

Reliability Analysis of Bulk Power Systems Using Swarm Intelligence

David G. Robinson, Ph. D., Sandia National Laboratories

Key Words: Bulk power, contingency analysis, reliability, artificial intelligence

SUMMARY & CONCLUSIONS

This paper documents research into the use of an adaptive cultural model and collective intelligence as a means of characterizing the reliability of bulk power networks. Historically, utilities support the reliable design and operation of bulk power networks through first-order contingency analysis. In contingency analyses the list of candidate elements for disruption are identified by engineers *a priori* based on the rate at which the elements failure through the course of normal grid operation.

The new method, an implementation of particle swarm analysis, a swarm of 'virtual power engineers' successfully identified the set of network elements which, if disrupted, would possibly lead to a cascading series of events resulting in the most wide spread damage. The methodology is technology independent: it can be applied on not only for reliability analysis of bulk power systems, but also other energy systems or transportation systems. The methodology is scale neutral: it can be applied to power distribution networks at the local, state or regional level.

1. INTRODUCTION

The objective of this research was the development of tools and techniques for the identification of critical nodes within the national bulk power system. These are nodes that, if disrupted would cause the most widespread, immediate damage.

Traditional contingency analyses performed by utilities are single point contingency analyses, focusing on identifying the single most critical element. In addition, analyses performed by utilities focus primarily on those elements which have a *naturally occurring* high failure rate, typically generation. Substations, transmission lines, etc. have low failure rates and so have low likelihood of inclusion in traditional investigations performed by utilities. The possibility of SCADA failures are also not considered in traditional analyses and are clearly vulnerable points. Finally, synergistic effects of multiple, simultaneous damage nodes that can amplify the impact are not considered in traditional contingency analyses. In rough orders of magnitude, there are 10,000 potential points of attack in the western U.S. grid, 45,000 points in the north-east and 5,000 points in the Texas area. These are only the major generation and transmission elements and do not include command and control elements or elements that might be critical to a particular region.

Through the identification of critical elements and the

quantification of the consequences of their failure, site/node specific vulnerability analyses can be focused at those locations where additional security measures could be effectively implemented. In particular, with appropriate sizing and placement within the grid, distributed generation in the form of regional power parks may reduce or even prevent the impact of widespread network power outages. Even without additional security measures, increased awareness of sensitive nodes can provide a basis for more effective national, state, and local emergency planning. Locations and types of critical nodes can be used to preposition spares, deploy security forces, or be points where additional site security measures can be employed.

Identification of critical nodes or points of vulnerability within such a large, complex system is a daunting computational task. This research was focused on those situations where simultaneous, multiple points within the system would be attacked. To overcome the computational difficulties associated with traditional methods of vulnerability analysis, an artificial intelligence (AI) method was developed and applied to a variety of bulk power test cases constructed by the Institute of Electrical and Electronic Engineers (IEEE). The approach has a foundation in a branch of cultural psychology that can be used to model adaptive group behavior similar to that observed in flocks of birds and schools of fish.

2. BACKGROUND

Historically, utilities identify single points of vulnerability through a first-order contingency analysis. The list of candidate elements for disruption are identified *a priori* typically based on the rate at which the elements fail through the course of normal grid operation. Based on their naturally occurring failure rates, elements are randomly chosen from the list and removed from service. The reaction of the grid to the disruption is analyzed using a computer model of the power redistribution that results. Reliability indices are collected and the simulation of contingencies continue until convergence is reached.

The number of contingencies can explode rapidly: for a simple system of 69 potential points of failure, there are nearly 6×10^{12} different possible contingencies that could occur. As noted previously, bulk power systems may involve on the order of tens of thousands of potential points of failure. For very large systems, the list is typically dominated by generation elements since transmission and substation failures are rare under normal bulk power system operation. The response of the bulk power system to these first-order

contingencies is the basis for scheduling maintenance, setting reserve margins, emergency planning, etc. Unfortunately, as evidenced by the outages in 1996 and 2003, transmission lines can easily be home to initiating events that cascades quickly.

The objective of this research was to explore alternative, nontraditional methods of identifying critical elements in a bulk power system to improve contingency analysis. With the computational power available today, enumeration of all possible single point contingencies is a possibility. However, nonlinearities and synergistic effects preclude simply exploring and ordering all first-order contingencies.

