Controlling cascading failures with cooperative autonomous agents

Paul Hines*

Department of Engineering and Public Policy
Carnegie Mellon University
Pittsburgh, PA 15213, USA
E-mail: phines@cmu.edu
*Corresponding author

Sarosh Talukdar

Department of Electrical and Computer Engineering and
Department of Engineering and Public Policy
Carnegie Mellon University
Pittsburgh, PA 15213, USA
E-mail: talukdar@ece.cmu.edu

Abstract: Cascading failures in electricity networks often result in large blackouts with severe social consequences. A cascading failure typically begins with one or more 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 dependent outages occur. Specifically, our approach places one autonomous software agents at each bus of a power network, each of which is tasked with solving the global control problem with limited data and communication. 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 allowing agents to cooperate can vastly improve system performance.

Keywords: cascading failures; autonomous agents; electrical power networks.

Reference to this paper should be made as follows: Hines, P. and Talukdar, S. (2007) 'Controlling cascading failures with cooperative autonomous agents', *Int. J. Critical Infrastructures*, Vol. 3, Nos. 1/2, pp.192–220.

Biographical notes: Paul Hines is currently a doctoral student in the Engineering and Public Policy Department at Carnegie Mellon University. Hines 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. Hines previously worked with Alstom ESCA Corporation (now Areva T&D) developing neural network load forecasting tools, with Black and Veatch Corporation designing electrical substation hardware and controls, and most recently with the Federal Energy Regulatory Commission studying interactions between Nuclear Power Plants and the transmission system in the USA.

Copyright © 2007 Inderscience Enterprises Ltd.

Sarosh Talukdar is Professor of Electrical and Computer Engineering and Engineering and Public Policy at Carnegie Mellon University. Talukdar 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 optimisation software for electric power circuits.

1 Introduction

In 1895, the Niagara Falls Power Company energised the first high capacity three-phase transmission line, connecting hydroelectric generators at Niagara Falls and consumers 22 miles away in Buffalo, NY. The line operated at 11 kV 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, synchronising 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 9 November 1965, the Northeastern USA 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 *et al.* (2003) and Carreras *et al.* (2004) show that the probability of large blackouts in the USA has a power-law tail (see Figure 1). Systems that have power-law probability distributions can have very high or even infinite expected consequences.

Figure 1 The risk distribution of blackouts in the USA has a fat tail

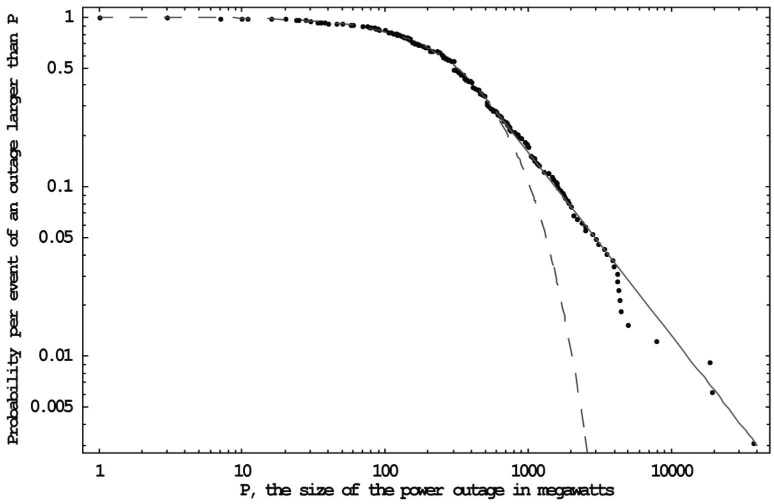

Notes: The dots indicate individual blackouts from NERC data. The dashed line is an exponential distribution (Weibull) fit to the set of blackouts 800 MW or less. The solid line fits the tail of the distribution with a power-law distribution.

Source: Talukdar *et al*. (2003)

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 (Chen *et al.*, 2005; Carreras *et al.*, 2002; 2004; Liao *et al.*, 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 Operations Committee, 1996).

The consequences of blackouts can be quite severe (see Table 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.

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 criteria[5] that most North American system operators use to manage the cascading failure risk under normal operating conditions. In the wake of the 2003 blackouts many in industry, government, and academia advocated that the industry's practice of self-regulation be replaced with a set of binding, enforceable reliability rules. Thus congress passed the Electricity Reliability Act of 2005, establishing that industry create an Electricity Reliability Organisation which would have federally regulated authority to establish and enforce reliability rules for participants in the US electricity network (US Congress, 2005).

Table 1 Several large cascading failures

Date	*Location*	*Notable consequences*
9 November 1965	Northeastern USA, Ontario	30 000 000 customers (20 000 MW) interrupted (USFPC, 1965).
13 July 1977	New York City	9 000 000 customers (6000 MW) interrupted. Widespread looting, chaos. Police made about 3000 arrests. (USDOE/FERC, 1978; Corwin and Miles, 1978).
2 July 1996	Western USA	2 000 000 customers (11 850 MW) interrupted (WSCC Operations Committee, 1996).
3 July 1996	Western USA	The disturbance from 2 July reoccurred. Operators interrupted load to most of Boise Idaho, vastly reducing the extent of the event (WSCC Operations Committee, 1996).
10 August 1996	Western USA	7 500 000 customers (28 000 MW) interrupted. Economic damage estimates range from $1–$3 billion (WSCC Operations Committee, 1996).
25 June 1998	Midwestern USA, Central Canada	152 000 customers (950 MW) interrupted.
November 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 and Pentayya, 2004).
14 August 2003	Midwestern and Northeastern USA, Southeastern Canada	50 000 000 customers interrupted. Estimates of the social costs range from $4 billion to $10 billion. Massive traffic jams in New York City (US-Canada Power System Outage Task Force, 2004; ELCON, 2004).
27 September 2003	Italy	57 000 000 customers interrupted. At least five deaths resulted. 30 000 passengers stranded in trains for hours (BBC, 2003; CNN, 2003; CRE/AEEG, 2004).

