Mission Planning for Joint Suppression of Enemy Air Defenses Using a Genetic Algorithm

Jeffrey P. Ridder
System of Systems Analytics Corporation
14900 Conference Center Drive, Suite 375
Chantilly, Virginia 20151
jridder@sosacorp.com

Jason C. HandUber
System of Systems Analytics Corporation
14900 Conference Center Drive, Suite 375
Chantilly, Virginia 20151
jhanduber@sosacorp.com

ABSTRACT

In this paper we present a genetic algorithm applied to the problem of mission planning for Joint Suppression of Enemy Air Defenses (JSEAD) in support of air strike operations. The stochastic nature of JSEAD scenarios and the complexity of JSEAD operations and interactions make this an especially challenging problem within the military domain. JSEAD planners and analysts stand to benefit from any advances in tools that address this problem. While our interest in this subject is broad, in this paper we are specifically investigating methods for developing robust plans that include routes for JSEAD assets, target types, firing ranges, and take off time, subject to multiple objective functions that capture different aspects of mission performance. The multi-objective optimization is performed by the Dynamic Non-Dominated Sorting GA (DNSGA), a non-elitist variant of NSGA-II. The objective functions are evaluated using a stochastic agent-based JSEAD simulation, and we assess the quality of mission plans produced by the GA in a set of test scenarios. The results from these tests indicate that our approach has significant promise as a component of a JSEAD mission planning tool.

Categories and Subject Descriptors: I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search – *Plan execution, formation, and generation*; I.6.3 [Simulation and Modeling]: Applications

General Terms: Algorithms

Keywords: Genetic algorithms, mission planning

1. INTRODUCTION

The role of Joint Suppression of Enemy Air Defenses operations is to apply lethal and non-lethal means of neutralizing enemy air defenses in order to enable air operations in enemy airspace. They are critical enabling capabilities of modern air power. Lethal means of air defense suppression include anti-radiation missiles (ARMs) that home on radio frequency emissions, standoff weapons, and guided or free-falling bombs. Non-lethal means include electronic attack — jamming of communications and radars. These capabilities are employed against an enemy integrated air defense system (IADS) in order to degrade its effectiveness, preventing or delaying detection and engagement of friendly aircraft. An IADS has a similar, but opposite role, to that of JSEAD forces — to prevent the destruction of targets by killing or deterring the aircraft that would threaten them. Because of the potentially large number of interacting units of various types, JSEAD is a complex problem that is difficult to analyze in significant depth. There are many potential applications for evolutionary computation tools in the military domain in general and the JSEAD domain in particular. In this paper we focus on the application of evolutionary computation to JSEAD mission planning.

Mission planning activities require an assessment of capabilities and threats in order to develop a mission plan. This assessment usually involves application of tools which apply limited analyses or time tested rules-of-thumb. Mission planning begins with an assumed level of *a priori* knowledge of the environment, including 1) The mix of friendly (Blue) capabilities is known, although the numbers of particular units and the course of action for employing the capabilities may be determined by the planning process, and 2) The locations of enemy (Red) units are partially known with some degree of uncertainty. For mission planning our goal is to create robust mission plans with good performance across a range of Red variations that reflect the uncertainty of the environment.

The purpose of our work is to develop a proof-of-concept prototype of an advanced JSEAD mission planning tool that uses a genetic algorithm to search for robust plans with good performance. We don't expect an operational tool to replace a human mission planner. Nor do we expect that a mission planner would implement the results of the GA without modification. Therefore, our goal is not to develop a GA that produces near perfect operational solutions, but rather to develop a genetic algorithm that rapidly and reliably produces solutions that provide the mission planner with useful insights to help construct a robust mission plan. The GA would assist the mission planner by helping to answer a number of questions, such as: How many JSEAD assets of each type do I need to use? Where should they be employed? What are their targets or target areas of interest? How should they approach and when should they arrive at their target areas of interest? How should they respond

GECCO'05, June 25–29, 2005, Washington, DC, USA.

to the environment? The answers to these and similar questions would greatly help the mission planner to develop a robust mission plan.

