

## Network Disruption White Paper

*Most real-world optimization problems have several, often conflicting objectives. Network Disruption is no different. Therefore, the result is not a single solution, but a set of solutions that trade off between objectives.*

### Networks Defined

Networks are essential infrastructures in society. Modern communication networks allow the transfer of communication at extremely high speeds and users to gather information in large volumes nearly instantaneously. Business and industry use networks to carry out critical functions, like the transfer of money, or to communicate sensitive business information or update databases.<sup>1</sup> Other examples of networks include computer, financial, transportation, power distribution, and social networks.

A network is comprised of links and nodes arranged so that messages or goods may be passed from one part of the network to another over multiple links and through various nodes. Interrupting or isolating the nodes to degrade network function is referred to as network disruption.

### Network Disruption

Network disruption attacks nodes and degrades networks. This may have application in various arenas, but the most common use is in the military to isolate enemy command centers from one another.

The problem lies within finding good network disruption attacks or target sets, i.e., a subset of vertices or edges that, once removed, impede communication or connections between particular nodes. Since multiple costs are associated when disrupting networks, success in bringing down networks is traded off against the costs of the attack plans.

Most real-world optimization problems have several, often conflicting objectives. Network Disruption is no different. Therefore, the result is not a single solution, but a set of solutions that trade-off between objectives. The Italian economist Vilfredo Pareto first formulated this concept in 1896, and it bears his name today. A solution is Pareto optimal if (for a maximization problem) no increase in any criterion can be made without

a simultaneous decrease in any other criterion.ii The line along the outer boundary of this Pareto optimal region is known as the *efficient frontier*. Solutions along this line represent options for which there is the lowest risk for a given level of return.

For various network disruption problems, the efficient frontier has value in bringing out a small fraction of worthwhile attack plans from among many possible plans, without asking the user for any difficult a priori judgments. Pareto genetic algorithms can discover good attacks rapidly among many possibilities. In later generations, further Pareto optimal plans are found at a decreasing rate. Diagrams make the results easy to survey and understand, and an aggregate view or center of gravity makes clear which of the targets are the most important for mission success.

### **Case Study**

Aetion applied a Pareto (dominance-based) genetic algorithm to generate network attack plans, and used cross-linked diagrams combined with an 'aggregate view' of the set of attack plans to study the results.


#### **A. Problem description**

The desired outcome of this experiment is to minimize communications among certain important nodes, maximize the number of pieces into which the network is broken, and minimize the various measures of cost. Aetion used the domain of communication networks as a concrete instantiation of a more abstract problem. The problem could be instantiated not only as a physically-realized communication network, but also, for instance, as a network of contacts between people, or as a transportation or distribution network. The non-dominated set of attacks that is found should be optimized in the sense the set should be as close to possible to the Pareto-optimal front, and the set of solutions should be as diverse as possible. This study concentrated on the first condition, since a goal is to defer value judgments of the tradeoffs on the optimal front, letting the user explore those tradeoffs.

#### **B. Scenario**

The scenario used involved randomly-generated connected graphs where the vertices represent communication nodes and the edges between them represent communication links. An attack plan is a set of nodes and/or links that are to be eliminated. Evaluation of

plans was performed by analyzing the post-attack network. Three of the nodes are distinguished: they represent command centers, and the mission is to completely isolate them from one another. The graphic below is the visualization of the links, nodes, and command centers.



C. The experiment introduces the problem domain, evaluation criteria, efficient frontier, and visualization. A network of 9 nodes and 18 links were considered, with command centers X, Y and Z. Links were annotated with a bandwidth, and Aetion only considered attacks that targeted links. The traffic routing protocol was assumed to be able to find a single highest-bandwidth path between nodes. With such a small network, each of the 31,180 attack plans that target six links or fewer could be rapidly considered.

In this experiment, each node and edge was randomly assigned a value for friendly-fire casualty probability and expected number of civilian casualties. The values of these two criteria for a given attack are, respectively, the combined probability and the sum of the values for the nodes and links that compose the attack. The random assignment of casualty values was done in such a way as to allow the demonstration of the capability of the Aetion solution and the optimization methods to find subtle yet salient tradeoffs.

The aim is to minimize the number of targets to be destroyed, the bandwidth remaining after the attack, and the projected casualties. A goal is to maximize the number of fragments into which the network is broken by the attack. Given these criteria, from the 31,180 attack plans, only 46 plans lie on the efficient frontier.