Source: NERC, (2005)

1.2 Operating power networks

Power systems are operated with many objectives, including:

- *Economics* – maximise the net economic benefit of service
- *Reliability* – minimise the risk of service interruption
- *Protection* – minimise 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.[6]

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 and LeReverend, 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 and Bertsch, 2001; Novosel *et al.*, 2004; Madani *et al.*, 2004). While almost all SPS designs currently in operation are operated out of a centrally located control center (Anderson and LeReverend, 1996), a few SPS design concepts use a more distributed architecture, though agents are generally organised hierarchically and are dependent on central facilities for planning activities (Jung and Liu, 2001; Jung *et al.*, 2002; Kamwa *et al.*, 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 electro-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, organised 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 and Lipson (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 optimisation 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 optimisation methods (Cohen and Zhu, 1984) can be applied to the Optimal Power Flow (OPF) problem and solved by distributed autonomous agents (Kim and Baldick, 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 and Talukdar, 2004). Another distributed optimisation technique (Modi *et al.*, 2004) organises agents hierarchically to solve discrete optimisation problems. Agent-based technologies have also been applied to the relay protection problem (Yanxia *et al.*, 2002; Coury *et al.*, 2002) and proposed as a means of improving distribution systems (Kueck and Kirby, 2003).

1.4 Cooperation

We define cooperation as the sharing of useful information and the utilisation 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 (1995) and Jennings and Bussmann (2003) discuss cooperative designs for an Energy Management System (EMS),[7] and automobile manufacturing plant control. These papers advocate that agents having clearly defined and known intentions and responsibilities. Camponogara *et al.* (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 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 optimisation 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 centralised applications. (Camponogara *et al.*, 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 optimisation 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 (GP) a formulation of the overall control problem, and therefore the problem that must be solved at each stage of an MPC algorithm.

(GP): Find and implement U_1 such that:

$$\underset{U_1,\ldots,U_K}{\text{minimise}} \quad C\left(U_{0,},U_1,\ldots,U_K\right) \tag{2.1a}$$

subject to:

$$X_{k+1} = \mu\left(U_k,\, X_k,\, \Omega_k\right);\; k = 1\ldots K \tag{2.1b}$$

$$G\left(U_K,\, X_K,\, \Omega_K\right) \le 0 \tag{2.1c}$$

where:

C is a function that evaluates the total social cost of a set of control actions

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, (GP), into sub-problems, (SPO_i), such that:

- Each sub-problem can be assigned to an autonomous agent.
- Each sub-problem is easier to solve than the overall problem.
- 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:

- breaking the control vector into disjoint parts
- making agent-*i* responsible for calculating only the *i*-th portion of the control vector
- allowing agent-*i* to assume that the other agents will calculate their parts optimally
- using predictive models that are simplified so as to de-emphasise 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 $U_k = [Z_{1k}, Z_{2k},.., Z_{Nk}]$, and Z_{nk} is assigned to agent-*i*
- 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 Equation (2.2b) is the sub-problem 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 $Y_{i1},\ldots,Y_{iK}$ and $X_1, X_2,..,X_K$. Simultaneously, solve the optimisation problem:

$$\underset{Z_{i1},\ldots,U_{iK}}{\text{minimise}} \quad C\left(U_0, U_1, \ldots, U_K\right) \tag{2.2a}$$

$$G\left(U_K,\, X_K,\, \Omega_K\right) \le 0 \tag{2.2b}$$

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

$$(SPO_{ik}):\ \underset{\{Z_{ik}, Y_{ik}\}_{1\le k\le K}}{\text{minimise}} \quad C\left(U_0, U_1, \ldots, U_K\right) \tag{2.3a}$$

subject to:

$$X_{ijk+1} = M_{ijk}\left(U_k,\, X_k,\, \Omega_k\right);\ j = 1 \ldots J,\ k = 0 \ldots K-1 \tag{2.3b}$$

$$G\left(U_K,\, X_K,\, \Omega_K\right) \le 0 \tag{2.3c}$$

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: $X_k = \{x_{ijk}\}_{0\leq j\leq J}$

$R_{i0}, R_{i1},\ldots,R_{iJ}$ are concentric and disjoint regions of the network, such that R_{i0} is centred 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 centred 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 (GP). 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 neighbours 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 neighbours, 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:

$$(SPC_{ik})\text{: } \underset{\{Z_{ik}\}x_{1\leq k\leq K}}{\text{minimise}} \; C\left(\hat{U}_0, \hat{U}_1, \ldots, \hat{U}_K\right) \tag{2.4a}$$

subject to:

$$X_{ij1} = M_{ijk}\left(\hat{U}_0,\, \hat{X}_0,\, \Omega_0\right);\ j = 1 \ldots J \tag{2.4b}$$

$$X_{ijk+1} = M_{ijk}\left(\hat{U}_k,\, \hat{X}_k,\, \Omega_k\right);\ j = 1 \ldots J,\ k = 1 \ldots K-1 \tag{2.4c}$$

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

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

3 Global control problem definition

This section adapts the general global problem formulation (GP) 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, Chap.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.

