7-27-2007

Infrastructure Hardening: A Competitive Co-evolutionary Methodology Inspired by Neo-Darwinian Arms Races

Travis Service

Daniel R. Tauritz
Missouri University of Science and Technology, tauritzd@mst.edu

William M. Siever

Recommended Citation

T. Service et al., "Infrastructure Hardening: A Competitive Co-evolutionary Methodology Inspired by Neo-Darwinian Arms Races," *Proceedings of the 31st Annual International Computer Software and Applications Conference (COMPSAC'06) (2007, Beijing)*, Institute of Electrical and Electronics Engineers (IEEE), Jul 2007.
The definitive version is available at https://doi.org/10.1109/COMPSAC.2007.129

Infrastructure Hardening: A Competitive Coevolutionary Methodology Inspired by Neo-Darwinian Arms Races

Travis Service, Daniel Tauritz, William Siever
Department of Computer Science
University of Missouri–Rolla
Rolla, MO, USA 65409
{tcsvw5,tauritzd,bsiever}@umr.edu

Abstract

The world is increasingly dependent on critical infrastructures such as the electric power grid, water, gas, and oil transport systems, which are susceptible to cascading failures that can result from a few faults. Due to the combinatorial complexity in the search spaces involved, most traditional search techniques are inappropriate for identifying these faults and potential protections against them. This paper provides a computational methodology employing competitive coevolution to simultaneously identify low-effort, high-impact faults and corresponding means of hardening infrastructures against them. A power system case study provides empirical evidence that our proposed methodology is capable of identifying cost effective modifications to substantially improve the fault tolerance of critical infrastructures.

I. General Methodology

The world is increasingly dependent on critical infrastructures such as the electric power grid, water, gas and oil transport systems. At the same time, these infrastructures are increasingly susceptible to faults in the form of natural disasters and intentional disruption. Due to both increasing demand, which is outpacing infrastructure expansion, and the increasingly interconnected nature of infrastructures, many critical infrastructures are becoming vulnerable to cascading failures, where a fault caused by an external force may induce a domino-effect of component failures.

These trends combined raise the specter of a well targeted attack bringing down an entire system of interconnected infrastructures, resulting in a devastating economic blow and potentially a significant loss of life. An important implication is that traditional infrastructure risk analysis methods, often relying on Monte Carlo sampling of disaster scenarios, are no longer sufficient. Instead, systematic analysis based on worst-case attacks by intelligent adversaries is essential.

The addition of control devices, such as adding a pump in a water system, coupled with intelligent control algorithms allow the effective use of spare system capacity. When working together, multiple control devices have the potential to protect components by better use of underutilized areas. The ability to balance system use in such a manner presents a means to stop cascading failures as well as better utilize system resources during normal operating conditions.

The problems of finding optimally balanced hardenings and worst-case fault scenarios are interdependent and their solution spaces share the characteristics of combinatorial complexity, making exhaustive search infeasible, and nonlinear dependencies between solution components, resulting in many local optima which defeat most traditional search algorithms. This paper provides a methodology, based on competitive coevolutionary algorithms inspired by Neo-Darwinian arms races, to simultaneously evolve near optimal hardenings and low-cost, high-impact faults. The effectiveness of our methodology is demonstrated on a power system case study.

Competitive Coevolution

Evolutionary Algorithms (EAs) are a class of stochastic population based optimization techniques inspired by principles of Neo-Darwinian evolution in which a population of solutions is evolved to maximize a fitness function. First, individual solutions are combined through recombination operators to create new offspring solutions which contain aspects of their parent solutions. Next, mutation operators are applied to the offspring solutions to provide a form of random search and a means of genetic variation. Finally,

31st Annual International Computer Software and Applications Conference(COMPSAC 2007)
0-7695-2870-8/07 $25.00

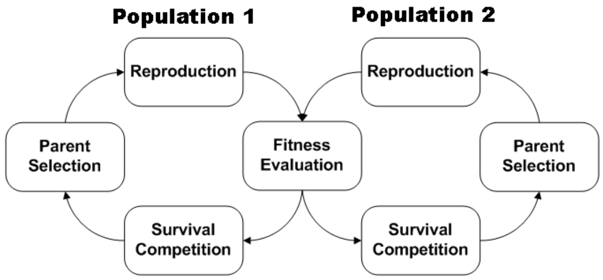

Figure 1. 2-population coevolution lifecycle

offspring are forced to compete for survival within the population. In this manner new areas of the fitness landscape are explored as new solutions are created from old solutions.