The focus of this research was the development of a methodology that can be used to identify the second-, third-, and higher order contingencies. Contingency analysis has been the focus of much research the past few years but the emphasis has remained on naturally occurring failures rates. It was decided that, rather than duplicating the direction of these research efforts, a more theoretical systems approach would be emphasized.

A trio of very different approaches was investigated. The first involves complexity theory and network theory in particular. Network theory has been applied extensively on a wider variety of complex networks including communication networks, the Internet, and even power systems. The goal of these efforts has been to characterize the robustness, fragility, and attack tolerance of complex networks. Large networks can be typically classified into two groups: homogeneous and non-homogeneous. Nodes in homogeneous networks typically have roughly the same number of connections. Random graphs and small-world networks are examples of homogeneous networks. Nonhomogeneous, or scale-free networks are largely homogeneous, but contain a few nodes that are highly connected. The Internet and World Wide Web are examples of scale-free networks. Scale-free networks typically result from an evolving system with preferential points of interconnection. The robustness (or inversely, the fragility) of large networks is characterized by their diameter, the average length of the shortest path between two nodes. The fragility of a network is investigated as highly connected elements are removed and the diameter of the network changes.

A second approach investigated involves characterizing the bulk power system as a complex series of n -dimensional polyhedra, similar in many ways to a large crystal. The vulnerable points on the grid expose themselves as 'cleavage' points in the crystalline structure. Polyhedral dynamics, a branch of set theory that deals with the topological relationships between finite sets, is employed to relate the various elements of the power network to a simplicial structure and to identify the vulnerable points in the network.

The final and most promising research focused on the use of an AI-based approach founded in cultural psychology. The swarming behavior of flocks of birds and schools of fish are used as a basis for finding the optimum combination of network nodes to be disrupted to cause the most damage.

This paper focuses on the most promising of these approaches. The fundamentals of particle swarm optimization are presented and the implementation is then discussed and the results from a number of test cases are provided. Finally,

conclusions of the research and recommendations for future efforts are presented. For additional details the reader is referred to the full report [1].

3. PARTIAL SWARM OPTIMIZATION

The fundamental concepts associated with particle swarms were developed in 1995 by Kennedy and Eberhart [2]. Particle swarm optimization (PSO) is regarded as being of the family of evolutionary strategies for problem solving. Other members of this family include, for example, genetic algorithms and evolutionary programming. While heavily influenced by the philosophy of evolutionary strategies, PSO differs significantly from these "survival of the fittest" algorithms in that it is based on a social cooperative perspective: individuals working with others in a common social group to solve problems. Contrary to the algorithms in Darwinism-based paradigms, individuals are not replaced by better performing individuals; rather, the individuals in a swarm model adapt to the environment by gathering information and processing that information as a group. In a swarm model it is not the individual who changes, but rather the knowledge of the individual that changes from iteration to iteration.

It should be noted that PSO is also closely related to the area of cellular automata, which is a discrete time, discrete state virtual machine where the current state of each cell of the system is determined by its most recent past and the states of those cells in the immediate proximity.

Since 1995 PSO has been used by a number of authors to address a wide variety of applications including optimization of reactive power and voltage control for a bulk power system [4,5]. Swarm intelligence is particularly suited for our problem due to its evasion of local minima.

As noted by Kennedy and Eberhart [140, p288], in a very simple sense an individual reacts and adapts to their environment, including other individuals, through three major principles: evaluating, comparing, and imitating. Individuals can identify desirable goals and objectives of their social group within the environment and have their own perception of their behavior relative to environment. They can also compare their behavior to the behavior of other individuals in achieving those goals and then imitate the behavior of those individuals who are seen as having behavior conducive to achieving those goals. By adapting in this fashion, individuals take advantage of the *experiences* of those individuals around them in much the same fashion as a school of fish takes advantage of the many eyes available to the group to warn of danger and the subsequent reaction of the group to avoid the danger.

Stepping quickly away from the metaphors, let an individual with a certain behavior set be described as a particle with a certain *position*. The change in the behavior of the individual as it seeks to imitate the behavior of successful individuals is characterized by the *velocity* of the particle.

A particle is distinguished by its:

- current position and velocity,
- value of that position,

- best position achieved thus far,
- best current position achieved by those particles in its neighborhood.

