Proceedings of the 2007 IEEE Symposium on Computational Intelligence in Multicriteria Decision Making (MCDM 2007)

Tradeoffs on the Efficient Frontier of Network Disruption Attacks

Mark T. B. Carroll
Aetion Technologies LLC
1275 Kinnear Road
Columbus, OH 43212, USA
Telephone: +1 614 340 1835
Fax: +1 614 487 3704
Email: Mark.Carroll@Aetion.com
Email: Carroll@CSE.Ohio-State.edu

John R. Josephson
Dept. of Computer Science and Engineering
The Ohio State University
2015 Neil Avenue
Columbus, OH 43210, USA
Telephone: +1 614 292 0208
Fax: +1 614 292 2911
Email: JJ@CSE.Ohio-State.edu
Email: John.Josephson@Aetion.com

James L. Russell
Aetion Technologies LLC
1275 Kinnear Road
Columbus, OH 43212, USA
Telephone: +1 614 340 1835
Fax: +1 614 487 3704
Email: James.Russell@Aetion.com

***Abstract*— A communications network is represented as a graph of flow capacities. We study the problem of finding good network disruption attacks or target sets, i.e., a subset of vertices or edges that, once removed, impede communication between particular nodes. Multiple costs are associated with removing vertices or edges. Success in disrupting communications is traded off against the costs of the attack plans: the efficient frontier of attacks is estimated, and the results are studied in cross-linked diagrams. A multicriterial genetic algorithm is used to discover good plans for disrupting the communications network, where the genes correspond to nodes or links to be attacked. The genetic algorithm is seeded with an initial population of single-target genomes, one for each potential target. Multi-target attacks may be generated by breeding. Being on the efficient frontier guarantees a genome's survival to the next generation, so the population size is allowed to vary. The results are studied in interactive diagrams and in an "aggregate view" of the resulting population. Good attacks were found relatively rapidly, and the aggregate view revealed significant targets.**

I. Background and Aims

A decision alternative A **dominates** B if, with respect to the criteria of merit, A surpasses B on at least one criterion and B surpasses A on no criterion. Among a set of alternatives, the efficient frontier is the subset of alternatives that are not dominated by any other. Markowitz described in [1] the efficient frontier, or Pareto-optimal set, in the context of investment portfolios.

The "Seeker-Filter-Viewer" (SFV) architecture is designed to support multicriterial decision making in large decision spaces [2], [3], [4]. 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.

For generating the set of decision alternatives, an alternative approach to the original version of the seeker is to use genetic algorithms to sample a large search space. The efficient frontier of generated alternatives can be used to decide which alternatives survive into the next generation. This technique has been applied in the context of investment portfolios as described in [5].

The SFV architecture has interesting features. The seeker can use distributed computing techniques to evaluate very many decision alternatives, millions or more, and use genetic algorithms to sample search spaces of billions. The filter is very effective, yet it is lossless in that it removes only dominated alternatives. As described in [6], as the number of alternatives increases, the fraction retained by the filter tends to decrease. The diagrams in the viewer are used to assist the user in making value judgments: plots of the efficient frontier show the tradeoffs among the best alternatives.

SFV exploits two synergies. First, that a large population of alternatives is generated and evaluated by the seeker is tolerable because the population's size is much reduced in number by the filter. Second, the survivors of the filter are exactly the alternatives between which tradeoff judgments must be made, and the viewer is an excellent interface for exploring tradeoffs.

The use of the filter in the genetic algorithm, and then use of the viewer to investigate the results, offers an important advantage: the user is not required to specify ranks, weights, or cut-offs for the criteria of merit in order to express their preferences; none of these would often capture the user's values exactly. Instead, ours is a compensatory approach wherein the value judgments are all postponed until the user

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). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Naval Surface Warfare Center, Army Research Laboratory, Defense Department, or the U. S. Government.

is faced with the concrete tradeoffs, at which point decision-making proceeds on the basis of comparing actual decision alternatives.

As well as SFV being useful for multicriterial decision problems, it assists model validation. In evaluating a wide range of decision alternatives, then allowing the user to graphically mine the results, it allows the user to discover anomalous patterns in the results that may indicate modeling errors. Thus, one of the things that we look for in the results of the experiments described below is that they are not inexplicably surprising.