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 behaviour. 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 and Wollenberg, 1996, Chap.13), though the problem can be solved in a variety of ways. In fact, the actions of a capable human operator could constitute a good solution to this problem, though arriving at truly optimal solutions would typically require some computational assistance. This global problem (GP) is given in Equations (3.1a–3.1h) below.

$$\underset{G,L}{\text{minimise}} \sum_{k=1}^{K} \sum_{n \in N} Cost_n \left(G_{nk-1} - G_{nk},\ L_{nk-1} - L_{nk} \right) \tag{3.1a}$$

subject to:

$$\left|V\right|^{\min} \le \left|V_K\right| \le \left|V\right|^{\max} \tag{3.1b}$$

$$\left|I_{nm,K}\right| = \left|y_{nm}\left(V_{nK} - V_{mK}\right)\right| \le \left|I_{nm}\right|^{\max},\ n,m \in N, n \ne m \tag{3.1c}$$

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

$$I_k = Y_{NN} V_k \tag{3.1d}$$

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

$$\mathrm{Re}\left(L_{nk} / L_{n0}\right) = \mathrm{Im}\left(L_{nk} / L_{n0}\right),\ n \in N \tag{3.1f}$$

$$G_n^{\min} \le G_{nk} \le G_n^{\max},\ n \in N \tag{3.1g}$$

$$0 \le L_{nk} \le L_{n0},\ n \in N \tag{3.1h}$$

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 and Wollenberg (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 Equation (3.1a)) are the social costs that would be incurred from the interruption of electrical supply or demand. If operators deem some customers as more valuable than others, the objective function Equation (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 (2.1c). 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 solvers[8] quickly find optimum solutions to this problem for small networks (< 200 buses). Figure 2 shows the result of one such calculation using the IEEE 39 bus test case.

Figure 2 Optimal load and generation shedding actions resulting from the solution of the global problem *P* on the IEEE 39 bus test case

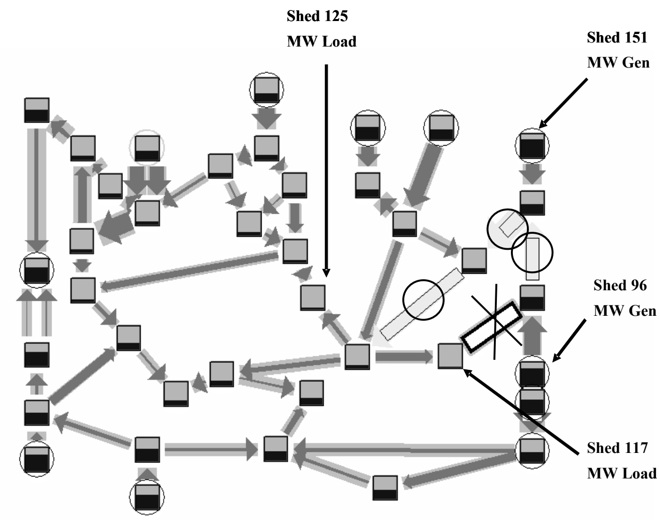

Notes: These are the minimum cost actions that eliminate the violations shown. Branch current violations are marked with circles. These violations occurred after a line outage was applied at the location marked with an X. The arrows indicate power flow magnitude and direction.

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, Equation (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 centralisation that is institutionally impractical. For these, among other, reasons a decentralised 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:

1 place a software agent at each load and generation bus

2 task each agent with controlling only its local load and generation

3 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

4 allow each agent to solve its local problem iteratively using MPC

5 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 3 provides a high level view of this method.

Figure 3 Feedback diagram of the system showing how operators and the agent network interact

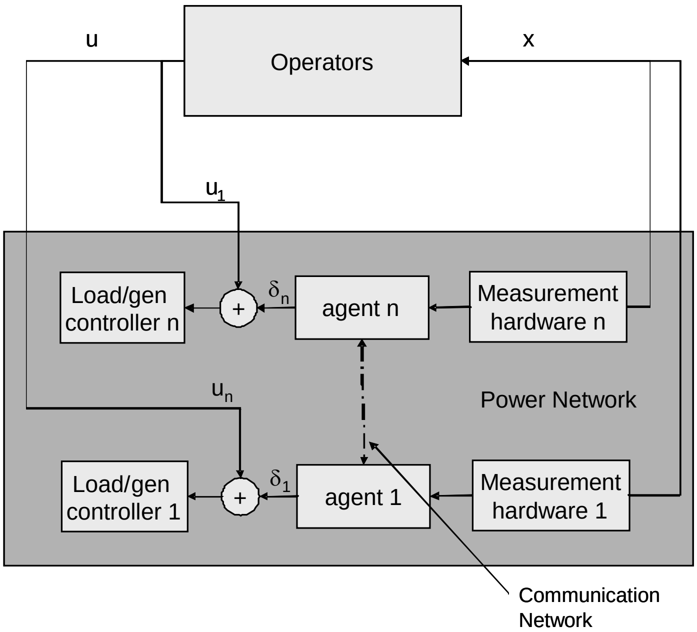

Notes: Under normal conditions operators control load and generator set points (U). When an agent detects a violation, it calculates a plan and implements the local portion of that plan effectively making adjustments (δ_1) to the operator set points.