A swarm is characterized by a set of particles and one or more neighborhoods describing the social structures of those particles. Each particle in the swarm changes position and velocity through a combination of its own past best experiences as well as the best experiences of its neighbors. This ability to gain and gather *experiences* individually and from the neighborhood provides the individual particle with *memory*. This capacity for memory is important since it allows the algorithm to exploit information via a local search (through the experience of each individual) and it also emphasizes exploration of the search space with a global search (through the combined experiences of the neighboring particles). This balance of local exploitation and global exploration results in a very robust search algorithm.

The traditional swarm equations take the form [129]:

$$v_{i,t+1} = c_1 v_{i,t} + c_2 (p_{i,t} - x_{i,t}) + c_3 (p_{g,t} - x_{i,t}) \quad (1)$$

$$x_{i,t+1} = x_{i,t} + v_{i,t+1} \quad (2)$$

where:

$x_{i,t}; v_{i,t} \equiv$ position; velocity of particle i at time t

$p_{i,t} \equiv$ position of best performance of particle i through time t

$p_{g,t} \equiv$ position of best performance of group through time t

$c_i \equiv$ coefficients

The meaning behind each of the terms in the above equations will be discussed in the following section, but to tie this all together, let us use the example of a flock of birds. The specific objective of this ‘social’ group is to minimize the distance between themselves and a source of food (such as a cornfield).

2. APPROACH

In the following discussion, the analogy between a team of engineers and a swarm will be exploited. It is important to understand that the analogy results from a serendipitous situation and not from any attempt to actually model the social behavior of a group of engineers. The algorithm operates quite distinctly from the analogy; however, the analogy provides a unique vehicle for discussion purposes.

In our application of a swarm paradigm the swarm will consist of a number of engineers linked together into a loose social structure (i.e. a team) with the goal of finding the combination of power grid elements, which if disrupted, would cause maximum disruption to the national bulk power system. A particle in the swarm will equate to an engineer and a swarm neighborhood or social network will equate to an engineering team. Note that this can be generalized further in the sense of having an engineering team modeled as a member of a larger organization.

The ‘position’ of each engineer is analogous to the choice of power grid elements each engineer has made from a long list of potential elements while ‘the ‘velocity’ of each engineer relates to the probability of the engineer choosing a

particular element. Each engineer on the team will have access to education, training and a variety of independent information sources. This knowledge base will be periodically queried and a decision on the suggested best course of action will be provided to the individuals. The phrase ‘suggested’ is used since there is a certain degree of interpretation and free will that lend uncertainty to the actual course of action taken by the individual engineer.

In general, the position of the particle at a particular time is a continuous variable. However, in our situation, the positions $x_{i,t}, p_{i,t}, p_{g,t}$ can take on only binary values $\{0,1\}$. The individual best $p_{i,t}$ will take on values of 1 if the individual best performance occurred when position $x_{i,t}=1$ and similarly, $p_{i,t}$ will take on values of 0 if the individual best performance occurred when position $x_{i,t}=0$. Following the example of Kennedy and Eberhart [129], we will assume that the velocity $v_{i,t}$ represents the probability that the position takes a value of 1. The probability of the null position $x_{i,t}=0$ is therefore $1-v_{i,t}$. The change in position is




Figure 1 Sigmoid Function

then given by evaluating: if $(rand() < S(v_{i,t}))$ then $x_{i,t}=1$; else $x_{i,t}=0$.

The transform expression $S(v_{i,t}) = 1/[1 + \exp(-av_{i,t})]$ is controlled by the slope parameter a , where the slope at the origin is $a/4$. Typically, from an application point of view, this sigmoid function is limited over the range $[-v_{max}, v_{max}]$ (Figure 1). This prevents the velocity from being driven to zero too quickly and forces exploration of new positions.

With the above formulation, the analysis can proceed equally from two directions with the overall goal of causing as much damage as possible. First, we can take the perspective of the engineering team composed of individual engineers. In this case, conceptually, the term $c_2(p_{i,t} - x_{i,t})$ represents the individual’s contribution of knowledge to the overall objective of the team. This knowledge may consist of such things as personal experience, unique training or specific educational background. On the other hand, the term $c_3(p_{g,t} - x_{i,t})$ represents the contribution of the individual engineers’ knowledge to the collective knowledge of the complete power grid analysis team including the team goals and objectives.