In Coevolutionary Algorithms (CoEAs) an individual's fitness is determined by interactions with other individuals rather than by a fixed metric as in typical EAs [4]. Fig. 1 shows the basic coevolutionary lifecycle for a two-population CoEA. Each population has its own lifecycle, which intersect during fitness evaluations.

In two-population competitive coevolution both populations are constantly evolved to exploit weaknesses in the opposing population. As a result both populations are constantly being exposed to more fit individuals from the opposing population. Ideally, this leads to incremental improvement, with each population continually evolving to meet the increasing pressure from the opposing population.

Proposed Methodology

A hardening of an arbitrary system is a set of modifications to that system that make it more resilient against faults. Examples of such modifications are adding transmission lines, generators, or power flow controllers to an electric power grid. While creating possible hardenings is easy, evaluating their effectiveness typically is not. Ideally, a hardening would effectively allow the system to withstand worst-case faults as well as minimize the damage caused by more probable faults. However, in many cases it may not be possible to find hardenings which, under all faults, can prevent a major system failure. Hardenings should then postpone the onset of a cascading failure allowing operators time to either repair the system or make intelligent decisions which will minimize further damage.

In practice those hardenings with the greatest cost-benefit ratio should be adopted. While determining this exactly in most cases is impractical, a heuristic can be used to approximate it. Estimates of the cost-benefit ratio of a given hardening and effectiveness of a fault are shown in Eq. 1 and Eq. 2 respectively, where H is the hardening, $C(H)$ the overhead cost associated with H, Ω the set of all faults, E_σ the expected number of times fault σ will occur throughout the lifetime of the system, and $L(H,\sigma)$ the expected monetary loss caused to the system by σ with hardening H in place.

$$F(H) = C(H) + \sum_{\sigma \in \Omega} E_\sigma \cdot L(H,\sigma) \quad (1)$$

$$G(\sigma) = \sum_{H \in \Sigma} [E_\sigma \cdot L(H,\sigma) + C(H)] \quad (2)$$

The minimal points of $F(H)$ represents the hardenings which perform best against all faults; likewise the maximal points of $G(\sigma)$ represent those faults which, on average, cause the most damage to any hardening.

Since it is not generally feasible to evaluate a hardening against all faults, nor a fault against all hardenings, a test set of probable damaging faults can be selected and hardenings found which are capable of lessening the damage of those faults. This, however, runs the risk of developing hardenings which protect well against only the test case faults. Competitive coevolution allows hardenings to be ever challenged by increasingly damaging faults, resulting in, ideally, hardenings which withstand a wide range of faults and cover general areas of the system in need of strengthening.

II. Power Grid Case Study

The electrical power grid is one of the largest interconnected critical infrastructures ever built and its correct functioning is a prerequisite for the correct functioning of most other critical infrastructures. One of the most promising power system control devices is the family of power electronics-based controllers known as Flexible AC Transmission System (FACTS) devices. While many different types of FACTS devices exist, one of the most powerful is the Unified Power Flow Controller (UPFC). Working in concert, UPFCs have been shown capable of preventing cascading failures [6]. This paper presents empirical evidence of the effectiveness of the proposed infrastructure hardening methodology by simultaneously optimizing the placement of UPFC devices and identifying worst-case power grid faults.

Related Work

Previous work has gone into optimizing UPFC placement and control [1, 2, 3]. In [2] a UPFC control algorithm based on the maxflow algorithm is presented in which UPFC set points are chosen such that the resulting system flow, under a fault, is close to the steady state flow under no faults. However, as the number of faults increases this flow model becomes increasingly inaccurate.

