Reducing the Costs of Disturbances to the Electric Power Network

Paul Hines[*] and Sarosh Talukdar[**]

Electric power networks are operated by human and mechanical agents, each with some autonomy. This paper describes two methods for coordinating the efforts of mechanical agents (relays, voltage regulators, and governors) using software agents distributed throughout the power system in order to reduce the social costs of disturbances. As evidenced by recent blackouts in Europe and North America, relatively minor disturbances can initiate cascading failures with catastrophic effects. We propose a method to limit such cascades, and thereby reduce the social costs of the resulting blackouts. The proposed methods have several important features. Firstly, they allow for varied interruption costs to be assigned to loads, thereby decreasing the social costs of the required load shedding. Secondly, the methods are distributed and cooperative, giving benefits in actuation speed and scalability and increasing the ability of the system to react well to problems that occur near seams between control areas.

The mechanical agents considered here are situated at the nodes of the network. Each has control over a few variables, such as the terminal voltage of a generator or the load fed from a substation. After a disturbance, software agents communicate and decide how to adjust their control-variables in order to minimize the spread and cost of the disturbance. They do this by solving a local optimization problem while cooperating (exchanging valuable information) with their neighbors. Preliminary results suggest that cooperation schemes exist such that locally optimum solutions can approach globally optimum solutions. That is, autonomous agents working with information from their local neighborhoods can find solutions that are nearly globally optimal.

1. INTRODUCTION

Blackouts are costly. Cost estimates of the August 14, 2003 North American blackout range between 4 and 10 billion USD (USCPSOTF, 2004). The September 2003 blackout in Italy led to three deaths: one from a traffic accident at an intersection where the lights had failed and two because elderly individuals fell down stairs in the dark (CNN, 2003). Additionally, blackouts occur relatively frequently. In the United States, a blackout large enough to darken half a million homes, occurs about once every four months (Apt et al., 2004). Large blackouts occur significantly more frequently than one would predict from a normal probability distribution (Talukdar et al., 2003).

From the above data, we can estimate that the expected yearly cost of large blackouts in the United States is about 1 billion USD. In the United States transmission grid there are approximately 10,000 high voltage transmission nodes. If a solution existed that could cut the frequency of large blackouts in half, but had no other benefits, its yearly cost would have to be

[*] Doctoral Student, Carnegie Mellon University, Engineering and Public Policy Department, Pittsburgh, PA 15213, USA, e-mail: phines@cmu.edu
[**] Professor, Carnegie Mellon University, Department of Electrical and Computer Engineering, Pittsburgh, PA 15213, USA, e-mail: talukdar@ece.cmu.edu

significantly less than 50,000 USD per node to be cost effective. Therefore solutions that require large quantities of expensive high voltage hardware (FACTS devices, large quantities of new transmission) are not likely to be cost effective. However, solutions that improve the use of existing hardware and controllers by making the controllers more cooperative could be cost effective.

This research draws from several areas in the existing literature. In the power systems literature, control methods intended to increase the system's ability to handle emergency conditions, are generally referred to as Special Protection Schemes (SPS) or system protection measures. While some system operators choose not to rely on such methods (PJM, the system operator in the Mid-Atlantic region of the US, for example), adoption of such schemes is widespread with generally positive industry experience. Anderson and LeReverend (1996) surveyed 111 system operators currently using SPS. They found that existing systems had fairly high effectiveness and reasonably low failure rates. It is interesting to note that all of the documented schemes were pre-programmed, rule-based schemes. Rules such as, "if line X trips and the system is at or near state Y, shed load at bus Z" are common. Additionally, with few exceptions, the schemes currently in use in industry require the involvement of a central control center.