Alternatively, it is possible to formulate the problem as a engineering organization with particular goals and objectives,

$c_3(p_{g,t} - x_{i,t})$, composed of committees/teams with their own unique knowledge base to draw upon, $c_4(p_{c,t} - x_{i,t})$. Finally, we can extend the velocity equation to account for all three levels of social dynamics:

$$v_{i,t+1} = c_1 v_{i,t} + c_2 (p_{i,t} - x_{i,t}) + c_3 (p_{g,t} - x_{i,t}) + c_4 (p_{c,t} - x_{i,t}) \quad (3)$$

(In the following analysis, only the first three terms are considered.) In all cases, the coefficients $c_j, (i \neq j)$ represent the value placed on the level of contribution of each social segment (individual, team, or group) to achieving the objective. Typically the contribution level can change dynamically as knowledge is lost/gained/obscured in the course of the search for the optimum course of action. These coefficients are therefore treated as random variables that are re-evaluated at each stage of the analysis. Any alternative course of action is therefore a weighted average of the individual best course and the group best course of action, e.g.

$$\frac{c_2 p_i + c_3 p_g}{c_2 + c_3}$$

The coefficient c_1 on the velocity term represents the momentum toward change in achieving the objective. In the simplest situation, the desire to achieve a particular goal or objective remains constant throughout the search for the best scenario. However, it is realistic to assume that this momentum may be greater the further away from the hoped for best solution and become smaller as the accumulation of individual and group knowledge begins to focus the alternative courses of action into the one that best in achieves the objective.

3. IMPLEMENTATION

3.1 Buzzard Software

The above equations for particle swarm optimization have been implemented in software. The Buzzard is a software program that acts as an instigator to computer models of large complex systems by introducing a failure in a set of components, e.g. a contingency. The algorithm embedded within Buzzard is independent of the particular system model or infrastructure, but the most recent application has been for Buzzard in the power system contingency analysis.

Buzzard uses the previously discussed AI-based swarm theory algorithm to develop a set of scenarios for disrupting the system. These scenarios are introduced into the system and incite a reaction from the system. The reactions that result are observed by Buzzard and a new set of scenarios are constructed automatically by Buzzard to stimulate the disrupting the critical nodes of the system. These new scenarios are constructed in an evolutionary fashion such that Buzzard seeks new and more effective provocations to disrupt the system.

The complexity of the scenarios is predetermined by the user along with the particular measures that characterize the impact of the scenarios on the system.

The evolutionary strategies within Buzzard differ significantly from approaches that commonly apply genetic-based algorithms as a basis for their search algorithms. Contrary to the algorithms in Darwinism-based paradigms, individuals are not replaced by better performing individuals. Rather, the individuals within Buzzard adapt to the environment by gathering information and processing that information as a group. In this approach it is not the individual who changes, but rather the knowledge of the individual that changes each time a new scenario is generated.

In addition, unlike genetic-based algorithms, the algorithms within Buzzard are less susceptible to being trapped within local minima. Buzzard algorithms are all coded in C/C++ and are scalable to the particular size of system being analyzed.

3.2 Power Flow Software

Implementation involved interfacing the Buzzard software directly with an actual power flow simulation program. This would provide the capability to observe (within the constraints of the simulation model) the impact of disrupting the power system. However, this presented some difficulties. Given the vast number of contingency scenarios to be investigated, the computational burden would still be substantial.

As an approximation, it was decided to make a number of simplifying assumptions. First, since a complete characterization of network reliability measure is not needed, only deterministic performance measures need to be considered. Second, after a network disruption, only the very immediate impact on the power flow in the grid would be characterized and collected for each scenario.

Since many performance assessment results must be compared during contingency analysis, there is a need to reduce the voluminous output of a power grid simulation to a manageable number of performance measures. All formal contingency assessments involve comparing a single index, or multiple indices, against some simple numeric standard. Also relevant to contingency analysis is the identification of failure criteria. These failure criteria include capacity deficiency, line overload, system separation with load loss, bus isolation with load loss, voltage collapse, MVAR limit violations, and non-convergent situations (which surrogate network instabilities). When a contingency fails, either an index is greater than some critical value or is outside of some believed-stable region of index values, or some failure criterion is met. These performance indices can themselves be the direct product of a contingency analysis model, without any intermediate derivation of more precise information such as power flow calculations, at the expense of precision and accuracy. Indices that avoid full power flow calculations to determine post-contingency voltage levels at each bus have shown promise in reducing computation time while still supporting ranking and screening procedures.

