood is r_1 and the radius of an agent’s extended neighborhood is r_2 .

4.2 Agent network models (M_{ijk})

As described earlier, in order to employ model predictive control each agent must use the data that it collects in combination with network models that preserve the essential functionality of the network and require minimal computation. This section describes the network models specifically used for the current DMPC application, extending the general description in section 2 and table 4.1.

4.2.1 Network model for region R_0 (M_{i0k})

The R_0 region for agent- i includes only bus i and the branches that are incident to this bus. Since the agent has good data for this region, and because the region is situated in the center of its other network models, it is generally able to make good predictions. The model includes two components. Firstly, a means of predicting future state variables and secondly a means of predicting future control variables. The

state variable prediction is a simple linear model that is used for all three of the local regions (R_0 through R_2):

$$\hat{x}_{ijk+1} = M_{ijk}(\hat{U}_k, \hat{U}_{k+1}, \hat{X}_k) = \hat{x}_{ijk+1} + A_{R_i, R_0} \hat{\delta}_{R_0, k} + A_{R_i, R_1} \hat{\delta}_{R_1, k} + A_{R_i, R_2} \hat{\delta}_{R_2, k} \quad (4.1)$$

where:

\hat{x} is a predicted or estimated value of x ;

A_{R_i, R_j} is a sensitivity matrix that gives the approximate sensitivity of the state variables in set R_i to changes in the control variables in the set R_j . See below for more details

$\hat{\delta}_{R_i, k}$ is a vector of predicted control variable changes such that:

$$\delta_{R_i, k} = u_{R_i, k+1} - u_{R_i, k} = \begin{bmatrix} \text{Re}(G_{R_i, k+1} - G_{R_i, k}) \\ \text{Re}(L_{R_i, k+1} - L_{R_i, k}) \end{bmatrix} \quad (4.2)$$

In our simulations we focus specifically on eliminating branch current magnitude excesses. Thus, the state variable prediction equation (4.1) is written in terms of branch current magnitudes. We can implement (4.1) using a Taylor's series expansion from the global problem, though for our problem this requires state, control, and configuration variable estimates for the entire network. Alternatively A_{R_i, R_j} can be estimated using a simplified model that is independent of the state and control variables. The DC power flow approximations commonly employed for power systems analysis provide us with such a model. Thus (4.1) becomes:

$$\left| I_{R_j, k+1} \right| = \left| I_{R_j, k} \right| + D_{R_j, M} \delta_M \quad (4.3)$$

where M represents the set of all control variables for which the agent has current measurements and $D_{R_j, M}$ represents a subset of the branch current (or power) distribution factor matrix (D) for the branch currents in region R_j and control variables in M . The full distribution factor matrix D is calculated from the bus admittance matrix ($Y_{NN} : I_N = Y_{NN} V_N$) and the branch incidence matrix ($Y_B : I_B = Y_B V_N$) as follows:

$$D = \text{Im}(Y_B) \text{Im}(Y_{NN}) \Lambda \quad (4.4)$$

where Λ is a matrix that translates the vector of control variables into a vector of bus power injection changes. Any unknown branch statuses are assumed to be connected.

Secondly the model includes a set of boundaries on the state and control variables. These include the boundaries that define the state variable violations that must be eliminated by the end of the control period:

$$x_{R_j}^{\min} \leq x_{R_j, k} \leq x_{R_j}^{\max}, \quad (4.5)$$

absolute boundaries on the control variables:

$$u_{R_j}^{\min} \leq u_{R_j, k} \leq u_{R_j}^{\max}, \quad (4.6)$$

and limits on the amount by which the control variables can change in a given period:

$$\delta_{R_j}^{\min} \leq \delta_{R_j, k} \leq \delta_{R_j}^{\max}. \quad (4.7)$$

Agent i uses this network model in combination with its local optimization problem (problem SPO_{*i*}, given in 4.8) to predict remote control actions and decide on local actions.

4.2.2 Network model for region R_1 (M_{i1k})

Model M_{i1k} differs from M_{i0k} only in that the measurements used to populate the model come through the communication network rather than being available locally.