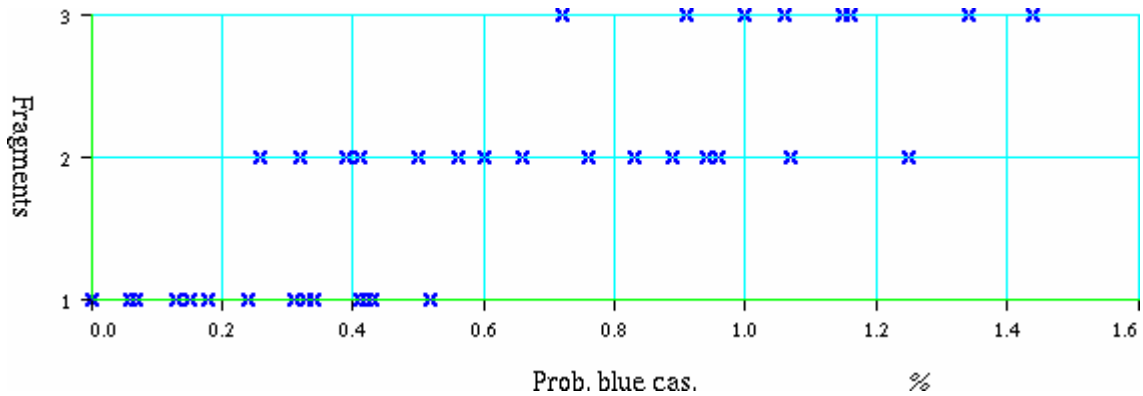



Figure 2. The 46 plans on the efficient frontier.

Color	# targets	Y-Z b/w	Fragments	Prob. blue cas. (%)	# exp. white cas.	Targets
Blue	4	12800	2	0.89	272	0, 8, 9, 16
Blue	5	12800	2	0.76	222	1, 7, 9, 14, 16
Blue	5	4800	2	0.96	552	0, 8, 9, 16, 17
Blue	5	3200	2	1.07	317	0, 2, 8, 9, 16
Blue	6	4800	2	0.83	502	1, 7, 9, 14, 16, 17
Blue	6	3200	2	0.94	267	1, 2, 7, 9, 14, 16
Blue	6	3200	2	0.66	739	2, 6, 9, 14, 16, 17
Red	6	0	3	1.15	425	0, 2, 3, 8, 9, 16

Figure 2. The plans that completely isolate X from Y and Z.

Visualization of these 46 plans reveals some reasonable results. For instance, figure 2 shows how the probability of friendly-forces casualties is correlated with the number of completely isolated fragments into which each plan (plotted with 'x's on the diagram) splits the network. Figure 3 shows the plans that completely isolate X from Y and Z. The one on the bottom line, marked in black, shows Y-Z having 0 bandwidth: it isolates X, Y, and Z from each other, but at the cost of the highest risk of friendly-force casualties.

This illustrates how the efficient frontier brings out the tradeoff decisions. Of course, to completely isolate the command centers from one another, the network must be split into at least three fragments. Figure 4 explores how many must be destroyed in order to achieve this. To interpret the figure, it is important to note that the area of each block is directly proportional to the number of attack plans that it represents. The plans that isolate X from Y and Z, those from figure 3, are those that are shown in black. So, just to achieve that isolation, destroying four links suffices. Figure 4 indicates redundancy in the network: it remains connected if only one link is removed, and it sometimes remains connected if five are removed.




**Figure 4. Shows number of targeted links.**

### Action Solution

Action is the leader in providing patented technology that integrates multi-objective optimization and interactive Pareto frontier decision making. The solution, Seeker-Filter-Viewer (SFV), offers a novel multi-objective portfolio optimization algorithm that uses

evolutionary computation for problems with multiple measures of risk and return, where the measures may be nonlinear and nonconvex.<sup>iii</sup>

Aetion's offering is broken down into three functional components, the Seeker, the Filter, and the Viewer, as shown in the following figure. The "Seeker" generates alternatives by rule-governed composition of components and evaluates them according to multiple criteria. Then, the "Filter" removes the dominated alternatives, and the efficient frontier is viewed interactively in the "Viewer" by means of cross-linked diagrams wherein the same alternatives are identically colored in each diagram, enabling the comparison of alternatives from multiple perspectives.<sup>iv</sup>



**The Seeker** – The Seeker generates alternatives and evaluates them according to multiple criteria using simulations and/or other computational methods. The Seeker can use any number of networked computers to evaluate alternatives in parallel. The number of alternatives evaluated is limited mainly by the number of computers available, the time available, and the time it takes to evaluate. If needed, the Seeker can use simulated evolution to generate alternatives by mutation and breeding. By considering millions of alternatives through this simulated evolution, users can effectively search among *billions of alternatives*.