2. RELATED WORK

Other work using genetic algorithms for military problems related to mission planning includes that of Yilmaz et al. [7], who sought to use a GA to design optimal teams of sensors to detect enemy radars, Grace et al. [2], who implemented a genetic algorithm as one of four optimization algorithms in a prototype GPS jammer planning tool, Soliday [6], who used a genetic algorithm for determining the paths of UAVs in order to complete a number of surveillance tasks, and Louis et al. [3] who used a genetic algorithm to dynamically reallocate strike force resources in order to assist real-time targeting and retargeting.

The current work is most closely related to a problem treated by Ridder [5] in the Advanced Reactive Electronic warfare Simulation (ARES). The ARES simulation is a simulation of airborne electronic attack (AEA), one of the critical components of JSEAD, that featured detailed models of terrain, radio frequency signal detection, jamming, networked communication, enemy radars, etc. It included a GA that would search a trade space of AEA platforms, systems, jamming targets, and strategies. However, due to ARES' level of detail, single complete GA runs could require over a day to execute on a 96 node Beowulf cluster. Clearly, mission planning applications require solutions in significantly less time.

Our current focus is on making advances in the mission planning domain by striking a balance between GAs using long-running, high resolution simulations like ARES and fast-running, low resolution simulations that yield operationally uninteresting results. We believe this balance can be achieved by sacrificing the notion that the GA-based mission planning tool will directly provide operationally useful mission plans in exchange for meaningful insights produced in a timely manner. This allows us to use a fast-executing agent-based simulation that emphasizes modeling of systems' interactions and their effects, rather than detailed computational models that attempt to predict systems' effectiveness. In addition to improvements in run time, this approach allows us to focus on a richer set of interactions than might otherwise be possible.

3. PROBLEM ENVIRONMENT

We have constructed a simulation of a JSEAD environment using the MASON agent-based modeling toolkit [4]. This simulation models two sides — Red and Blue — where each side has two basic types of entities: targets and air defense units for Red, and strikers and JSEAD units for Blue. Red targets are stationary throughout the simulation and carry point values to support objective function calculation. Red's air defense units are of one of three types: early warning sites, surface-to-air-missile (SAM) sites, and surface-to-air missiles. Early warning sites contain a single, long range radar which emits periodically, detecting any aircraft within its detection range. For all radars, an aircraft is considered tracked if it is detected on two consecutive scans. Early warning site tracks are passed to SAM sites as a cue. SAM sites contain two radars: 1) a target acquisition radar that may emit periodically or in response to cues from early warning sites, and 2) a target tracking radar that remains off until cued by the co-located target acquisition radar. A SAM site may fire a SAM at tracked aircraft that are within a given range (limited by the lethal maximum and minimum ranges of the SAM). The target tracking radar must continuously track aircraft in order for the SAM to reach the endgame (i.e., the point at which a random number generator determines whether the aircraft is killed based on the SAM's probability of kill). Jamming aircraft are an exception and may also be fired upon using the home-on-jam feature of the SAMs. That is, if a SAM site detects jamming it can fire a SAM in a mode that allows it to home on jamming without need for any additional tracking guidance.

Blue strikers fly to user-assigned targets using five route points. JSEAD units include weasels and jammers. Weasels are aircraft similar to strikers, except that their targets are only air defense sites and they fire anti-radiation missiles (ARMs) at emitting targets of a designated type. Jammers also have five route points, of which any but the first could be a terminal orbit point (i.e., the jammer will stop and orbit at that point for the remainder of the simulation). Jammers impact the detection range of enemy radars based on proximity and detection. They reduce the detection range of all radars as a function of $1/R_j^2$, where R_j is the range from the jammer to the victim radar. Additionally, narrowband reactive jamming is used to focus twenty times more jamming energy at radars that are detected by the jammer. Radars are detected when they are within a specified electronic support measures detection range of the jammer. Finally, if a jammer detects a SAM site and determines that it is within maximum lethal range of the SAM, it will disengage and move to an orbit position outside of the SAM's lethal range for the remainder of the mission.

The simulation is stochastic with a number of random characteristics. Both SAMs and ARMs have attributes of probability of kill which determine their effectiveness on an event-by-event basis. The initial emission state of radars is randomized. Also, all Red units are assigned a location centroid and uncertainty such that the actual position of any Red unit in any simulation run is randomly generated with a Gaussian distribution using the location uncertainty as the standard deviation relative to the centroid.