4.2.3 Network model for region R_2 (M_{i2k})

Model M_{i2k} differs from M_{i0k} and M_{i1k} in that the state variable measurements and subsequently constraints are only included in the model if they are obtained through the cooperation process. The agent solicits occasional control variable measurements in order to estimate the control abilities at remote locations. We assume that the agent knows the control costs at these locations as well.

4.2.4 Network model for region R_3 (M_{i3k})

Model M_{i3k} is the most remote and least accurate of the four network models. In this area agent i does not take measurements or make predictions, except regarding the configuration of the network. It assumes that all of the branches in this region are in some default state (connected typically) unless it obtains other data during the cooperation process.

4.3 Full DMPC agent sub-problem

In order to build the agent sub-problems we integrate the above models with the general formulation (SPP₁). The result is an optimization problem with which the agent determines the actions to take locally during the current time period (the control variables that fall within R_0 and at time t_l) and obtains predictions for the remaining control variables that fall within R_0 , R_1 , and R_2 , and over a time horizon $k = \{1, \dots, K\}$. This integration results in the sub-problem for agent- i at time t_0 given in (4.8) below.

$$\underset{\Delta_M}{\text{minimize}} \sum_{k=1}^K e^{-\rho k} c_M^T \delta_{Mk} \quad (4.8a)$$

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

$$|I_{Bk}| \equiv |I_{Bk-1}| + D_{BM} \delta_{Mk} \leq f(I_{B,0 \dots k-1}) |I_B|^{\max} \quad (4.8b)$$

$$\sum_{g \in G_M} \delta_{gk} = \sum_{l \in L_M} \delta_{lk} \quad (4.8c)$$

$$RR_g \leq \delta_{gk} \leq 0, \quad g \in G_M \quad (4.8d)$$

$$-u_{M0} \leq \sum_{k=1}^K \delta_{Mk} \leq 0 \quad (4.8e)$$

where:

M is the index set of all control variables that the agent includes in its problem. This will include only the control variables in R_0 , R_1 , and R_2 ;

δ_{Mk} is the matrix of predicted control variable changes over the entire control variable set (M) and time horizon $\{1 \dots K\}$;

δ_{Mk} is the vector of predicted control variables changes for time t_k ;

c is a vector of costs associated with load and generation reductions;

ρ is a discount factor such that $0 < \rho < 1$;

B is the set of all branches for which the agent has current measurements. B includes all of the branches in R_0 , and R_1 , and those branches in R_2 for which the agent has data from the cooperation process;

f is a function that evaluates to a scalar and determines the amount by which the given state variable must be reduced during a given period;

G, g represent the set of all generator locations in the control vector and an index into that set;
 RR_g is the ramp rate for generator g : i.e. the amount by which the generator can be reduced between time steps.

In the above formulation, the cost function (4.8a) 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. Simulations indicate that the solution is independent of the discount rate chosen for discount rates in the range $0 < \rho < 1$. Equation (4.8b) gives the combined state variable (branch current¹⁰) prediction and limits for each time period (t_k). This constraint includes a scaling function f that adds some slack to the constraints so that the violations need not be entirely eliminated during the first period, but can be eliminated gradually over the time horizon. Finally, we add ramp rate constraints on the generators (4.8d) since there are natural limits to how fast a generator can decelerate, and constraints (4.8e) to ensure that the agent does not enact or predict more load or generation shedding than that which is feasible.

In order to build the branch current scaling function (f in 4.8b), system operators will need 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 must 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. In this paper we use four period simulations and define f such that agents seek to reduce current magnitude violations linearly from 130% of the limit in the first period to 100% in the final period (figure 5.1 shows this control goal). 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 an estimate of the global control plan Δ_M^* . This plan is illustrated in figure 4.3.




Figure 4.3—Diagram of the two dimensional decision plan calculated by agent- i . The agent obtains this control plan by solving (4.8) 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.

4.4 Cooperation

According to our earlier definition of cooperation as sharing goals and exchanging useful information, an agent that merely solves (4.8) and acts is not cooperative. Such an agent uses an overlapping objective function (4.8) but does not exchange useful information with its neighbors before taking action. There are many ways to design cooperation into an agent network. This paper presents results from only one of many possible methods.