Pinkstaff describes in [7] the multicriterial optimization problem of deciding how best to attack a communications network. In the work we describe below, inspired by the problem domain described by Pinkstaff, we applied a Pareto (dominance-based) genetic algorithm to the generation of network attack plans, and used the cross-linked diagrams from the SFV architecture, combined with an 'aggregate view' of the set of attack plans, to study the results.

II. Experiments

A. Problem description

In this study we use 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. A description of the abstract problem follows.

Given a graph $\mathcal{G} = (\mathcal{V}_\mathcal{G}, \mathcal{E}_\mathcal{G})$ with an indexed set of distinguished vertices $\mathcal{X}_{\mathcal{G}_i}$, find a non-dominated set of attacks, where each attack $\mathcal{A} = (\mathcal{V}_\mathcal{A}, \mathcal{E}_\mathcal{A})$, creates an induced subgraph $\mathcal{G}' = (\mathcal{V}_\mathcal{G} - \mathcal{V}_\mathcal{A}, \mathcal{E}_\mathcal{G} - \mathcal{E}_\mathcal{A})$. The set of criteria used for dominance filtering the set of attacks includes:

- given a bandwidth function $bw(\mathcal{V}_i, \mathcal{V}_j)$ for some model of network flow, independently minimize the objectives $bw(\mathcal{X}_{\mathcal{G}_i}, \mathcal{X}_{\mathcal{G}_j})$, for $i > j$ (for undirected graphs) or $i \neq j$ (for directed graphs), where $bw(\mathcal{V}_i, \mathcal{V}_j) = 0$ if $\mathcal{V}_i$ or $\mathcal{V}_j$ are not vertices in $\mathcal{G}'$.
- maximize the number of (strongly-)connected-components (fragments) of $\mathcal{G}'$.
- independently minimize/maximize some additional set of (real-valued) objectives $f_i(\mathcal{V}_\mathcal{A}, \mathcal{E}_\mathcal{A})$.

Intuitively, we desire 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.

The non-dominated set of attacks that is found should be optimized in the sense of Deb [8], viz. the set should be as close to possible to the Pareto-optimal front, and the set of solutions should be as diverse as possible. We will concentrate here on the first condition, since our stated goal is to defer value judgments of the tradeoffs on the optimal front, letting the user explore those tradeoffs using the viewer component of the SFV architecture. Furthermore, since this is an initial exploration of this problem domain, it is not necessarily clear that there is an *a priori* best method for maintaining diversity in the solution set. However, we will argue that diversity preservation in our methods is sufficient for the problem.

B. Scenario

The scenario that we used for experiments 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.

C. Experiment #1 – introduces the problem domain, evaluation criteria, efficient frontier, and visualization

We first considered a network of 9 nodes and 18 links, with command centers X, Y, and Z. Links were annotated with a bandwidth and we only considered attacks that target links. The traffic routing protocol was assumed to be able to find a single highest-bandwidth path between nodes. With such a small network, we can very rapidly consider each of the 31,180 attack plans that target 6 links or fewer. The criteria of evaluation for attack plans are described in Table I and the result of plotting each plan against each criterion is shown in figure 1. Additionally, each plan has an attribute named "Targets" that lists the unique identifier of each link to be attacked in that plan.

The last two criteria correspond to the set of additional (unspecified) criteria discussed in the abstract problem description. In this and the following experiments, each node and edge of the graph was randomly assigned a value for blue 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 SFV architecture and the optimization methods to find subtle yet salient tradeoffs.