The primary objective of this work is to create means of enabling software agents without global knowledge, to solve network control problems and act in a way that approaches global optimality in real time. It draws from existing work on multi-agent systems and distributed optimization. Durfee (1999) gives a useful general discussion of coordinating the work of distributed intelligent agents. Talukdar and Camponogara (2000) present a method of solving a network control problem using agents distributed through a network. They found that cooperating agents can effectively solve network control problems near optimally. The distributed optimization method described in section 3 is an adaptation of the distributed optimal power flow method presented by Kim and Baldick (2000).

This paper is organized as follows: section two presents the global optimal load and generation shedding (OLGS) problem, section three discusses methods of solving this problem using distributed controllers, section four presents some results demonstrating the functionality of the proposed methods, and section five outlines some preliminary conclusions and directions for future work.

2. CASCADING FAILURES

Power system operators are generally expected to operate their system according to what is commonly known as the "N-1 security criterion." The North American Reliability Council defines this expectation as follows, "All control areas shall operate so that instability, uncontrolled separation, or cascading outages will not occur as a result of the most severe single contingency" (NERC, 1998). A single contingency is a failure or outage of a single component (generator, transmission line, transformer, etc.) of the network. While this criterion is useful as a general guide, it is not a sufficient means of preventing cascading failures. Even when this criterion is being followed, the failure of two components simultaneously can trigger a cascading failure. Additionally, due to the voluntary nature of compliance with NERC policies, regulatory bodies have little ability to enforce compliance with this and other reliability policies. The result is a system that is vulnerable to cascading failures.

Figure 1 shows the effect of a transmission line outage on a heavily loaded system. Since a single contingency results in a cascading failure, the pre-contingency state violates the "N-1" criterion.

Figure 1 – Cascading failure demonstration using the IEEE 39 bus test system. Arrow widths indicate flows through transmission lines and the boxes indicate network nodes with voltages. After the fourth stage voltage collapse occurs, likely causing a blackout over the entire system.