We've implemented a simple rule system for many of our agents. Although the rules are fixed, several parameters in the rules are subject to manipulation by the GA, significantly impacting behavior of units. For Red air defense units the rules concern conditions that would cause radars to initiate or terminate emission, such as the amount of elapsed time, the range of the closest track, or the range to the closest cue. In addition, SAM sites have rules for the conditions under which they will fire a SAM. For Blue JSEAD units, weasels have rules with conditions that determine when to return to base, when to fire an ARM, when to approach a potential target, and when to orbit a location in anticipation of a target emitting. Jammers have a single rule (with no parameters) that causes them to move outside of a detected SAM's lethal range and initiate orbit. The effectiveness of Red air defense units is governed by the amount of location uncertainty and settings of rule parameters, including elapsed time and range limits. The effectiveness of Blue JSEAD units is influenced by routes, timing, and weasel rule parameters, including target types and firing ranges.

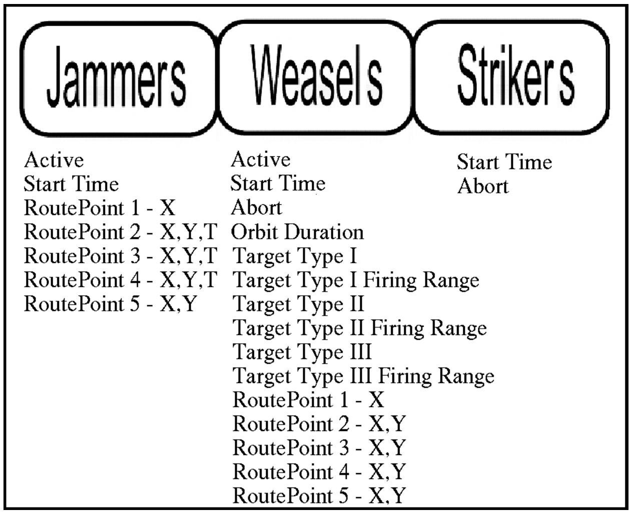

Figure 1: Problem representation.

4. GENETIC ALGORITHM

4.1 Problem Representation

Each individual in the population specifies all Blue mission variables as a vector of real-valued genes. As shown in Figure 1, the first segment of the genome contains the genes for each of the jammers, grouping each jammer's genes contiguously. The second segment contains the genes for the weasels and the third segment contains genes for the strikers. For example, a simulation using one jammer, two weasels, and three strikers would have a genome with 58 genes. Since our current focus is on planning JSEAD in support of a strike, most aspects of the strikers' mission plan are fixed, allowing some flexibility in the strikers' start times (to permit phasing with JSEAD units) and abort conditions. The JSEAD units (weasels and jammers) contain route point genes which determine their flight path through the scenario. Jammers have additional route point type genes that determine whether a point is a waypoint (fly through) or orbit (terminal point). Weasels, like strikers, have an abort threshold gene which determines the number of SAMs that must be actively tracking them before they'll abort their current target and either divert to their secondary target or return to base. Weasels have additional genes for determining their target types (radar classification), their relative priority (target type 1 will be pursued if available before target type 2), and the range at which they will fire at that target type (closer means less time to impact, reducing the time the radar has to turn off and increasing the chance of success, but also increasing the risk that a SAM will be fired in return). All JSEAD units also have an activation gene that determines whether they will be used at all in the mission plan.