31st Annual International Computer Software and Applications Conference(COMPSAC 2007)
0-7695-2870-8/07 $25.00 © 2007 IEEE

In [7], a Sequential Quadratic Programming control algorithm is used to minimize a performance index:

$$PI = \sum_{i \in lines} \left(\frac{S_i}{S_i^{max}} \right)^2 \qquad (3)$$

where S_i is the current power flow through line i and S_i^{max} is the maximum rated capacity of line i.

In [5] an EA is used to evolve UPFC placements to minimize the number of line overloads over all single line faults, but does not include secondary, induced failures.

Power System Model & UPFC Control

A steady state model of the IEEE 118 bus system[1] is used as the testbed for all experiments. A Newton-Raphson solution method is used to solve the polar form of the steady state loadflow equations. The method in [8] is used to restore solvability in unsolvable situations such as when demand exceeds system capacity.

To simulate cascading failures several iterations of loadflow are performed to account for additional contingencies induced by the initial fault. A model of an overcurrent relay is employed to determine the times at which individual lines fail in response to an original fault. The work here simulates load served over the course of a day.

Long term UPFC set points are determined by minimizing the performance index in Eq. 3, as done in [7]. Since lines fail at a rate proportional to their degree of overload, choosing set points in such a manner causes the next line failure to be delayed

Experimental Setup

Cost, $L(H, \sigma)$, is proportional to the demand which cannot be delivered, and the expected number of times a given fault will occur, E_σ, is assumed to be $\frac{1}{1+|\sigma|}$. To evaluate the fitness of the individuals in the UPFC placement population and the fault population, subsets of both populations are selected to be used as test cases. Eq. 1 and Eq. 2 are simplified by assuming the initial cost of all hardenings/placements to be zero and are adapted to be fitness functions for the UPFC placement and fault populations by changing the sums over all possible placements and faults to sums over the members of the test set. The fittest individual from the previous generation is included in the test set along with three randomly selected members in the current population.

Individuals from both populations are represented as bit strings. Semi-generational survival selection is used in which only the three best individuals in the current population are allowed to survive into the next. Binary tournament selection is used for parent selection. Offspring are created by randomly selecting from a set of possible reproduction operators and performing the operation on the selected parents.

Recombination for both populations is accomplished by one-point crossover, with 40% probability. Mutation is accomplished for both populations by both a bit flip operator, with 20% probability, and by either moving the UPFC or the outaged line to a neighboring line, with 40% probability. Individuals are initialized such that, on average, two lines are outaged or two UPFCs are placed randomly in the system.

Due to the high computational cost of fitness calculation, small populations consisting of 10 individuals are used. Coevolution is terminated after 50 generations and five runs are performed.

Results

The evolved UPFC placements are successful in delaying the evolved faults' cascades from occurring, but in the cases seen are unable to entirely prevent them. Fig. 2 shows a graph of the percentage of load served at each time step for evolved UPFC placements from generations 0, 10, and 20, and that of the system when UPFCs are not used under the best evolved fault. The best evolved UPFC placement significantly delayed the cascade by 30 minutes.

In general it is desirable for current individuals to be able to defeat opponents from earlier generations and to perform better against current opponents then their ancestors. As the generations increased, seen in Fig. 2, the best placements from those generations were better able to withstand the best evolved fault, until about generation 20 when a placement was evolved which performed the best against the fault.

The final evolved faults, on average, forced the system to only serve about 30% of the demand. This may indicate that finding good faults is easier than good UPFC placements in the test configuration.

The evolved faults ranged from removing only two lines in the system to removing as many as eight. The fact that faults of as little as two lines can cause serious damage may indicate that the spare capacity available in the test system was insufficient.