We aim to minimize the number of targets to be destroyed, the bandwidth remaining after the attack, and the projected casualties; we aim 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 lie on the efficient frontier. Visualization of these 46 plans reveals some

TABLE I
CRITERIA OF EVALUATION FOR EXPERIMENT #1 ATTACK PLANS

# targets:	The number of links that are destroyed in the plan
X-Y b/w:	Bandwidth between X and Y after the attack
X-Z b/w:	Bandwidth between X and Z after the attack
Y-Z b/w:	Bandwidth between Y and Z after the attack
Fragments:	Number of isolated fragments into which the attack splits the network
Prob. blue cas.:	The probability that blue forces will suffer any losses
# exp. white cas.:	The expected number of civilian casualties

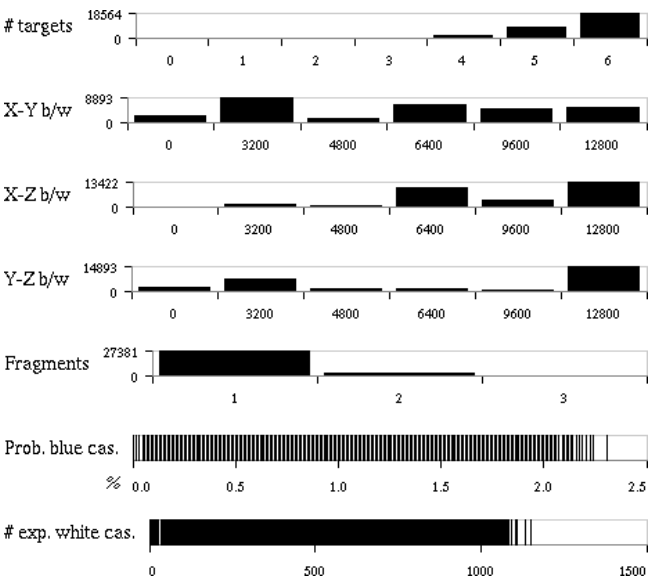

Fig. 1. All attack plans from experiment #1

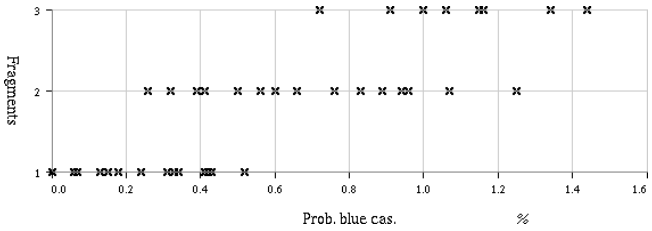

Fig. 2. The 46 plans on experiment #1's efficient frontier

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

Fig. 3. Experiment #1 plans that completely isolate X from Y and Z

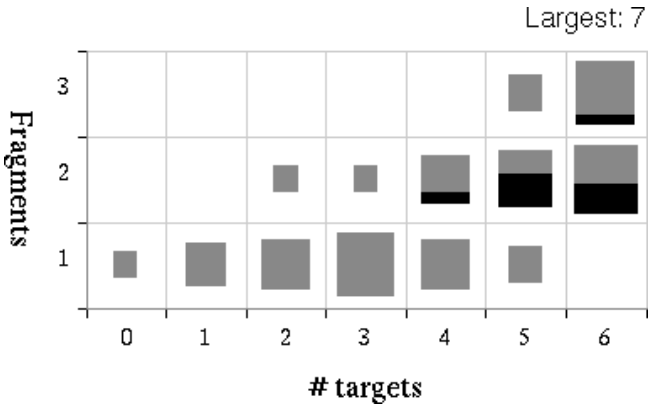

Fig. 4. Number of targeted links in experiment #1 plans

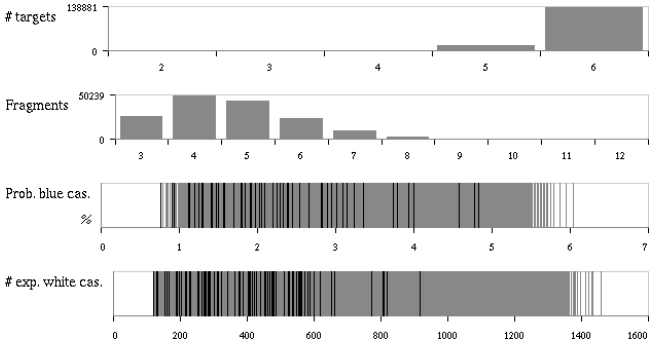

Fig. 5. Successful attack plans from experiment #2

reasonable results. For instance, figure 2 shows how the probability of blue (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 $Y \leftrightarrow Z$ having 0 bandwidth: it isolates X, Y, and Z from each other, but at the cost of the highest risk of blue 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 4 links suffices. Figure 4 indicates redundancy in the network: it remains connected if only 1 link is removed, and it sometimes remains connected if 5 are removed.