4.2 Multi-Objective GA

The problem of mission planning is fundamentally multi-objective, where the goal is to strike a balance between the components of mission success: 1) maximizing operational objectives achieved, 2) minimizing attrition (loss of friendly forces), and 3) minimizing resource cost. In principle, these objectives could be combined into a single fitness function. However, this would require the mission planner to perform the difficult task of assigning relative importance to each of these components of mission success — a difficult judgement call with potentially lethal consequences. Furthermore, in practice we found the combined objective approach to be problematic since quite often one component of mission success would dominate early in the GA run, resulting in solutions clustered near an undesirable local optimum. For example, in one of our less successful early runs we found that it was difficult to get aircraft to fly on their assigned missions. The GA quickly rewarded the notion that the best way to survive was to not fly at all, resulting in a solution where all aircraft survived, but no operational objectives were achieved. Upon modifying the objective function to reward aircraft for flying we found that all left their airfields then immediately declared an abort and returned to base. Apparently the survival instinct is alive and well in digital warfare. Further modification of the combined objective function to provide incentives for achieving various phases of a mission proved to be a futile exercise in attempting to achieve a balance between the many sub-objectives. The difficulty of achieving this balance in practice, combined with the benefits for a mission planner to be able to choose from multiple plans along the Pareto front led us to abandon this approach in favor of a multi-objective GA.

We modified the NSGA-II [1] non-dominated sorting GA to make it suitable for use in uncertain, noisy environments such as that modeled by the JSEAD simulation. The NSGA-II algorithm performs a non-dominated sort of a combined population of elites with the evaluated population (creating a double-sized population), then discards the worst half, keeping the remainder as both the next elite population and as the parent population for the next generation. This works well as long as the fitness values assigned to each member of the elite population remain constant throughout the GA run. However, the elite population becomes problematic for noisy or dynamic problems since elite individuals retain their approximated scores from a single evaluation that may inflate their true fitness. Instead, our modification to NSGA-II, which we call the Dynamic Non-Dominated Sorting GA (DNSGA), is to implement a steady-state approach with 50 percent replacement. All individuals must be evaluated in each generation followed by a non-dominated sort (there is no separate elite population). The lowest ranking half of the population is discarded and replaced by children produced by parents selected from the best half. For non-noisy problems DNSGA requires evaluation of populations that are twice the size of that required by NSGA-II. All other aspects of the algorithm are the same as NSGA-II.

4.3 Fitness Evaluation

The results presented in this paper were produced using the DNSGA algorithm with three objective functions. These are:

1. Targets at risk — the accumulated points of all targets that had weapons released at them.

2. Attrition — the accumulated points of all aircraft that were destroyed by enemy air defenses.

3. SEAD cost — the accumulated points of all JSEAD aircraft that participated in the mission.

Before each generation we generate a set of N variations of a JSEAD simulation scenario. Each Blue individual is evaluated using the same set of scenario variations and the fitness assigned to each individual for each objective function is the average for the set.

One technique that we found helps the development of good solutions is to not implement the SEAD cost objective until a sufficient number of generations has elapsed. Prior to this point we force activation of each JSEAD unit through the "active" gene so that they are forced to fly and find solutions that contribute to, or at least do not detract from (via attrition), mission success. The cost phase then reduces the number of JSEAD units to find cost-effective solutions.

4.4 Selection and Genetic Operators

We use Deb's multi-objective tournament selection that is integral to NSGA-II. The NSGA-II algorithm uses a tournament of size 2, performing selection based on Pareto rank or, if both individuals are of the same rank, the individual with the highest "crowding distance" (a diversity measure integral to NSGA-II) is selected. We also use single-point crossover, but restrict it to non-route point segments of the genome so that, in effect, the GA treats the route point segments as a single gene. We use Gaussian mutation on continuously real-valued genes (e.g., ARM firing range for weasels), flip mutation on discretely valued genes (e.g., target emitter type), and a special mutator for route points. If the mutator determines that a route point gene is to undergo mutation (the only way that route point genes are modified since crossover is prohibited), then it executes a route nudger that randomly selects a single route point and nudges it using a Gaussian distribution. In nudging the point it ensures that the maximum range constraint of the aircraft is enforced (i.e., no infeasible routes are generated). Finally, if the nudged route point is for a jammer, then the route point type (waypoint or orbit) is determined by flip mutation.