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 modelling, 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 neighbourhood. 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 neighbourhood 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 2 describes the agent regions and models in more detail, and Figure 4 illustrates the four regions on a graph of the IEEE 118 bus network model.

Table 2 Summary of agent measurement and prediction for the four agent regions

	Measurement			*Prediction*		
Region	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	Constant measurements	Assumed optimal given M_{ijk}	Predicted with M_{ijk}	$\Omega_{ijk+1} = \Omega_{ijk}$
R_2	Estimated	Obtained through cooperation	Obtained 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 Illustration of the spatial decomposition of the overall problem on the IEEE 118 bus network

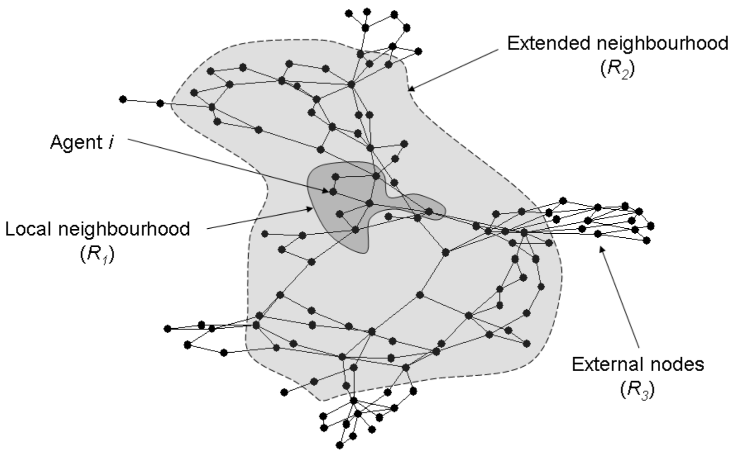

Notes: An agent gathers measurements and watches for violations frequently within its local neighbourhood and occasionally within its extended neighbourhood. The radius of an agent's local neighbourhood is r_1 and the radius of an agent's extended neighbourhood is r_2.

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 2.

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 centre 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):

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

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:

$$\delta_{R_i,k} = u_{R_i,k+1} - u_{R_i,k} = \begin{bmatrix} \mathrm{Re}\left(G_{R_ik+1} - G_{R_ik+1}\right) \\ \mathrm{Re}\left(L_{R_ik+1} - L_{R_ik+1}\right) \end{bmatrix} \tag{4.2}$$

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 Equation (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, Equation (4.1) becomes:

$$\left|I_{R_jk+1}\right| = \left|I_{R_jk}\right| + D_{R_jM}\delta_M \tag{4.3}$$

where M represents the set of all control variables for which the agent has current measurements and D_{R_jM} 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 $\left(Y_{NN} : I_N = Y_{NN}V_N\right)$ and the branch incidence matrix $\left(Y_B : I_B = Y_BV_N\right)$ as follows:

$$D = \operatorname{Im}\left(Y_B\right)\operatorname{Im}\left(Y_{NN}\right)\Lambda \tag{4.4}$$

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:

$$x_{R_j}^{\min} \le x_{R_jK} \le x_{R_j}^{\max}, \tag{4.5}$$

absolute boundaries on the control variables:

$$u_{R_j}^{\min} \le u_{R_jk} \le u_{R_j}^{\max}, \tag{4.6}$$

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

$$\delta_{R_j}^{\min} \le \delta_{R_jK} \le \delta_{R_j}^{\max}. \tag{4.7}$$

Agent-*i* uses this network model in combination with its local optimisation problem (problem SPO_i, given in Equation (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 (SPP_i). The result is an optimisation 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_1) and obtains predictions for the remaining control variables that fall within R_0, R_1 and R_2, and over a time horizon $k = \{1, ...K\}$. This integration results in the sub-problem for agent-*i* at time t_0 given in Equation (4.8) below.

$$\underset{\Delta_M}{\text{minimise}} \sum_{k=1}^{K} e^{-\rho k} {c_M}^T \delta_{Mk} \tag{4.8a}$$

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

$$|I_{Bk}| \cong |I_{Bk-1}| + D_{BM}\delta_{Mk} \leq f\left(I_{B,0 \ldots k-1}\right)|I_B|^{\max} \tag{4.8b}$$

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

$$RR_g \leq \delta_{gk} \leq 0, \; g \in G_M \tag{4.8d}$$

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

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 $\{1 \ldots K\}$

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 Equation (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 current[9]) 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 Equation (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 minimise 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 6 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 5.