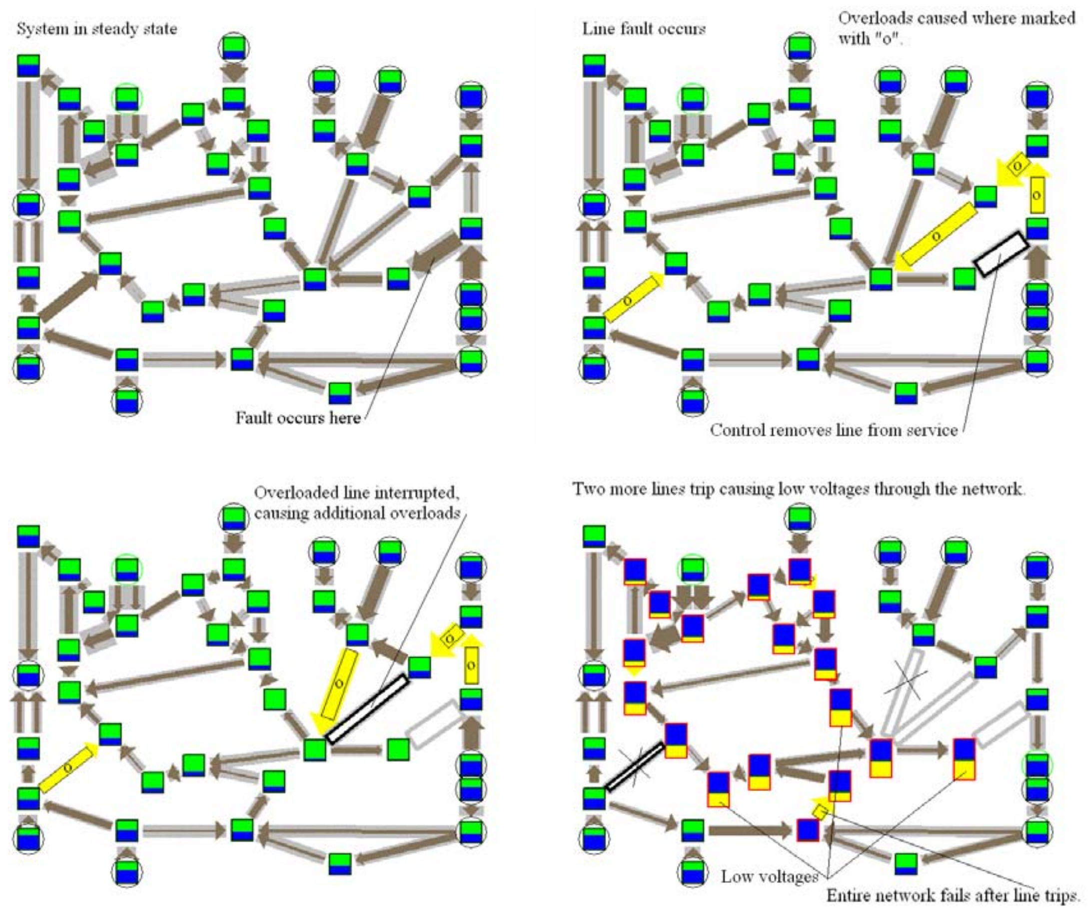

3. THE OPTIMAL LOAD AND GENERATION SHEDDING PROBLEM

This section presents the optimal load and generation problem as used in this paper. For the sake of this paper, we define the problem as follows: minimize the total cost of shedding firm load and changing generator outputs, while ensuring that the resulting node voltages and line currents are within some pre-specified limits.

In this section the following symbols and operators are used:

N	The index set of all nodes in the system, n indicates a single node
NM	The index set of all branches in the system; nm indicates a connected node pair
G	A complex vector of generator electric power output (supply) at each node/bus in the network. For nodes with no generation, $G_n = 0$.

L	A complex vector of energy consumption rates (demand) at each node. For nodes with no load, $L_n = 0$.
$C(\Delta G_n)$	The cost of changing the generation at node n by quantity ΔG_n
$C(\Delta L_n)$	The social cost of changing the load at node n by quantity ΔL_n
V	A complex vector of node voltages. The voltage at node n is V_n.
I	A complex vector of current injections for all nodes. The current injection at bus n is therefore I_n.
I_{nm}	The current injected into the transmission line or other branch between nodes n and m at node n
Y	The complex, n by n, matrix representing the network configuration (node and line admittances)
y_{nm}	The admittance of the branch between nodes n and m
$x(0)$	The pre-contingency value of x
$\|x\|$	The absolute value of x
$Re(x)$	The real portion of x
$Im(x)$	The imaginary portion of x
$conj(x)$	The complex conjugate of x

Using the above symbols the OLGS problem can be written as follows:

$$\underset{L,G,V}{\text{minimize}} \sum_{n\in N} C(\Delta G_n) + C(\Delta L_n) \qquad \text{Control action cost} \qquad (1a)$$

subject to:

$$I = YV \qquad \text{Node current injection} \qquad (1b)$$

$$G_n - L_n = V_n conj(I_n), \qquad \text{Conservation of power} \qquad (1c)$$

$$\operatorname{Im}(L_n) = \frac{\operatorname{Im}\left(L_n(0)\right)}{\operatorname{Re}\left(L_n(0)\right)} \operatorname{Re}(L_n), \qquad \text{Constant power factor at loads} \qquad (1d)$$

$$G_n^{\min} \le G_n \le G_n^{\max}, \qquad \text{Generation output limits} \qquad (1e)$$

$$L_n^{\min} \le L_n \le L_n(0), \qquad \text{Load control limits} \qquad (1f)$$

$$\left|V_n\right|^{\min} \le \left|V_n\right| \le \left|V_n\right|^{\max}, \ \forall n \in N \qquad \text{Node voltage limits} \qquad (1g)$$

$$\left|I_{nm}\right| = \left|y_{nm}(V_n - V_m)\right| \le \left|I_{nm}\right|^{\max}, \ \forall nm \in NM \qquad \text{Line current limits} \qquad (1h)$$

The result is a relatively simple sparse mathematical programming problem that can be solved using standard optimization tools. The problem is mostly linear, with a few quadratic terms, and could be implemented in any existing control center with good state estimation and control abilities. With minimal effort the problem could be augmented to include additional controllers such as transformer tap changers, FACTS devices, and static VAr compensators.