**Filter** - The total number of alternatives might be quite large, so it is desirable to use the computer to select a relatively smaller number of alternatives that are worth further examination. The Filter can employ a method based on a dominance rule so that filtering is lossless, i.e., so that there are guarantees that no good alternatives are excluded. In other words, an alternative A *dominates* another alternative B if A is superior or equal to B with respect to every criterion of evaluation and distinctly superior with respect to at

least one criterion. The alternatives generated by the seeker are systematically compared, and any candidate that is dominated by any other candidate is removed. The survivors comprise the *Pareto subset*. No member of the Pareto subset dominates any other element in the subset. Thus, deciding among alternatives in the Pareto subset is always a matter of tradeoffs.

After obtaining a Pareto subset, the filter can augment this space with additional metrics to reflect additional constraints for use in the tradeoff process. A graphical tool, called the Viewer, is used to visualize 2-D projections of the Pareto front.

For the experiment discussed within this paper, nine criteria were used by the dominance filter. Six criteria were minimizing the maximum flows after the attack, measured in both directions between the three command centers - named X, Y and Z. Two other criteria were estimated civilian losses and estimated casualties, which were to be minimized. A final criterion was the number of targets, which was to be minimized.

Dominated plans were removed; undominated plans survived from generation to generation as long as they remain undominated. After dominance filtering, a mutated copy of the undominated set was added to the population, along with an equal-sized set of offspring bred from the undominated set. This type of filtering does not require weighting or any other form of scalarization. Thus, the population size is not fixed and is free to increase or decrease from generation to generation. After dominance filtering, a mutated copy of the undominated set is added to the population, along with an equal-sized set of offspring bred from the undominated set.

**Viewer** – The Viewer enables decision-makers to perform interactive visual trade-off analyses of the elements of the Pareto set and to narrow to a subset for further exploration. The Viewer is used to explore the surviving design alternatives, and to visualize the trade-offs between pairs of evaluation criteria. The Viewer enables the planner to visualize performance characteristics and other attributes of sets of alternatives using various types of interlinked displays, including *trade-off diagrams*, which are two-dimensional scatter plots with points that represent alternatives and axes that represent criteria. The user can graphically select a subset of the alternatives and explore their ranges of values in multiple dimensions. In this way, the users gain an

understanding of the decision space and can act to successively narrow the focus to reduce the number of alternatives under consideration. The user is able to zoom in on subsets with desirable tradeoff characteristics, and reduce the number of alternatives toward making a final selection. The Viewer is graphical, mouse driven, and browser-embeddable so it can be easily deployed over the Internet or over an intranet. An example of a trade-off diagram appears below.

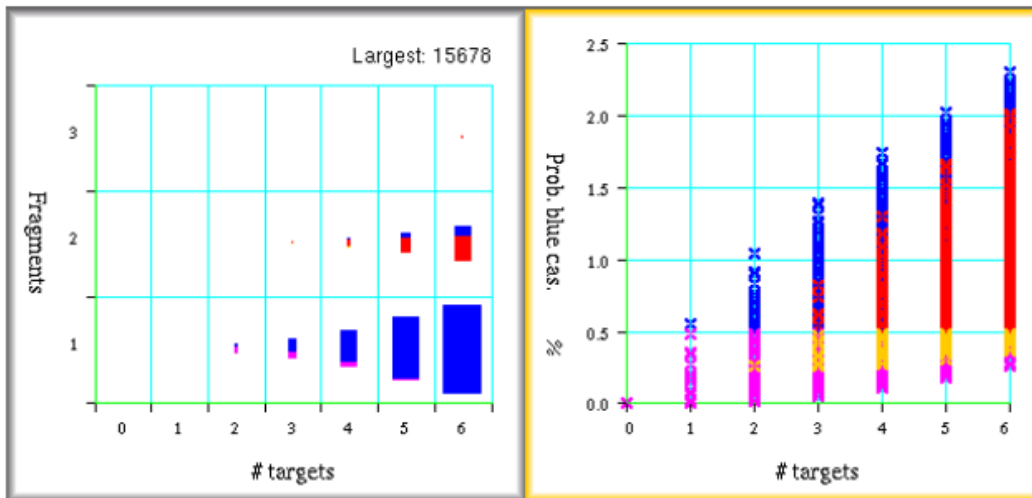


Figure 1: Cross-linked displays in the viewer

### Key Benefits

Evidence suggests that people are poor at multiple-criteria decision making, even when there are only a few alternatives, often picking alternatives that are clearly sub-optimal. The SFV solution is effective even when the number of alternatives is small. However, the technology works particularly well when the number of alternatives is very large, say, in the millions. In such situations, the technology empowers decision makers to explore many more possibilities - orders of magnitude more possibilities - than would otherwise be possible. Some additional advantages in using SFV technology includes its ability to:

- *Extend human capabilities* -- SFV technology takes advantage of human cognitive strengths, and provides computer assistance to mitigate human cognitive weaknesses. The objective is to improve human decision-making, not supplant it. This is *cognitive augmentation*, or *intelligence amplification*. The computer and the human do what each is best at: the human makes the ultimate

value judgments, while the computer provides for acquiring, scoring, keeping track, and routine comparison of all the alternatives.

- *Provide insight into the decision space* - A major benefit of the approach is the *understanding* of the decision space that decision makers get as a result of exploring a vastly extended range of possibilities. Then, the SFV technology allows users to explore tradeoffs and dependencies interactively and successively narrow the choice set.
- *Manage uncertainty* -- SFV technology provides many ways in which decision makers can manage or eliminate uncertainty in achieving their goals. Very large numbers of alternatives can be considered, thus reducing the chance that surprises are lurking undiscovered. The proverbial needle in the haystack can be brought to light. Parallel processing supports this large-scale exploration, so that it may be accomplished quickly if large numbers of processors can be employed. The study of many evaluation results in the Viewer will highlight surprising ones that are worth further investigation.
- *Enable out of the box thinking*-- Decision makers are encouraged to consider many more possibilities than they would otherwise be able to handle. This makes possible discoveries outside the range of usual thinking.

### **Take Aways**

- The case study shows how network disruption can isolate command centers from each other while minimizing civilian and friendly-force casualties and the number of targets attacked, and maximizing the number of fragments into which the communication network is broken. The problem lies within finding good network disruption attacks or target sets, i.e., a subset of vertices or edges that, once removed, impede communication between particular nodes. Since multiple costs are associated when disrupting networks, success in bringing down communications is traded off against the costs of the attack plans.
- SFV technology is useful for many multi-criterial decision making tasks, not simply network disruption planning. Evidence suggests that people are poor at

multiple criteria decision making, even when there are only a few alternatives, often picking alternatives that are clearly sub-optimal. The SFV solution is effective even when the number of alternatives is small. However, the technology works particularly well when the number of alternatives is very large, say, in the millions. In such situations, the technology empowers decision makers to explore many more possibilities - orders of magnitude more possibilities - than would otherwise be possible.

- Most real-world optimization problems have several, often conflicting objectives. Network Disruption is no different. Therefore, the result is not a single solution, but a set of solutions that trade-off between objectives.

### **For More Information**

Aetion Technologies

[Info@Aetion.com](mailto:Info@Aetion.com)

(614) 340-1835

(614) 487-3704 (fax)

Aetion Technologies LLC

1275 Kinnear Road

Columbus, OH 43212-1155

---

<sup>i</sup> Alberto Leon-Garcia, Indra Widjaja, *Communication Networks Fundamental Concepts and Key Architectures*, 2007.

<sup>ii</sup> R. Subbu, P. Bonissone, S. Bollapragada, K. Chalermkraivuth, N. Eklund, N. Iyer, R. Shah, F. Xue, and W. Yan, "A Review of Two Industrial Deployments of Multi-criteria Decision-making Systems at General Electric," *General Electric Global Research*, One Research Circle, Niskayuna, NY 12309, USA.

<sup>iii</sup> Ibid.

<sup>iv</sup> M. Carroll, J. Josephson, J. Russell, Tradeoffs on the Efficient Frontier of Network Disruption Attacks, This research was prepared under sponsorship from the U.S. Naval Surface Warfare Center under contract N00178-02-C-3063, through participation in the Advanced Decision Architectures Collaborative Technology Alliance sponsored by the U.S. Army Research Laboratory under Cooperative Agreement DAAD19-01-2-0009, and by federal flow-thru by the Department of Defense under contract FA8652-03-3-0005 (as a subcontract from Wright State University and Wright Brothers Institute).