Figure 5 Diagram of the two dimensional decision plan calculated by agent-*i*

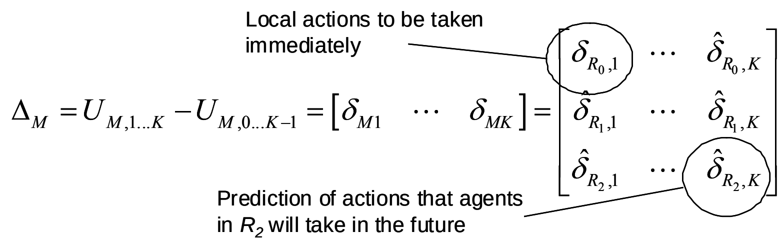

$$\Delta_M = U_{M,1\ldots K} - U_{M,0\ldots K-1} = \begin{bmatrix}\delta_{M1} & \cdots & \delta_{MK}\end{bmatrix} = \begin{bmatrix}\delta_{R_0,1} & \cdots & \hat{\delta}_{R_0,K} \\ \hat{\delta}_{R_1,1} & \cdots & \hat{\delta}_{R_1,K} \\ \hat{\delta}_{R_2,1} & \cdots & \hat{\delta}_{R_2,K}\end{bmatrix}$$

Notes: The agent obtains this control plan by solving Equation (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 Equation (4.8) and acts is not cooperative. Such an agent uses an overlapping objective function Equation (4.8) but does not exchange useful information with its neighbours 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 neighbourhood (R_1). 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 neighbourhood (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 provides a more general description of this cooperation algorithm.

Table 3 Agent algorithm with cooperation

Condition	*Step*	*Action*
No violations detected	1	Collect data from neighbours.
	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.

Note: Unilateral agents skip Steps 6–10

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 and Gan, 1997). One agent is placed at each bus and has the capabilities specified in Section 4. Table 4 summarises the important model input parameters and assumptions.

Table 4 Model input parameters

Input	*Description*
Network data	Modified from the IEEE 118 bus test case (Zimmerman and Gan, 1997)
Load shedding costs	Randomly assigned between \$500/MW and \$1500/MW
Generator shedding costs	Assigned uniformly at \$30/MW
Solution horizon (K)	Four time steps
External neighbourhood radius (r_e)	Ten branches
Local neighbourhood radius (r_l)	Varies between one and six branches
External data estimation error	15% coefficient of variation $\left(\sigma_x / \overline{x}\right)$
Initiating disturbances	Chosen randomly from a set of 100 violation inducing double branch outages

A simulation is initiated by choosing a disturbance, a local neighbourhood 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 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 Simulation procedure

Step	*Action*
1	Choose a disturbance randomly from the set of double contingencies.
2	Choose an internal neighbourhood 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 neighbourhoods.
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 neighbourhoods.
8	Allow the agents to calculate control plans for the control horizon ($k + 1 \ldots K$).
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 4. Each simulation is repeated for agents with and without cooperation. Figure 6 shows the trajectory of the most severe violation resulting from a typical disturbance for both cooperative and unilateral agents. Figures 7 and 8 show the relationship between the quantity of communication (internal neighbourhood size) and two measures of performance: control error and completion time (see Figure 6 for definitions).

Figure 6 Violation trajectories that result from agents reacting to a disturbance in the IEEE 118 bus test case

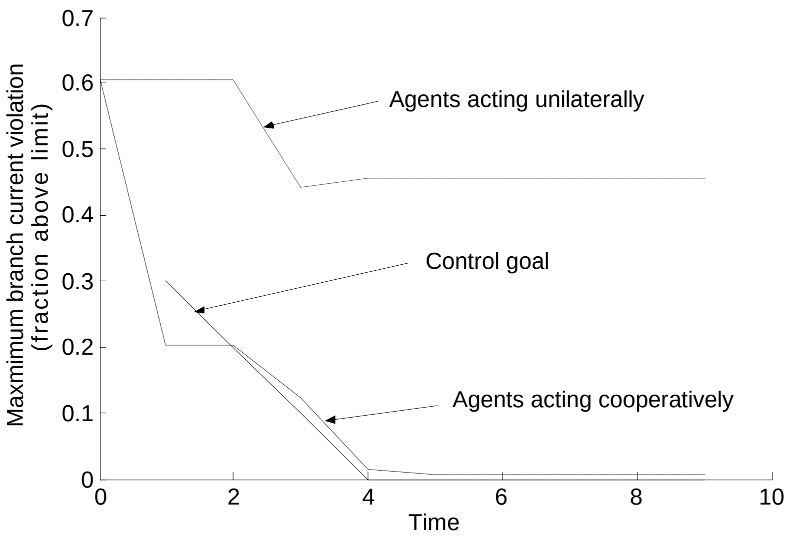

Notes: The agents have a small local neighbourhood ($r_l = 2$). The disturbance consists of outages on branches 8 and 40. The cooperative agents eliminate the violation nearly along the control goal. The control error is the area of the space between the control goal and the actual trajectory. For the 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 four. For the unilateral case, the completion time is set to ten (beyond the solution horizon).