The system would likely operate as follows. A disturbance occurs, such as a sudden line or generator outage. In response the energy management software at the control center updates its network model, runs the above optimization problem, and sends control signals to load and generator controllers.

The system operators could determine the cost functions for both load and generation shedding by a variety of means. Most system operators should be able to determine costs associated with shedding generation in their system without much difficulty. In a regulated system, the system operator should have full knowledge about the costs of emergency controls. In a restructured network, the generators available for this type of emergency control would likely be those participating in regulation or reserve markets. The costs associated with emergency control actions could be determined through adjustments to existing regulation and spinning reserve markets.

The costs associated with load shedding are somewhat more difficult to determine. For the system operator in a regulated control area, load values can be assigned based on priority lists and using Value of Lost Load (VOLL) estimates for different customers. Many utilities already have load shedding priority lists for their area. Load values should be assigned in relation to customers' willingness to pay for avoided interruptions. Reasonable estimates for the value of unserved load are 1 to 3.30 USD / kWh for residential customers, and 2 to 47 USD / kWh for commercial & industrial customers (Graves, 2003).

If the system was able to disconnect every individual load in the network, every load could potentially be assigned a different value. In such a system it is feasible to structure energy tariffs such that customers with higher willingness to pay for reliability could pay more for energy, and would be assigned a higher load value. This higher load value would lead to a lower probability of being disconnected in the case of an emergency. It can be shown that this type of system can lead to an increase in transfer capability in the transmission network.

4. DISTRIBUTED METHODS OF SOLVING THE OLGS PROBLEM

While the above method can interrupt many cascading failure sequences, it has several shortcomings. Firstly, it relies on a central energy management system that is vulnerable to failure. Computer failure was one of the factors that led to the August 14, 2003 blackout (USCPSOTF, 2004). Secondly, it is dependent on the operation of communication channels between the control center and the distributed mechanical agents. Thirdly, and perhaps most importantly, it is only useful if the effects of the disturbance can be mitigated by control actions entirely within a single control area. Due to the interconnected nature of electricity networks, and especially with the increase of long-distance transfers, it is unreasonable to assume that failures, and the resulting control actions, will be contained within a single control area. With this in mind we present two methods of enabling software agents, distributed through the system, to solve the OLGS problem. The first method we present in some formal detail, whereas the second is discussed in conceptual form only.

The distributed optimal power flow method

For the first method, we adapt the distributed optimal power flow method presented by Kim and Baldick (2000) to the optimal load and generation shedding problem. For this method to work, the network must be divided into small groups of interconnected nodes. For testing purposes we use a version of the IEEE 39 bus system that is divided into three areas each with 10 to 15 nodes. This division is shown in figure 2. A software agent is located in each area. The agent is able to monitor the state (voltages and currents) of and send control signals to actuators at each node. Additionally, the agent can exchange information with the area-agents in neighboring areas.