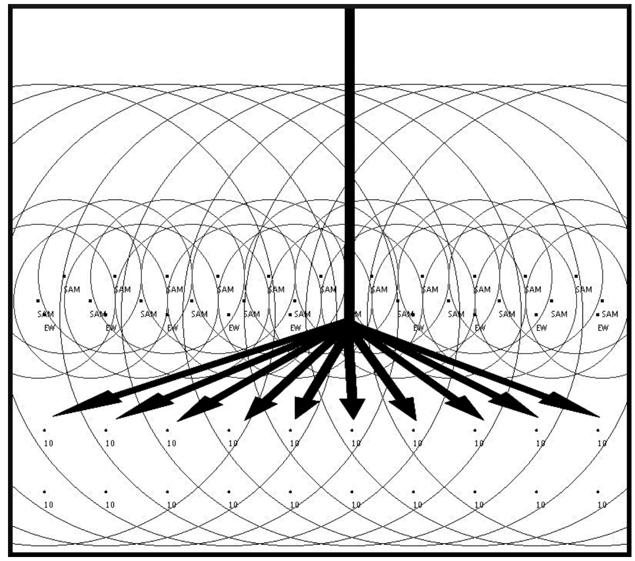

Figure 2: Gauntlet scenario.

Table 1: Scenario cases

Case	Air Defense Location Uncertainty (Nautical Miles)	SAM Site Aggressiveness
A	0	Low
B	20	Low
C	40	Low
D	0	High
E	20	High
F	40	High

5. SCENARIOS

In this paper we consider six cases for each of two scenarios. The first scenario we refer to as the Gauntlet, illustrated in Figure 2. It is an operationally unrealistic, but challenging scenario that features a line of staggered air defense units (23 SAM sites and 10 early warning sites) providing overlapping coverage to each other and protecting the point-valued targets in the rear. The Blue strike plan is to fly through the center of the air defenses, get into their rear, then disperse to their assigned targets. This requires the JSEAD units to create a corridor of protection for the strikers through the middle of the Gauntlet. A perfect solution would allow the ten Blue strikers to achieve a targets at risk score of 100 points with no attrition and require no JSEAD units.

The second scenario, illustrated in Figure 3, we refer to as Point-defense. The Point-defense scenario is more operationally realistic than the Gauntlet and features ten locations, each with two targets, two SAM sites, and one early warning radar in close proximity. In addition, there are three SAM sites acting as "free agents" that have a large location uncertainty (100 nautical miles) in every case. The Blue strike plan is for ten strikers to enter on a single axis, slanting in from the upper left, and departing from that axis to approach targets with minimal exposure to SAMs. The JSEAD units are significantly challenged by this scenario since they must protect both the main ingress axis and the individual target approach axes of the strikers.

We ran six different cases for each of these scenarios, where each case models a particular location uncertainty and level of aggressiveness of Red air defense units. Table 1 lists the scenario cases. Regardless of the location uncertainty for all other air defense units, all Point-defense cases include 3 SAM sites with a location uncertainty of 100 nautical miles (free agents). For the low aggressiveness cases, the SAM sites do not emit until they receive a cue from the early warning radars that a target aircraft is within detection range of the SAM site target acquisition radar. The high aggressiveness cases differ in that the SAM site target acquisition radars will also emit periodically in an attempt to detect target aircraft independently of the early warning radars.

6. RESULTS

The results presented here are from one typical run of each case for both scenarios to illustrate how a mission planner would use this information to gain insights for constructing the mission plan. For each case, Red targets and Blue strikers and weasels are valued at ten points each, and each jammer is valued at fifty points (consistent with the notion that jammers are high value, high demand assets). For example, a plan that results in eight targets placed at risk, loss of one striker and one jammer, and use of two jammers and

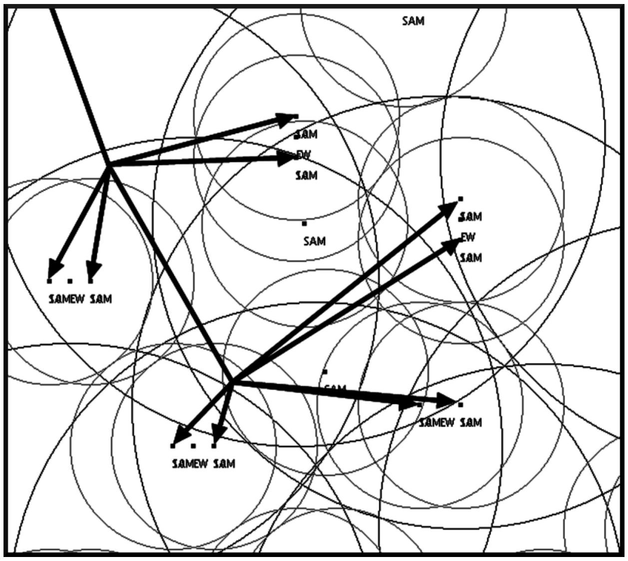

Figure 3: Point-defense scenario.