Figure 7 The relationship between r_l and control error (see Figure 6 for definition)

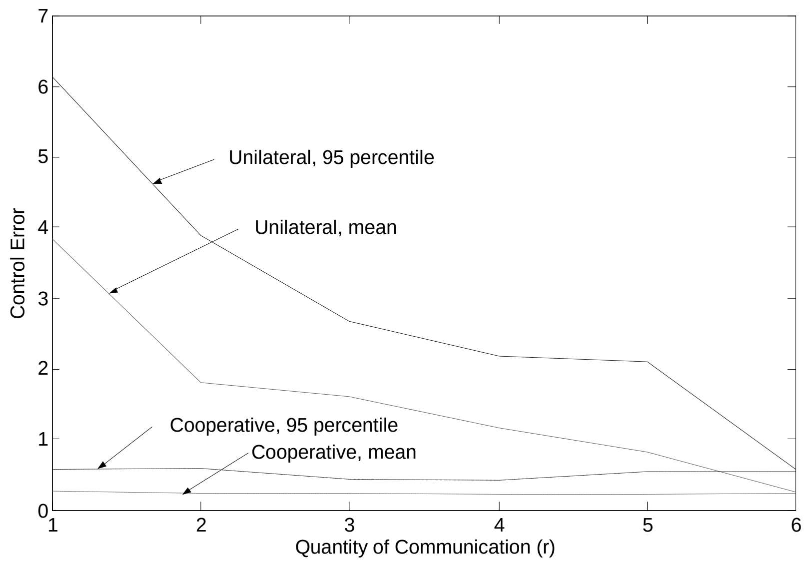

Figure 8 The relationship between r_l and completion time (see Figure 6 for definition)

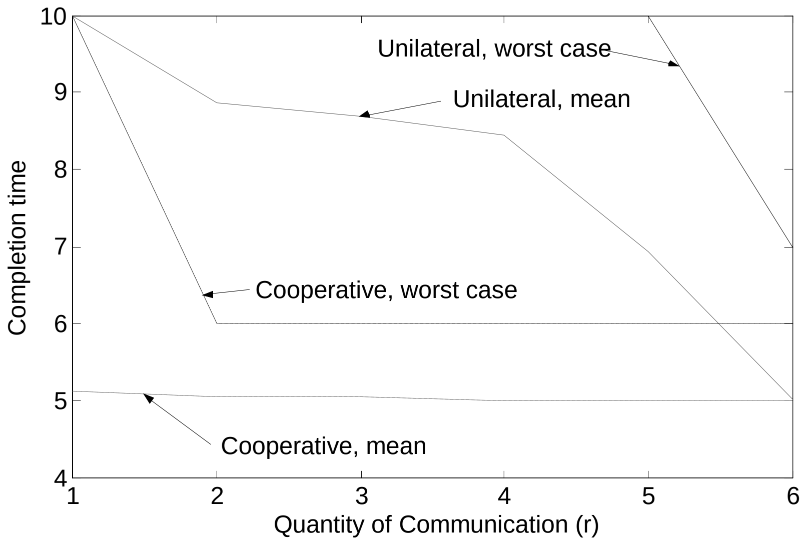

Note: On average the cooperative agents require one additional time step to remove their violations, independent of the local neighbourhood size.

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 (Fumagalli *et al.*, 2004), 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 *et al.*, 2004), but far less than the total cost of service interruptions in the USA (about $80 billion/year according to LaCommare and Eto, 2004). 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 9 illustrates the potential effects of improved grid control and altered operating criteria.

Figure 9 Illustration of the desired benefits of improved grid control

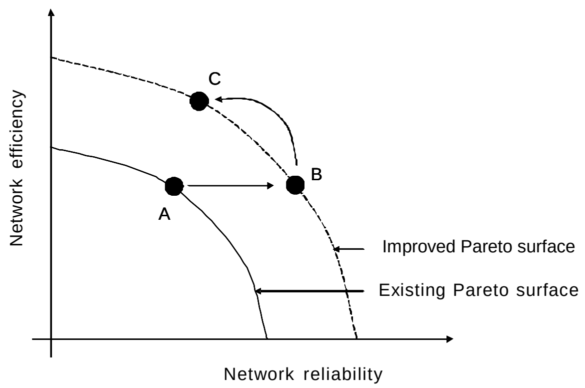

Notes: An effective control scheme will increase network reliability, and thereby create a new Pareto surface of potential tradeoffs between reliability and efficiency. If the existing tradeoff is at Point A, the effect of improved control alone will be to move to Point B along the new Pareto surface. 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).

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 decrease as well. Table 6 shows an order-of-magnitude approximation of the cost associated with implementing this technology for the US eastern interconnect.[10]

Table 6 A preliminary cost estimate for installing a network of autonomous control agents at every node of the eastern interconnect

Agent hardware	(Rack mounted PC)	50 000 buses × \$3000 each = \$150 million
Agent software	(Assumed fixed cost)	\$1,000,000
Installation	(Engineering, labour, *etc.*)	50 000 buses × \$10,000 each = \$500 million
Maintenance		50 000 buses × \$1,000 every five years = NPV: \$120 million
Communications	(For half of all buses)	25 000 buses* × \$3000 each = \$75 million
Total		\$850 million or \$68 million/year

Notes: Assumed discount rate is 7% with a 30 year planning period

*Assumes that half of the buses in the system have an existing communication system that can be connected to.

Naturally, a project of this scope would need to be implemented incrementally, beginning in a few control areas and expanding through the interconnect. Nevertheless, this calculation allows us to compare the relative magnitudes of the benefits and costs. Relative to the large social cost of blackouts and the expense of massive amounts of additional transmission it seems likely that the benefits of this technology could substantially outweigh its costs.

6.3 Risks

One interpretation of the No Free Lunch Theorem of Optimisation 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 and Pepyne, 2001). It is possible that wide-spread implementation of a distributed 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 characterising 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, and to develop the agent designs so that some performance guarantees are possible.