Figure 2 – The IEEE 39 bus, 3 area system

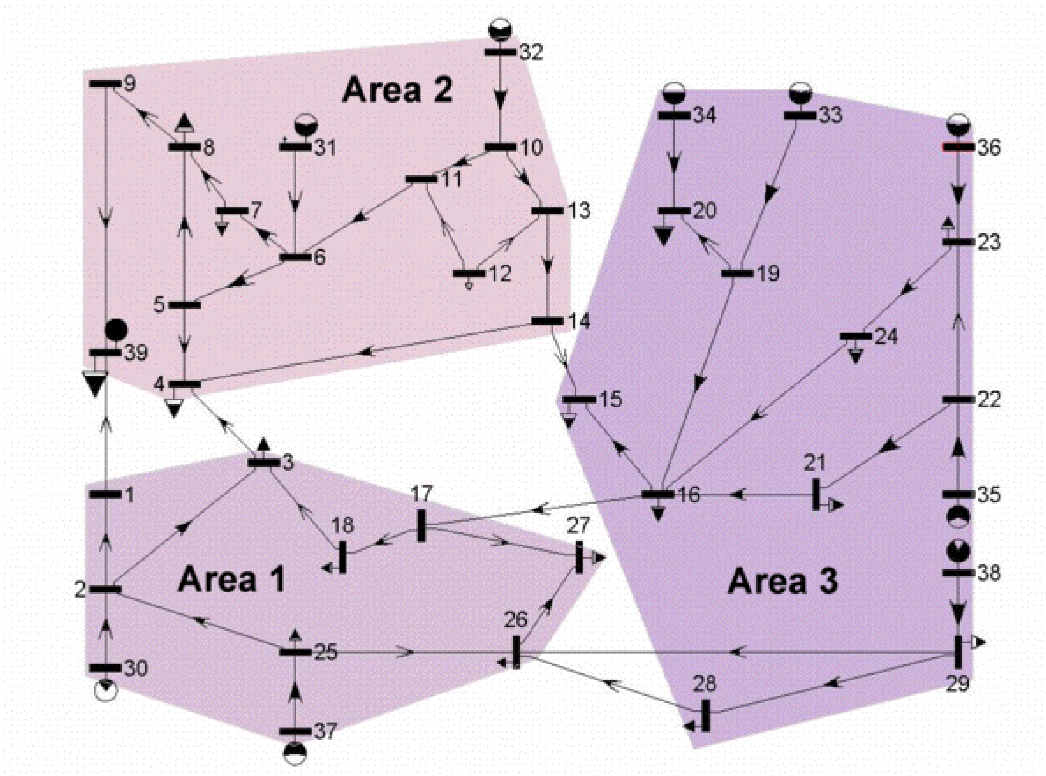

When a disturbance occurs in area controlled by an agent, it immediately solves an optimization problem with only local variables, and exchanges a subset of the results with its neighbors. The neighbors then iteratively solve sub-problems and exchange answers until the solutions converge. Once an agent determines that its solution has converged, it can execute the required control action.

This problem decomposition method uses a penalty function technique, derived from the Augmented LaGrangian Method. It requires the creation of artificial nodes on the borders between the agent areas. Each agent carries its own version of the voltage at the border nodes. The sub-problem assigned to each area-agent is nearly the same as the full OLGS problem (1). The differences are the addition of the coordination penalty function $D(y_A)$, and that only the constraints within the agent's neighborhood are considered. In the following, capital A indicates the index set of nodes within agent a's local area and NM_A indicates the set of branches within the area (including the branches connected to the artificial border nodes). Also, y_A is a vector of the voltage magnitudes and phase angles at the artificial nodes.

$$\min_{L_A,G_A,V_A,y_A} f_A = \sum_{n\in A}\left(C(\Delta L_n)+C_n(\Delta G_n)\right)+D(y_A) \qquad (2a)$$

subject to:

$$I_A = Y_A V_A \qquad (2b)$$

$$G_n - L_n = V_n conj(I_n), \qquad (2c)$$

$$\mathrm{Im}(L_n) = \frac{\mathrm{Im}\left(L_n(0)\right)}{\mathrm{Re}\left(L_n(0)\right)} \mathrm{Re}(L_n), \tag{2d}$$

$$G_n^{\min} \le G_n \le G_n^{\max}, \tag{2e}$$

$$L_n^{\min} \le L_n \le L_n(0), \tag{2f}$$

$$\left|V_n\right|^{\min} \le \left|V_n\right| \le \left|V_n\right|^{\max}, \; \forall n \in A \tag{2g}$$

$$\left|I_{nm}\right| = \left|y_{nm}(V_n - V_m)\right| \le \left|I_{nm}\right|^{\max}, \;\; \forall nm \in NM_A \tag{2h}$$