three weasels, has fitness scores of (80, 60, 130) for targets at risk, attrition, and SEAD cost. Each case was run using a population size of 200 individuals for 300 generations with fitness being approximated as the average score using ten scenario variations, where the variations differ from each other in terms of pseudo-randomly generated locations of Red units. The pool of available JSEAD units consisted of five jammers and ten weasels, for a maximum SEAD cost of 350. The first 150 generations forced all JSEAD units to be active, resulting in a constant, maximum SEAD cost during this period. The remaining generations added SEAD cost as a fitness measure by allowing the GA to manipulate SEAD resource selection. The crossover probability was 0.8 and the mutation probability was 0.01.

6.1 Gauntlet Scenario

Figures 4–6 illustrate the results from runs of cases A–C of the Gauntlet scenario. Each of the figures plots the non-dominated individuals in the final population by their targets at risk and attrition scores versus the SEAD cost. To identify the non-dominated individuals, we first evaluated each individual in the final population against 1000 random scenario variations to determine the average and standard deviation scores for targets at risk and attrition (SEAD cost is invariable for a given individual). We then did a non-dominated sort using the average objective function scores to identify the individuals on the Pareto front. The average scores are marked by symbols and the standard deviation by error bars. The standard deviation provides a measure of robustness of a particular individual in that objective function, where smaller error bars indicate a greater tolerance to variation. Each attrition score for a given SEAD cost corresponds to a targets at risk score from the same non-dominated individual at the same SEAD cost. There is usually only one non-dominated individual for each SEAD cost and, therefore, the targets at risk and attrition scores at that SEAD cost correspond directly to that individual.

The figures show an expected trend adequately captured by the GA. That is, as scenario uncertainty increases, the

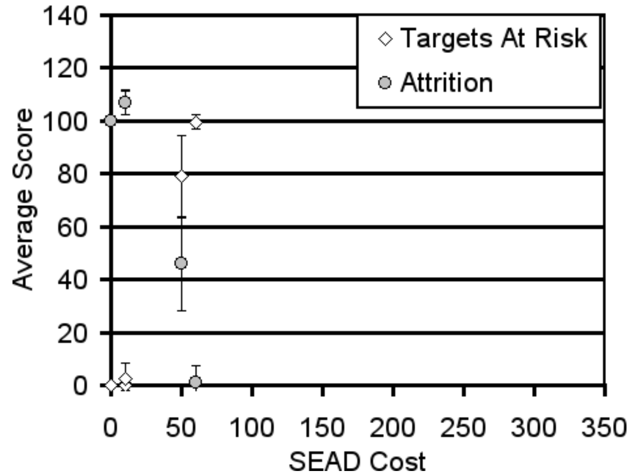

Figure 4: Non-dominated individuals from Gauntlet scenario case A.

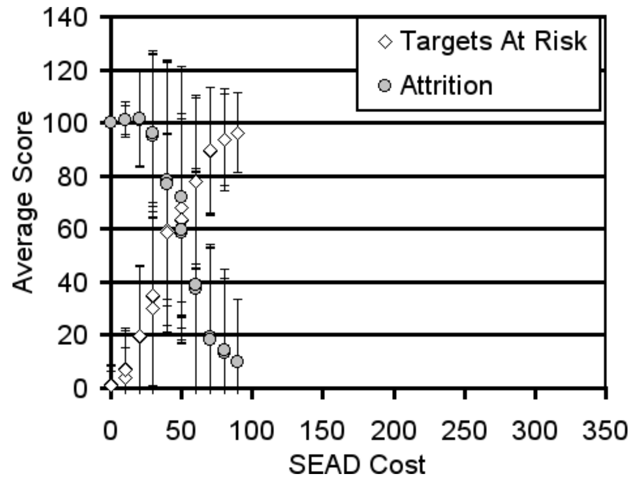

Figure 5: Non-dominated individuals from Gauntlet scenario case B.