7 Conclusions

Cascading failures and blackouts result from violations that persist in the network long enough to trigger 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 through the use of improved network models and cooperation 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 quantify accurately, and widely distributed. Under current regulation in the USA, transmission owners, regional transmission organisations, and state regulated utilities are responsible for funding transmission investments. Large investments are unlikely to occur unless regulatory bodies ensure that investors can recover their capital investments. Additionally, the installation of a distributed agent network that spans multiple control areas in a synchronous grid will require substantial coordination among system operators. This is likely to require at least some regulatory oversight. On the other hand our work demonstrates that it is possible to implement a SPS for load and generation shedding without the centralised architecture that is typical of such schemes. A decentralised design relaxes the requirement that a large, centrally managed, Regional Transmission Organization (RTO) operate such a system. This is an important result, because it implies that it is possible to solve some global power network problems without centralisation. This property could ease the implementation of a global load and generation shedding scheme where it is institutionally infeasible to organise the grid into large, centrally operated, control areas.

Finally, a coordinated method for determining the social cost of load shedding must be adopted for this method to operate correctly. If operators can set local load-shedding costs unilaterally, neighbouring 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, such as a price-based mechanism that affects individual loads, can be implemented.

Acknowledgements

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

Funding for this work comes from the Carnegie Mellon Electricity Industry Center and ABB Corporate Research in Raleigh, NC.

References

Anderson, P.M. and LeReverend, K.K. (1996) 'Industry experience with special protection schemes', *IEEE Transactions on Power Systems*, Vol. 11, No. 3.

Antsaklis, P.J. (2000) 'A brief introduction to the theory and applications of hybrid systems', *Proceedings of the IEEE*, July, Vol. 88, No. 7.

Apt, J., Lave, L.B., Talukdar, S., Morgan, M.G. and Ilic, M. (2004) 'Electrical blackouts: a systemic problem', *Issues in Science and Technology*, Summer.

BBC (2003) 'Italy launches blackout inquiry', *BBC News Online*, 30 September, news.bbc.co.uk.

Blackburn, J.L. (1998) *Protective Relaying: Principles and Applications*, 2nd ed., New York: Marcel Dekker.

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

Camponogara, E., Jia, D., Krogh, B.H. and Talukdar, S. (2002) 'Distributed model predictive control', *IEEE Control Systems Magazine*, Vol. 22, No. 1.

Carreras, B.A., Lynch, V.E., Dobson, I. and Newman, D.E. (2002) 'Critical points and transitions in an electric power transmission model for cascading failure blackouts', *CHAOS*, Vol. 12, No. 4, pp.985–994.

Carreras, B.A., Newman, D.E., Dobson, I. and Poole, A.B. (2004) 'Evidence for self-organized criticality in a time series of electric power system blackouts', *IEEE Transactions on Circuits and Systems – I: Regular Papers*, September, Vol. 51, No. 9.

Chen, J., Thorp, J.S. and Dobson, I. (2005) 'Cascading dynamics and mitigation assessment in power system disturbances via a hidden failure model', *International Journal of Electrical Power and Energy Systems*, May, Vol. 27, No. 4, pp.318–326.

CNN (2003) 'Italy recovering from big blackout', *CNN*, 28 September, www.cnn.com/world.

Cohen, G. and Zhu, D.L. (1984) 'Methods in large scale optimization problems', *Advances in Large Scale Systems*, Vol. 1.

Commission de Regulation de l'Energie (CRE, France) and Autorita per l'energia elettrica e il gas (AEEG, Italy) (CRE/AEEG) (2003) *Report on the Events of 28 September Culminating in the Separation of the Italian Power System from the other UCTE Networks*, http://www.pserc.org/Resources.htm#European%20Blackouts.

Corwin, J.L. and Miles, W.T. (1978) 'Impact Assessment 1977 New York City Blackout', *Systems Control, Inc. (Palo Alto) Report*, Prepared for Lester H. Fink, US Department of Energy, July.

Coury, D.V., Thorp, J.S., Hopkinson, K.M. and Birman, K.P. (2002) 'An agent-based current differential relay for use with a utility intranet', *IEEE Trans. on Power Delivery*, January, Vol. 17, No. 1.

Durfee, E.H. (1999) 'Practically coordinating', *AI Magazine*, Vol. 20, No.1.

Electricity Consumers Resource Council (ELCON) (2004) *The Economic Impacts of the August 2003 Blackout*, 9 February, www.elcon.org/Documents/EconomicImpactsOfAugust2003Blackout.pdf.

Fink, L.H. and Carlsen. K. (1978) 'Operating under stress and strain', *IEEE Spectrum*, March.

Fisher, D.A. and Lipson, H.F. (1999) 'Emergent algorithms: a new method for enhancing survivability in unbounded systems', *Proc. of the 32nd Hawaii Int. Conf. on System Sciences*, IEEE.

Fumagalli, E, Black, J.W., Vogelsang, I. and Ilic, M. (2004) 'Quality of service provision in electric power distribution systems through reliability insurance', *IEEE Transactions on Power Systems*, August, Vol. 19, No. 3.

Hines, P. and Talukdar, S. (2004) 'Reducing the costs of disturbances to the electric power network', *Presented at the Int. Assn. for Energy Economics Meeting*, Washington, DC, June.