Such an approach was chosen here; the measure chosen, line voltage and current over-rating, is commonly used in contingency analysis of bulk power systems. Line over-rating is expressed as a percentage of the allowable load, either

No. of Nodes	Team Size	Penalty Cost	Momentum	Vmax	Iterations	Cost	Critical Node Set	First Top 7	Alt Cost
3	5	10000	1.0	5	175	90103	31,29,40	175	90103
	7	10000	1.0	5	125	90095	31,32,42	125	90095
	9	10000	1.0	5	3525	90342	31,33,40	450	90105

No. of Nodes	Team Size	Penalty Cost	Momentum	Vmax	Iterations	Cost	Critical Node Set
3	7	10000	1.0	5.0	150	90151	31,40,42
	7	10000	1.0	5.5	150	89095	31,35,63
	7	10000	1.0	6.0	150	88359	9,31,51

Table 2. Summary of Investigations for Various Momentum, V_{max} Values

voltage or current, that is placed on the system. Under normal operation, a line rating of 100% is typical. It was felt that the change in line over-rating immediately subsequent to a disruption would provide at least a qualitative measure of the severity of the disruption.

To compare contingencies, a single performance measure or cost function was developed: the sum of all line over-ratings which exceed a particular criteria. Critical over-ratings vary slightly from area to area (105%-110%) but for the purposes of this study a single criterion is used. Unless specifically noted otherwise, a limit of 110% is used as the critical level for all the cases discussed.

To characterize the performance of a network before and after disruption, two open source power flow packages were employed. The first was developed by New Mexico State University under contract to Sandia National Laboratories. The second is a product with a long history that has been developed by the Bonneville Power Association (BPA). Over the past 20 years, BPA has been actively involved with the development of power system analysis software. In 1991, BPA, in partnership with WECC and the Electric Power Research Institute (EPRI), began development of an enhanced power flow package referred to as the Interactive Power Flow (IPF) program. Both the NMSU and BPA/IPF program were used in the analyses.

3.3 Test Case

The IEEE 300 bus test system was chosen as the basis for the initial investigation. The IEEE 300 bus test case was initially developed by the IEEE Test Systems Task Force in 1993 based on data from a northeast power pool. The particular data set used in this analysis is available from the University of Washington Power System Test Case Archive. The site provides World Wide Web access to power system data (test cases) and is maintained by Richard D. Christie, a Professor at the University of Washington, Seattle, Washington, USA (christie@ee.washington.edu). The system consists of three connected regions as depicted in Figure 2 with 69 generators and 298 busses, transformers, etc. available for disruption.

Cost	Critical Set: 3 node scenario		
342	31 (191)	40 (236)	33 (213)
151	31 (191)	40 (236)	42 (239)
105	31 (191)	40 (236)	41 (238)
104	31 (191)	40 (236)	29 (187)
095	31 (191)	32 (198)	42 (239)
082	31 (191)	39 (233)	42 (239)
050	31 (191)	40 (236)	28 (186)
030	31 (191)	32 (198)	43 (241)

Table 1. Optimal 3 Node Sets: Truth

4. APPLICATION

For the initial investigation it was decided that a subset of the IEEE 300 RTS would be sufficient. A reduced test system was developed, focusing only on the 69 generators in the network. This network is depicted in Figure 2, with the generators identified and numbered. (Generators 65-69 are included in the computer model as reserve generators.)

The number of possible node combinations is incredibly large even when limiting the analysis to the 69 generators. When limiting the allowable number of potential critical nodes to 2 there are 2,346 possible contingencies, while for 3 critical nodes the possible scenarios explodes to 52,394 and for 9 there are 56,672,074,888. Validation and verification of the algorithms is obviously very difficult and verges on being computationally impossible.

The focus is on identifying the best combination of modeling parameters to identify the best 2 or 3 critical nodes since it was possible to actually enumerate all possible node combinations of these sizes. Table 1 presents the results of an exhaustive search for the best combination of 3 nodes from a set of 69 possible nodes. The cost (damage) associated with each node combination is provided as well as the names and bus reference number (in parentheses) of each node.