Different formulations for the function $D(y_A)$ exist. We use what is known in the literature as the Auxiliary Problem Principle formulation as defined in (3) and the update function (4).

$$D(y_A) = \frac{\beta}{2}\left\|y_A - y_A(k)\right\|^2 + \gamma y_A^{\;T}\left(y_A(k) - y_B(k)\right) + \lambda(k)^T y_B \tag{3}$$

where y_A is the vector of voltages at the border nodes, y_B is the vector of border node voltages according to the neighbor agent(s), λ is a vector of penalty multipliers associated with the border variables, and *(k)* indicates the value of a variable after the previous iteration (assuming that the current iteration is k+1). For a general discussion of this formulation and convergence proofs, see (Losi, 2003). After an agent and each of its neighbors solve (2) they share the results for the border variables (y_A) and then update the multiplier vector using the difference between the border node voltage estimates:

$$\lambda(k+1) = \lambda(k) + \alpha\left(y_A(k+1) - y_B(k+1)\right) \tag{4}$$

The method requires the use of three tuning parameters (α, β, and γ). The convergence conditions for the method require that $\alpha < 2\gamma < \beta$. We found that the combination $\alpha=1.5$, $\beta=2$, $\gamma=1$, works well in most cases. We also found that the local objective function may need some scaling in order to obtain good results with this method.

This Auxiliary Problem Principle has been shown to effectively solve large optimal power flow problems divided into many sub-networks (Kim and Baldick, 2000). In section 5 we present several representative results demonstrating that this method can be used to find solutions to the OLGS problem, that are often near optimal.

Overlapping areas with linearized sub-problems

The above method is useful, but has some limitations. The convergence properties of the algorithm are quite sensitive to the tuning parameters, and optimal results are not guaranteed. Additionally, the method depends on the proper functionality of each area-agent during each iteration. A method that placed a software agent at each node, with control over only those decision variables at that node, would increase the autonomy of the controllers. In general, higher levels of autonomy in control lead to decreased actuation time and increased robustness to failures.

The following is a conceptual description of a method that would eliminate the need for area control agents and allow the software agents to be co-located at the same node as the actuators. Each agent maintains a relatively large model of the network. If the network is small, the agent can maintain a model of the entire network. During normal operation the agent exchanges data with agents in its neighborhood to maintain a relatively accurate network model. Preliminary experiments show that the solution sensitivity due to data errors decreases with distance from the control variables. The agent can be designed to gather data most frequently from nearby nodes, and with decreasing frequency from nodes further from the agent. With the addition of some estimation and learning capabilities, agents should be able to obtain good estimates of remote variables with minimal remote communication.

Figure 2 – Overlapping network model. Each agent in the system maintains a model of a portion of the network. The agent shares data with agents in it neighborhood to keep the model updated.