D. Experiment #2 – increasing realism

We now expand our network to include 12 nodes connected by 24 links. Attacks may target nodes and/or links. The command centers, now designated 1, 5, and 9, are well-defended (reflected in casualty estimates), making it harder to find a simple, safe plan to completely isolate the centers from each other. Again, we consider all plans that attack up to 6 targets. Of these 2,391,496 attack plans, 157,926 completely isolated the command centers from each other. Of those, 86 lie on the efficient frontier, and are highlighted in black in figure 5.

In our interactive visualization environment, we now mark the least risky plans on the efficient frontier in black, and the plans with the fewest targets in dark gray. These two sets are disjoint: the command centers being well-defended, we must either attack them directly at high risk, or attack more targets around them in order to completely isolate them. The efficient frontier is shown in figure 6 and the plans with least risk or the fewest targets are detailed in figure 7. Even among the dark gray high-cost plans, although they all target command centers

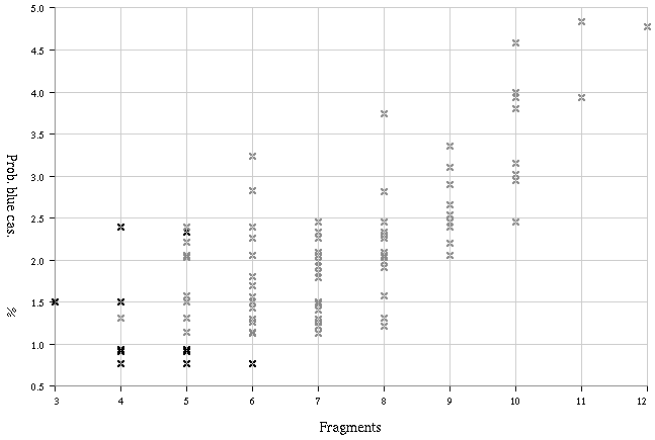

Fig. 6. Experiment #2 efficient frontier

Color	# targets	Fragments	Prob. blue cas. (%)	# exp. white cas.	Targets
	6	5	0.77	454	links 7, 22, 23, nodes 1, 2, 10
	5	4	0.77	442	links 7, 22, 23, nodes 1, 10
	6	5	0.91	453	links 8, 22, 23, nodes 1, 2, 10
	5	4	0.91	441	links 8, 22, 23, nodes 1, 10
	6	5	0.94	291	links 15, 22, 23, nodes 1, 2, 6
	5	4	0.94	279	links 15, 22, 23, nodes 1, 6
	6	6	0.77	549	links 22, 23, nodes 1, 2, 8, 10
	5	5	0.77	537	links 22, 23, nodes 1, 8, 10
	3	4	1.50	132	nodes 1, 2, 5
	3	5	2.34	230	nodes 1, 3, 5
	3	4	2.40	122	nodes 1, 4, 5
	2	3	1.50	120	nodes 1, 5
	3	5	2.06	272	nodes 1, 5, 7

Fig. 7. Interesting attack plans from experiment #2

1 and 5, none of them target all three command centers because destroying two of them successfully leaves each (remaining) one isolated from the others.

E. Experiment #3 – further increase in problem difficulty, introduction of genetic algorithms

We expanded the network to have 20 nodes and 50 links, precluding exhaustive search. Attacks may still target nodes and/or links. Links are now unidirectional, although a bidirectional link may be represented as a pair of unidirectional links. The asymmetry of unidirectional links increases the number of criteria as, to stop communication between P and Q, we wish to minimize both the bandwidth from P to Q and also from Q to P. The traffic routing protocol is now assumed to be sophisticated enough to use multiple paths at once for increased bandwidth, so we calculate bandwidth between nodes as maximum flows, using the Edmonds-Karp algorithm of [9], a Ford-Fulkerson method. In this network, the command centers are designated 1, 6, and 12.