(It is also important to note that the set of seven solutions summarized in the above tables reflect (effectively) multiple,

optimal solutions, since the differences between the costs are negligible. In retrospect, the choice of the standard IEEE 300 RTS was unfortunate. There are a large percentage of network elements that can have very similar operational impact on the performance of the system. In addition, these sets differ in their value (cost) by only roughly 5%. This made validation a bit more challenging than would be expected from an actual bulk power system.)

Those nodes belonging to a critical set of 3 are highlighted in Figure 2. In addition to identifying sets of critical nodes, the physical location of the nodes in the network can provide insight into sensitive areas within the network. Table 2 presents various summaries for scenarios involving a engineering teams of size 3, 7, and 9 and a potential set of 3 nodes. The number of iterations was artificially capped at 150; no exhaustive attempt at identification of the final optimum set was conducted.

It became clear that by reducing the momentum factor, the convergence to the optimum set of nodes is slower. The number of potential critical nodes with high probability of selection as a critical nodes is still rather large for a momentum of 0.90 and decreases rapidly as the momentum factor is raised to 0.95 and then finally to 1.0. *This can be useful if, rather than attempting to identify a specific set of X critical nodes, it is simply desired to identify a larger set of important critical nodes perhaps with the intent of identifying vulnerable regions within a bulk power network.*

Alternatively, the impact of v_{max} on the identification of the optimal set critical nodes is less distinct. For a specific

momentum, in this case 1.0, the cost function quickly focuses on the selection of 3 nodes. Larger values of v_{max} allow the search for the optimum to extend over a broader region of support and the algorithm is less likely to be 'stuck' in a local minimum.

Figure 3 provides a snapshot of the velocities for a few of the 69 potential nodes from a typical simulation. Recall that the velocity of a node is a measure of the probability that the node will be selected for membership into the set of X critical nodes. The velocities provide insight into not only the optimum set of nodes but also can be used to identify other nodes who are just outside the selection criteria. These nodes may warrant additional attention when intangible selection criteria are included, e.g. military or economic value of node.

5. CONCLUSIONS

It is clear that the Buzzard software coupled with a traditional power flow analysis program can be used to identify critical elements within large complex bulk power systems. The algorithms are consistent with traditional methods that identify critical single point contingencies in the sense that the new approach can also be used to characterize single point contingencies. For small numbers of potential attack sites (e.g. 2-3) on relatively small systems (69 nodes) it was possible to enumerate all possible contingencies and in every case the nodes identified through enumeration corresponded to the nodes identified using the Buzzard algorithm.