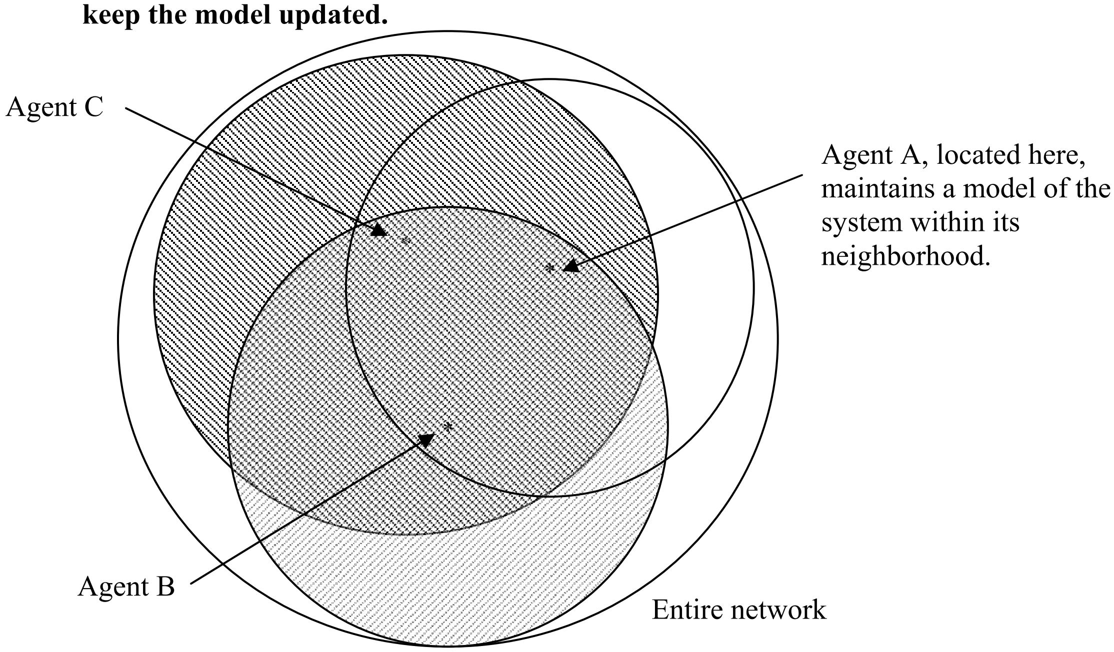

When a disturbance occurs in an agent's neighborhood, it solves a linearized version of the OLGS problem for its network model. The decision space of the sub-problem includes all control and state variables within the agent's network model. After solution, the agent signals the local mechanical agents to act. It may be desirable to enact only a portion of the solution and measure the system response before subsequent actions.

One of the important features of the design is the linearization of the agent sub-problems. Due to the difference nature of the OLGS objective function, the problem is well suited to a linearized formulation. The linear sub-problems should be trivial to solve with limited computing capabilities. In future work we plan to formalize this design and study its properties.

5. DEMONSTRATIVE RESULTS

In order to test the above problem, trials were run on the IEEE 39 bus test system shown in figure 2. The network includes nine generators and 19 loads. For case 1 and 3 we assign all the loads a uniform value of $1000. Similarly assume that the cost of shedding generation is uniform at $30 / MW. In this section we present the results of three tests with the above system.

Case 1 shows the result of the loss of a single transmission line (between nodes 21 and 22) on the heavily loaded system, as shown in figure 1. Case 2 shows the result of losing the same line but with the load value at bus 21 increased to $10,000 / MW. Case 3 shows the result of a double contingency; the loss of lines 21-22, and 6-11 simultaneously

Table 1 – Globally optimal results

	Case 1 Line 21-22 out		Case 2 Line 21-22 out, node 21 load value increased		Case 3 Lines 21-22 and 6-11 out	
Node	Re(ΔL)	Re(ΔG)	Re(ΔL)	Re(ΔG)	Re(ΔL)	Re(ΔG)
4			16.2	0		
8	0	0	0	0	104.8	0
15	125.6	0	227.3	0	29.3	0
21	117.7	0	0	0	112.4	0
35	0	150.7	0	150.8	0	151.2
36	0	95.7	0	95.7	0	96.2
Costs	\$243.3k	\$7.299k	\$243.6k	\$7.307k	\$246.5k	\$7.394k
Total	\$ 250,700		\$ 250,900		\$ 253,900	

Table 2 -- Area-agent results

	Case 1 Line 21-22 out		Case 2 Line 21-22 out, node 21 load value increased		Case 3 Lines 21-22 and 6-11 out	
Node	Re(ΔL)	Re(ΔG)	Re(ΔL)	Re(ΔG)	Re(ΔL)	Re(ΔG)
15	0	0	0	0	212.7	0
16	17.2	0	58.8	0	0	0
21	232.0	0	193.0	0	35.7	0
32	0	0	0	0	0	2.6
35	0	150.6	0	151.6	0	150.5
36	0	95.6	0	96.9	0	95.2
Costs	\$249,200	\$7,386	\$1,989,000	\$7,455	\$248,400	\$7,449
Total	\$256,600		\$1,996,000		\$255,900	
Agent Penalty	\$5,100 (2.35 %)		\$1,745,000 (696 %)		\$2000 (0.759 %)	