Because the size of the network now makes exhaustive search infeasible, a genetic algorithm (GA) was introduced to guide the search. A relatively simple GA was used which has a combination of features that makes it distinct from approaches that are currently prominent ([8], [10] provide surveys).

Our method is an elitist GA which maintains a set of all non-dominated plans, where the population of each generation is mostly created anew and comprises all plans on the efficient

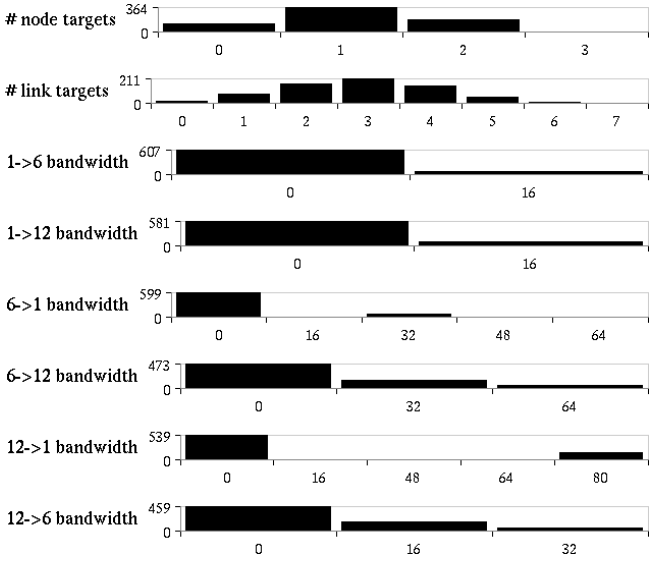

Fig. 8. All attack plans from experiment #3

frontier, as well as the results of random crossover and mutation of those plans. A minimum population size was also established as a diversity preserving measure, and dominated plans had a probability of surviving into the next generation that decreased as the population size increased, going to zero when the number of non-dominated plans exceeded the minimum population threshold.

Thus, our method can be viewed as having a variable and unconstrained population size. Growth of the population size remained manageable over a variety of parameter settings for the problem domain. Also, given that this study is, as previously mentioned, an initial exploration of the problem domain, we did not want to *ad hoc* artificially constrain the solution set before the user could explore it in the SFV viewer. Furthermore, retention of the entire non-dominated set of plans in conjunction with random population mutation lessened the need for any additional, explicit diversity-preservation mechanism, although it may still be worthwhile to include such a mechanism. It is possible that other GA methods could be substituted for our method, although elitism is a necessary characteristic (the benefits of which are discussed in [8] and [10]). We also feel that the strict elitism of keeping the entire non-dominated set, thus requiring a variable population size (at least so long as it is computationally feasible), is most consonant with the SFV architecture.

The diversity of the initial population of attack plans was ensured by seeding the GA with the set of all plans that have exactly one target, either a node or a link. In each generation, plans could breed or mutate. Breeding of plans was accomplished by taking a random subset of the union of the parents' individual node and link attacks. Mutating a plan consisted of taking a random subset of the plan's existing node or link attacks augmented with additional randomly-selected node or link attacks.