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 neighborhood (R_l). 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

¹⁰ The current formulation includes only branch currents in the constraint set. Future editions will include voltage constraints as well.

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 4.2 provides a more general description of this cooperation algorithm.

Table 4.2—Agent algorithm with cooperation. Unilateral agents skip steps 6-10.

<i>Condition</i>	<i>Step</i>	<i>Action</i>
No violations detected	1	Collect data from neighbors
	2	Update network models
	3	Check for violations. If violations found go to 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

While better cooperation algorithms certainly exist, 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.

5 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.

5.1 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 4. Table 5.1 summarizes the important model input parameters and assumptions.

Table 5.1—Model input parameters

<i>Input</i>	<i>Description</i>
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 (σ_x / \bar{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 5.2 describes this simulation procedure in more detail. For every disturbance/radius combination steps 5-11 were repeated for both cooperative agents and unilateral agents.

Table 5.2—Simulation procedure

<i>Step</i>	<i>Action</i>
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 ($k+1 \dots K$)
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.

5.2 Results

What follows are results from 771 simulations using the above procedure and sampled using the assumptions in table 5.1. Each simulation is repeated for agents with and without cooperation. Figure 5.1 shows the trajectory of the most severe violation resulting from a typical disturbance for both cooperative and unilateral agents. Figures 5.2 and 5.3 show the relationship between the quantity of communication (internal neighborhood size) and two measures of performance: control error and completion time (see figure 5.1 for definitions).

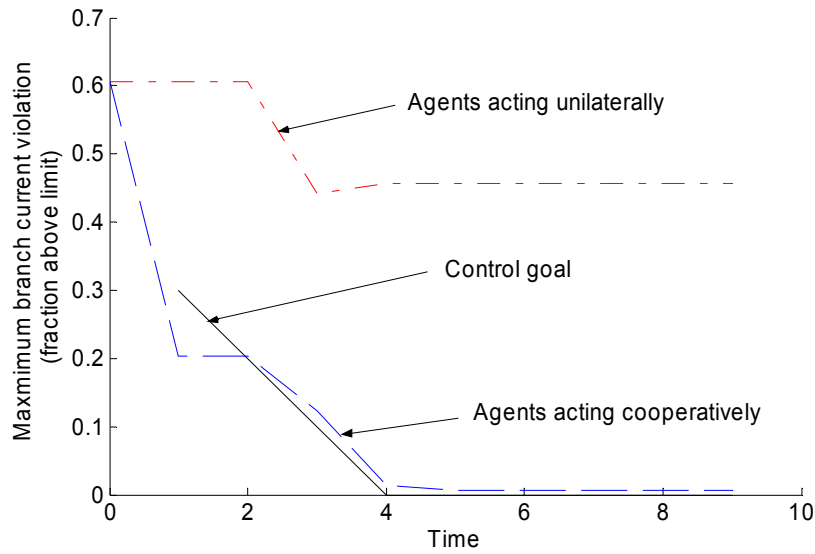


Figure 5.1—Violation trajectories that result from agents reacting to a disturbance in the IEEE 118 bus test case. The agents have a small local neighborhood ($r_l = 2$). The disturbance consists of outages on branches 8 and 40. The cooperative agents eliminate the violation nearly along the control 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).