Hines, P., Jia, D. and Talukdar, S. (2004) 'Distributed model predictive control for electric grids', *Proc. of the Carnegie Mellon Transmission Conf.*, Pittsburgh, December, www.ece.cmu.edu/~tanddconf_2004.

Hines, P., Liao, H., Jia, D. and Talukdar, S. (2005) 'Autonomous agents and cooperation for the control of cascading failures in electric grids', *Proc. of the IEEE Conf. on Networking, Sensing, and Control*, Tuscan, March.

Ho, Y-C. and Pepyne, D.L. (2001) 'Simple explanation of the no free lunch theorem of optimization', *Proc. of the 40th IEEE Conf. on Decision and Control*, Orlando, December.

Jennings, N.R. (1995) 'Controlling cooperative problem solving in industrial multi-agent systems using joint intentions', *Artificial Intelligence*, Vol. 75, No. 2.

Jennings, N.R. and Bussmann, S. (2003) 'Agent-based control systems,' *IEEE Control Systems Magazine*, Vol. 23, No. 3.

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

Jung, J., Liu, C., Tanimoto, S.L. and Vittal, V. (2002) 'Adaptation in load shedding under vulnerable operating conditions', *IEEE Transactions on Power Systems*, November, Vol. 17, No. 4.

Kamwa, I., Grondin, R. and Herbert, Y. (2001) 'Wide-area measurement based stabilizing control of large power systems – a decentralized/hierarchical approach', *IEEE Transactions on Power Systems*, February, Vol. 16, No. 1.

Kim, B.H. and Baldick, R. (2000) 'A comparison of distributed optimal power flow algorithms', *IEEE Trans. on Power Systems*, May, Vol. 15, No. 2.

Kueck, J.D. and Kirby, B.J. (2003) 'The distribution system of the future', *The Electricity Journal.*

LaCommare, K.H. and Eto, J.H. (2004) 'Understanding the cost of power interruptions to U.S. electricity consumers', *Report No. LBNL-55718*, Ernest Orlando Lawrence Berkeley National Laboratory, September.

Liao, H., Apt, J. and Talukdar, S. (2004) 'Phase transitions in the probability of cascading failures', *Proc. of the Carnegie Mellon Transmission Conf.*, Pittsburgh, December, www.ece.cmu.edu/~tanddconf_2004.

Madani, V., Adamiak, M. and Thakur, M. (2004) 'Design and implementation of wide area special protection schemes', *Proc. 57th Annual Texas A&M University Conference for Protective Relay Engineers*.

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

Neil, C.E. (1942) 'Entering the seventh decade of electric power', *Edison Electric Institute Bulletin,* September.

NERC (2005) 'Examples of major bulk electric system power outages', *NERC*, www.nerc.com. 'BlackoutTable.pdf' accessed January.

Novosel, D., Begovic, M.M. and Madani, V. (2004) *IEEE Power and Energy Magazine*, January–February, Vol. 2, No. 1.

Rawlings, J.B. (2000) 'Tutorial overview of model predictive control', *IEEE Control Systems Magazine*, Vol. 20, No. 3, pp.38–52.

Rehtanz, C. and Bertsch, J. (2001) 'A new wide area protection system', *IEEE Power Tech Conf. Proc.*, September.

Roy, A. and Pentayya, P. (2004) 'Experience of blackouts and restoration practices in the western region of India', *Slides Presented at the IEEE Power Engineering Society General Meeting*, Denver, July.

Talukdar, S.N., Apt, J., Ilic, M., Lave, L.B. and Morgan, M.G. (2003) 'Cascading failures: survival versus prevention', *The Electricity Journal*, November.

United States Congress (2005) *Energy Policy Act of 2005*, April, http://energycommerce .house.gov/.

US Department of Energy: Federal Energy Regulatory Commission (USDOE/FERC) (1978) *The Con Edison Power Failure of July 13 and 14, 1977*, Washington, DC: US Government Printing Office, June.

US Federal Power Commission (USFPC) (1965) *Northeast Power Failure: November 9 and 10, 1965*, Washington, DC: US Government Printing Office, 6 December.

US-Canada Power System Outage Task Force (2004) *Final Report on the August 14, 2003 Blackout in the United States and Canada*, April.

Vassell, G.S. (1991) 'Northeast blackout of 1965', *IEEE Power Engineering Review*, January.

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

WSCC Operations Committee (1996) *Western Systems Coordinating Council Disturbance Report*, WSCC, September.

Yanxia, C. Yin, X., Zhang, Z., Chen, D., Zeng, X. (2002) 'Research on protective relaying systems based on multi-agent system', *Proceedings of International Conference on Power System Technology*, PowerCon.

Zimmerman, R.D. and Gan, D. (1997) 'MATPOWER User's Manual, Version 2.0', *PSERC*, December, www.pserc.cornell.edu/matpower.

Notes

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 *et al.*, 2005, Antsaklis, 2000).

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

3 Within one minute the western grid had separated into five islands (WSCC Operations Committee, 1996).

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 and Carlsen, 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 criterion, in short, requires that a system be operated such that no single contingency will affect a cascading failure.

6 This coincides with recommendation 21 from the 14 August 2003 blackout report, in which the authors recommend that US system operators, "Make more effective and wider use of system protection measures" (US-Canada Power Outage Task Force, 2004).

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

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

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

10 See (Apt *et al.*, 2004) for a related calculation.