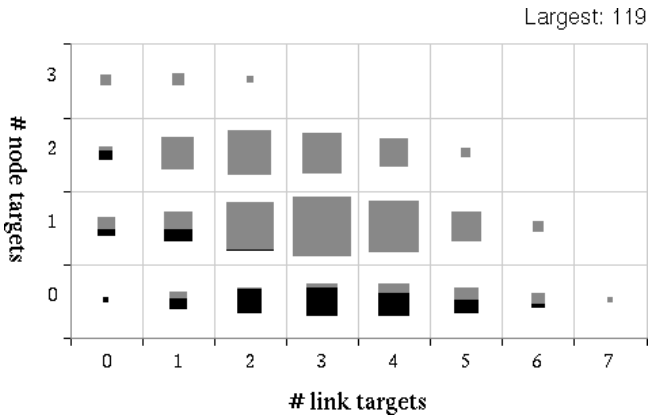

Fig. 9. Highlighting the experiment #3 efficient frontier

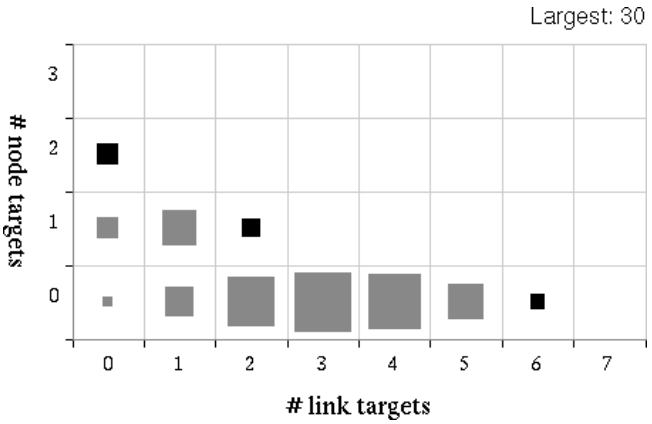

Fig. 12. Experiment #3 tradeoff between node attacks and link attacks

Node targets	Link targets	# node targets	# link targets
6,12	none	2	0
1,6	none	2	0
4,12	none	2	0
1,12	none	2	0
12	1->4,6->19	1	2
6	1->4,14->1	1	2
12	1->4,14->1	1	2
none	1->4,3->6,4->6,6->9,6->19,14->1	0	6
none	1->4,3->6,4->6,6->19,9->12,14->1	0	6

Fig. 10. Experiment #3 attacks that completely isolate command centers

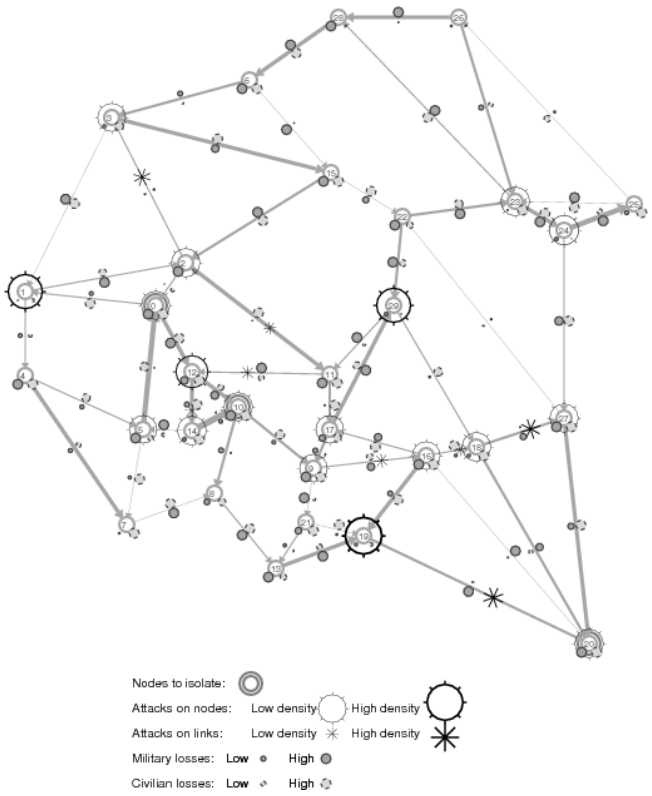

Fig. 11. Aggregate view of a population of attack plans

We ran the evolution for ten generations, resulting in a population of 702 plans. The criteria and plans are plotted in figure 8, by analogy with experiment #1's figure 1. The 118 plans on the efficient frontier are highlighted in black in figure 9, and clearly tend to be those that attack fewer targets.

The plans on the efficient frontier are particularly interesting. Figure 10 shows that, in only ten generations, the GA found each plan that completely isolates each of the three command center from the others by attacking two of them directly. These plans are shown in black in figure 12 which vividly illustrates the tradeoff between attacking nodes directly and attacking the links around them. It also shows that if we attack fewer targets than required by any of the figure 10 attacks, then at least one command center will remain in contact with another.