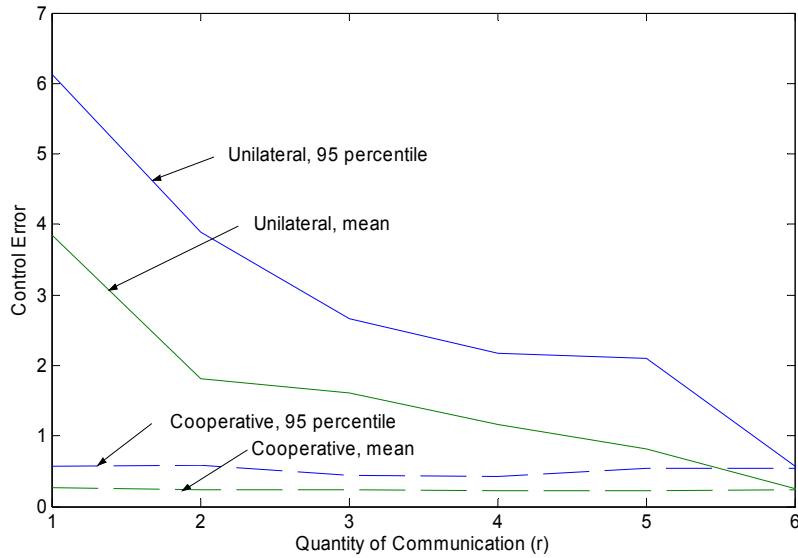


Figure 5.2—The relationship between r_l and control error (see figure 5.1 for definition).




Figure 5.3—The relationship between r_l and completion time (see figure 5.1 for definition). On average the cooperative agents require one additional time step to remove their violations, independent of the local neighborhood size.

5.3 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.

6 Benefits, costs and risks

Special protection schemes can have substantial benefits to a system, but these benefits will generally come in terms of services such as reliability of the bulk electric grid and transmission capacity. In the case of bulk network reliability, the benefits 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. Without instruments such as reliability insurance, 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.

6.1 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. This estimate is close to the \$1 billion/year cost used in (Apt, 2004), but far less than the total cost of service interruptions in the United States (\$20-\$100 billion – LBL paper <get this paper>). Given this assumed risk distribution, 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.

The net effect of an effective control system will be to create a new Pareto surface of potential tradeoffs between network reliability and efficiency. Figure 6.1 illustrates the potential effects of improved grid control and altered operating criteria.

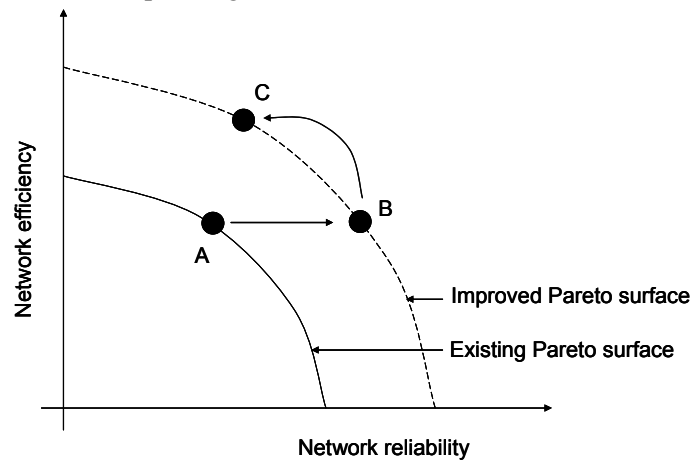


Figure 6.1—Illustration of the desired benefits of improved grid control. An effective control scheme will increase network reliability, and thereby create a new Pareto surface of potential tradeoffs between reliability and efficiency. If the existing tradeoff is at point A, the effect of improved control alone will be to move to point B along the new Pareto surface. If a set of new operating criteria were used in addition to the new control system, the effect could be to move to another point along the Pareto surface (point C for example).

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

6.2 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 6.1 shows an order-of-magnitude approximation of the cost associated with implementing this technology for the US eastern interconnect.¹¹ Relative to the large social cost of blackouts and the cost of building new transmission it seems likely that the benefits of this technology will far outweigh its costs.

¹¹ See (Apt, 2004) for a related calculation

Table 6.1—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

6.3 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.

7 Conclusions

Cascading failures and blackouts result from violations that persist in the network long enough to trigger the protection system actions. 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 protection system. Experiments also 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 and communication methods, 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 in the United States, transmission owners, regional transmission organizations, and state regulated utilities are responsible to fund 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 over a large portion of a synchronous grid will require substantial coordination among system operators across state boundaries. This is likely to require at least some regulatory oversight. On the other hand our work demonstrates that improving the power network control system, can be done using a decentralized architecture, and therefore does not require large centralized regional transmission operators (RTOs). 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

The authors wish to thank Dong Jia for substantial help and advice in the development of this method.

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.