III. Discussion & Impact

One problem inherent with coevolution is measuring evolutionary progress. Fitness values of individuals in different generations can no longer be directly compared as they were judged against different sets of individuals. In

[1] http://www.ee.washington.edu/research/pstca/pf118/pg_tca118bus.htm

31st Annual International Computer Software and Applications Conference(COMPSAC 2007)
0-7695-2870-8/07 $25.00 © 2007 IEEE

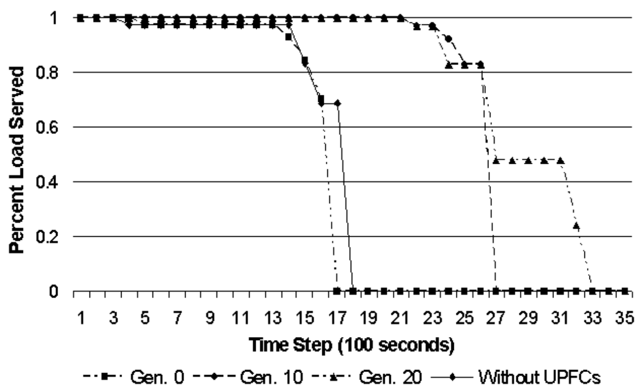

Figure 2. Comparison of placements from multiple generations against the same fault

typical EAs a plot of the population's maximum and average fitness values over time will indicate if the EA's solution quality is improving; however, in CoEAs this is not the case. Since CoEAs have no explicit objective function being optimized there may not be a natural metric by which to measure progress. One method to determine coevolutionary progress, used in the above case study, is to compare the performance of the current individuals against past and present opponents to the performance of their ancestors against the same opponents.

In the power system case study, faults are able to be evolved to exploit specific weaknesses in the hardenings, specifically a weakness in the UPFC control. It may be necessary to provide a means to either speed the evolution of the hardening population or slow the evolution of the fault population in order to allow enough time for the hardening population to evolve effective countermeasures to the faults or vice versa.

Conclusions and Future Work

The proposed coevolutionary critical infrastructure hardening methodology is successfully demonstrated by applying it to a power system case study. Placements of UPFC devices are evolved which demonstrate the ability to slow the progress of a cascading failure, allowing operators time to intercede and restore the system before a blackout occurs. Our proposed methodology is shown to effectively evolve UPFC placements which outperform the initial randomly selected placements.

To increase real-world applicability of this methodology to power systems, more accurate models and control algorithms will be used in future studies. By combining multiple forms of system hardenings, such as control devices and topological changes, more comprehensive hardenings of critical infrastructures can be identified. The versatility of our methodology needs to be tested on additional critical infrastructures.

References

[1] H. Abdelsalam, G. Aly, M. Abdelkrim, and K. Shebl. Optimal location of the unified power flow controller in electrical power systems. In *Power Systems Conference and Exposition*, volume 3, pages 1391 – 1396. IEEE, 2004.

[2] A. Armbruster, M. Gosnell, B. McMillin, and M. Crow. Power transmission control using distributed max-flow. In *29th Annual International Computer Software and Applications Conference*, volume 1, pages 256 – 263. COMPSAC, 2005.

[3] L. Cai, I. Erlich, and G. Stamtsis. Optimal choice and allocation of FACTS devices in deregulated electricity market using genetic algorithms. In *Power Systems Conference and Exposition*, volume 1, pages 201 – 207. IEEE, 2004.

[4] R. C.D. *Coevolutionary Search Among Adversaries*. PhD thesis, University of California, San Diego, 1997.

[5] J. Chaloupek, D. R. Tauritz, B. McMillin, and M. Crow. Evolutionary Optimization of Flexible AC Transmission System Device Placement for Increasing Power Grid Reliability. In *Proceedings of FEA 2005, the 6th International Workshop on Frontiers in Evolutionary Algorithms*, pages 516–519, July 2005.

[6] A. Lininger, B. McMillin, M. Crow, and B. Chowdhury. Analysis of Max-Flow values for setting FACTS devices. Technical report, FACTS Interaction Laboratory, 2007. `http://filpower.umr.edu`.

[7] W. Siever, D. Tauritz, and A. Miller. Improving grid fault tolerance by optimal control of FACTS devices. In *Proceedings of First International ICSC Symposium on Artificial Intelligence in Energy Systems and Power - AIESP 2006*, 2006.

[8] O. T.J. Computation of a practical method to restore power flow solvability. *IEEE Transactions on Power Systems*, 10(1):280 – 287, 1995.

0-7695-2870-8/07 $25.00 © 2007 IEEE