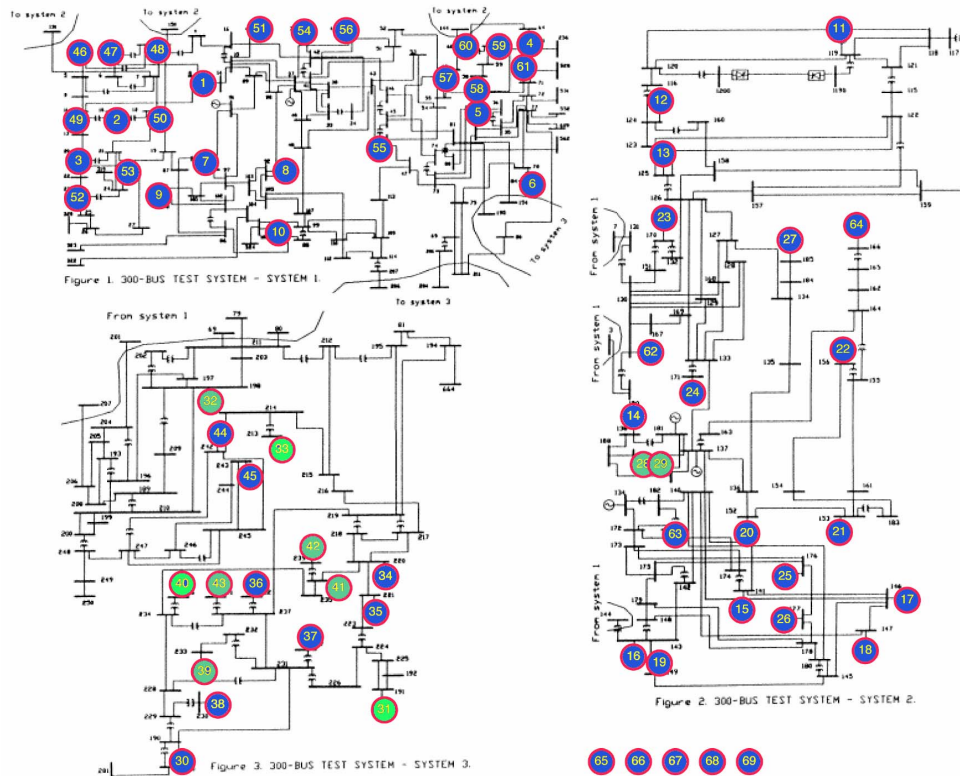


Figure 2: Node Designations for Simple Example




Figure 3 Final Velocities for Team Size 7, $V_{max}=5.0$, Momentum=0.9

The software very quickly identifies one of the multiple 'best' solutions. As mentioned previously, for the given cost function, there are a number of best solutions that are very close in value. The algorithm typically finds a solution that is within 0.3% of the value of the true optimum, but may require 5000 iterations to find the final 'best' combination of critical nodes.

The existence of multiple, sub-optimal solutions with very similar total critical node values poses a bit of a dilemma: it is important to be aware of similar 'optimal' solutions, but it clouds identification of the 'best'. It is suggested that a small number of additional searches be conducted with various initial seed values before identifying a specific set of critical node sets. In addition, the velocity vector provides considerable insight into the existence of these potential members of the optimal set. High residual velocities at the completion of the simulation are key indicators of potential optimal set membership.

Two key considerations that need to be understood and possibly investigated in a more formal fashion in future efforts: sensitivity to v_{max} and the penalty associated with exceeding the allowable number of critical node locations. The choice of v_{max} impacts the search algorithm by constraining the search to be either more locally focused or allowing the search to extend to a more global solution space. Typically, $3 < v_{max} < 6$, with smaller values being associated with local search and larger values allowing the search to broaden. Penalties over the range of 1000 to 10,000 were used to force the number of selected critical nodes to be approximately the user specified values. High penalties coupled with low values of v_{max} resulted in lengthy simulations until convergence.

ACKNOWLEDGEMENTS

This research was conducted with support from the Laboratory Directed Research and Development program at Sandia National Laboratories. Sandia is a multi-program laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000. The author wishes to particularly thank Lee Eubanks, GRAM, Inc. who did the majority of the programming associated with the above analyses.

REFERENCES

1. D. Robinson, R. Cox, "Vulnerability of Critical Infrastructures: Identifying Critical Nodes," SAND 2004-2696, Sandia National Laboratories, Albuquerque, NM (Jun) 2004.
2. R. Eberhart, J. Kennedy, "A New Optimizer Using Particle Swarm Theory," Int. Symp. Micromechatronics and Human Science, Nagoya, Japan.
3. H. Yoshida, K. Kawata, Y. Fukuyama, S. Takayama, and Y. Nakanishi, "A particle swarm optimization for reactive power and voltage control considering voltage security assessment," *Transactions of the Institute of Electrical Engineers of Japan, Part B*, vol. 119/B, pp. 1462-9, 1999.
4. H. Yoshida, K. Kawata, Y. Fukuyama, S. Takayama, and Y. Nakanishi, "A particle swarm optimization for reactive power and voltage control considering voltage security assessment," *IEEE Transactions on Power Systems*, vol. 15, pp. 1232-9, 2000.

BIOGRAPHY

David G. Robinson, PhD,
Risk and Reliability Analysis Department
Sandia National Laboratories
PO Box 5800, MS 0748
Albuquerque, NM 87185 USA

e-mail: drobin@sandia.gov

Dave Robinson is currently a Distinguished Member of the Technical Staff at Sandia National Laboratories in Albuquerque, NM. He earned his doctoral degree in System Engineering from the University of Arizona in 1986, a Masters Degree in System Engineering from the Air Force Institute of Technology in 1981 and a BSME in Mechanical Engineering in 1977 from Colorado State University. Dr. Robinson's research interests focus on Bayesian methods for the quantification of uncertainty in complex systems.