From tables 1 and 2 we can make several observations. Firstly, we observe that the OLGS problem, as formulated, is solvable with standard optimization tools. If appropriate voltages and current thresholds are chosen this method could likely stop many cascading failure sequences, thereby reducing the costs of cascading failures. Secondly, Case 2 demonstrates that the global OLGS method is able to avoid shedding high value loads when specified, further reducing the social costs of power outages. Thirdly, we observe that, while the area-agent method generally finds satisficinig solutions, it does not necessarily find optimal solutions. In the uniform load value case, the algorithm's solutions were very nearly optimal, but in the case of the high-value load, the algorithm only reduced the amount of shedding at the high value load, instead of not shedding any load at all at this location. This is likely because of the steepest descent method used to update the multipliers. It is well known that steepest descent methods have poor convergence properties as they approach optimal solutions.

6. CONCLUSION

One of the recommendations from the task force studying the August 14th blackouts was that the electricity industry "make more effective and wider use of system protection measures." When well designed, system protection measures can reduce the spread of cascading failures. Most current system protection technologies are rule-based and therefore apply only to the system for which they were designed. Additionally, current technologies are primarily centralized, and therefore only useful to an entity which can observe and control large portions of the grid. While regional transmission organizations which observe and control large portions of the grid are helpful, failures will frequently extend beyond the view and reach of the area. Additionally, as system operators become larger, the difficulty involved in optimally mitigating failures increases substantially.

The methods presented in this paper make progress toward designing distributed network agents such that locally optimum actions are near globally optimum. We implemented a method based on the Auxiliary Problem Principle, and show that this method can produce near globally optimum solutions, but does not do so under all conditions.

In future work we plan to formalize the linearized overlapping area method and study its properties. We also hope to study different public policy approaches to improving the national transmission control system.

WORKS CITED

Anderson, P.M. and LeReverend, B.K. "Industry Experience with Special Protection Schemes." *IEEE Transactions on Power Systems*. May 1996.

Apt, Jay, Lester B. Lave, Sarosh Talukdar, M. Granger Morgan, and Marija Illic. "Electrical Blackouts: A Systemic Problem." *Issues in Science and Technology.* Summer 2004.

(CNN) "Italy recovering from big blackout." [http://www.cnn.com/]. Sept. 28, 2003.

Durfee, Edmund H. "Practically Coordinating." AI Magazine. Spring, 1999.

Graves, Frank and Lisa Wood. "Economic Cost of the August 14th 2003 Northeast Power Outage: Preliminary Estimate." The Brattle Group. August 18, 2003.

Kim, Balho H. and Ross Baldick, "Course-Grained Distributed Optimal Power Flow." *IEEE Transactions on Power Systems*. May 1997.

Kim, Balho H. and Ross Baldick, "A Comparison of Distributed Optimal Power Flow Algorithms." *IEEE Transactions on Power Systems*. May 2000.

(NERC) North American Electricity Reliability Council. "Operating Policies: Policy 2 – Transmission." January 5, 1998, available from: http://www.nerc.com/.

Talukdar, Sarosh, Jay Apt, Marija Ilic, Lester B. Lave, and M. Granger Morgan. "Cascading Failures, Survival versus Prevention." The Electricity Journal. Nov. 2003.

Talukdar, Sarosh and Eduardo Camponogara. "Callaborative Nets." *Proceedings of the 33rd Hawaii International Conference on System Sciences.* 2000.

(USCPSOTF) U.S.-Canada Power System Outage Task Force, "Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations." [http://www.ferc.gov/]. April 2004.