F. Later Experiments – larger networks, rate of progress of the Pareto GA, aggregate view

A larger connected network was randomly created, with 50 nodes and 100 directional links. The initial population again consisted of the possible one-target attacks. After 54 generations the population consisted of 128 attack plans. In instrumenting the GA we noted that the rate of progress slowed markedly: after the 20th generation, 87 of the 54th-generation plans had already been found.

Subsequently, an even larger example was constructed, with 100 nodes and 232 links. The GA was set to terminate when two consecutive generations find no new attacks that are on the efficient frontier. After 43 generations the GA terminated with 140 attack plans in the population.

For when the efficient frontier contains many attack plans, an 'aggregate view' was created so that the important targets can be identified. Figure 11 shows such a view (suffering somewhat from the reduction to grayscale), where the thicker links carry more bandwidth, and the thicker spiked circles around nodes and the thicker stars drawn across links illustrate that more of the attack plans target those nodes and links. These would be a fruitful starting point for searching for 'centers of gravity,' the critical sources of the enemy's strength,

as described in [11]. That more plans include these targets is circumstantial evidence that the targets are more significant determiners of mission success: in a less deterministic attack where more targets are eliminated, each with less reliability, this aggregate view may provide a strong hint that it is worth giving combat forces a relative focus upon eliminating those particular targets.

III. Conclusions

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

This effort was performed as a feasibility study where aspects changed to allow various situations to be considered. To better understand the phenomena, it would be worth varying single dimensions of the communications network, or of the plan generation, while holding others constant, so that an "apples to apples" comparison reveals the contribution of various aspects of the approach. Furthermore, although the relatively simple GA described above worked well for the problems we studied, it is possible that as the size of the network scales up, some of the decisions embedded in the design of the GA might need to be revisited, and techniques borrowed from other evolutionary algorithms that have been proposed.

That the above techniques worked well against various network-disruption problems provides encouragement for extending the approach to a wider set of combinatorially-complex military planning problems.

References

[1] H. M. Markowitz, "Portfolio selection," *Journal of Finance*, vol. 7, no. 1, pp. 77–91, March 1952.

[2] J. R. Josephson, B. Chandrasekaran, M. Carroll, N. Iyer, B. Wasacz, G. Rizzoni, Q. Li, and D. A. Erb, "An architecture for exploring large design spaces," in *AAAI/IAAI*, 1998, pp. 143–150.

[3] J. R. Josephson, B. Chandrasekaran, and M. Carroll, "System for multi-criterial decision making," U.S. Patent 6 771 293, August 3, 2004.

[4] J. R. Josephson, B. Chandrasekaran, M. Carroll, and N. S. Iyer, "Multi-criterial decision making system and method," U.S. Patent 7 155 423, December 26, 2006.

[5] R. Subbu, P. Bonissone, N. Eklund, S. Bollapragada, and K. Chalermkraivuth, "Multiobjective financial portfolio design: A hybrid evolutionary approach," in *Proceedings of the 2005 IEEE International Congress on Evolutionary Computation*, 2005.

[6] N. S. Iyer, "A family of dominance filters for multiple criteria decision making: Choosing the right filter for a decision situation," Ph.D. dissertation, The Ohio State University, 2001.

[7] M. S. Pinkstaff, "An approach to disrupting communication networks," Master's thesis, Air Force Institute of Technology, March 2001.

[8] K. Deb, *Multi-Objective Optimization using Evolutionary Algorithms*. Wiley, 2001.

[9] J. Edmonds and R. M. Karp, "Theoretical improvements in algorithmic efficiency for network flow problems," *Journal of the ACM*, vol. 19, no. 2, pp. 248–264, 1972.

[10] E. Zitzler, L. Thiele, and K. Deb, "Comparison of multiobjective evolutionary algorithms: Empirical results," *Evolutionary Computation*, vol. 8, no. 2, pp. 173–195, 2000.

[11] J. Strange and M. T. Hopgood, Jr, *Centers of Gravity and Critical Vulnerabilities: Building on the Clausewitzian Foundation So That We Can All Speak the Same Language*. DIANE Publishing Company, May 2001.