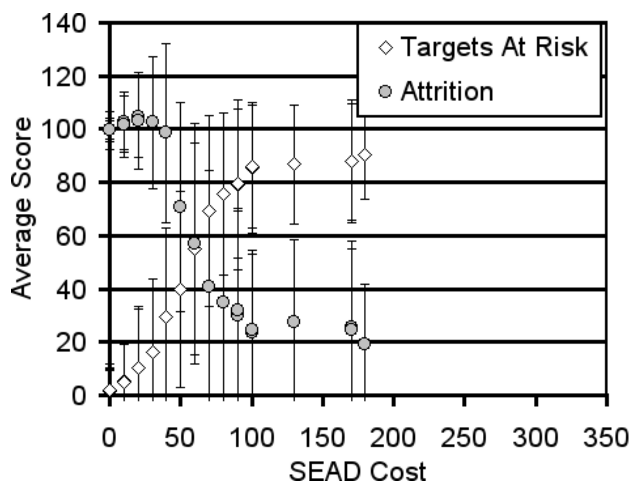

Figure 6: Non-dominated individuals from Gauntlet scenario case C.

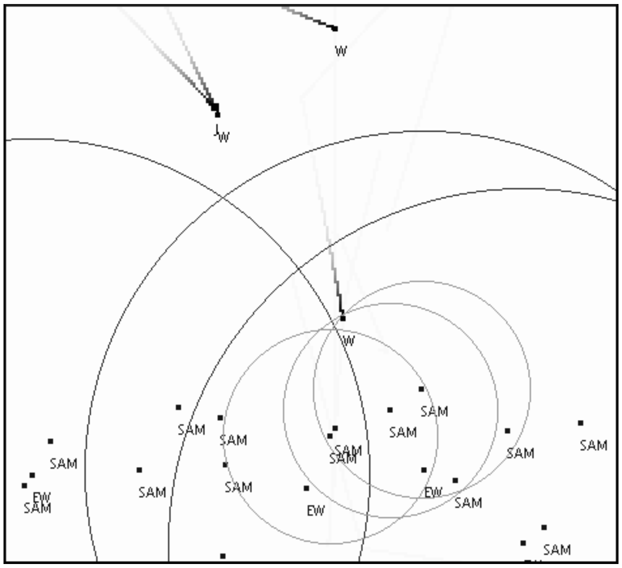

Figure 7: Case B of the Gauntlet scenario before attack by SEAD forces implementing a plan with SEAD cost of 80.

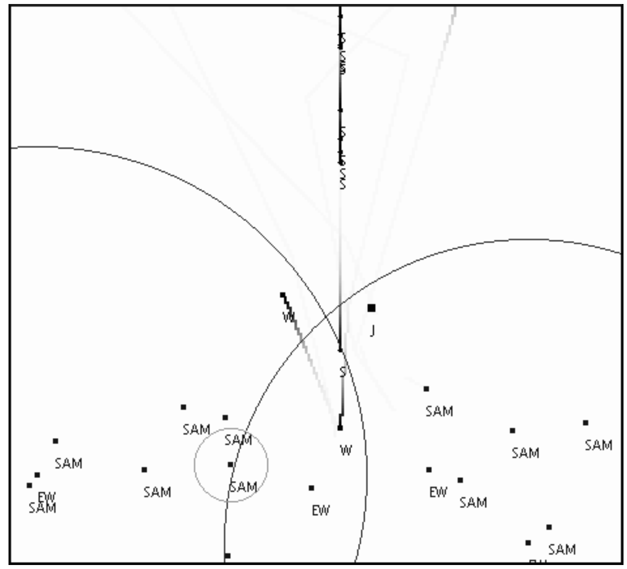

Figure 8: Case B of the Gauntlet scenario after successful attack by SEAD forces.

SEAD cost to handle that uncertainty also increases. Figure 4 shows that when Red locations are known (case A) the Gauntlet scenario requires a SEAD cost of at least 50 points (corresponding to one jammer) in order for Blue to be effective. With a SEAD cost of 60 points — one jammer and one weasel — Blue can expect excellent mission performance resulting in low attrition and high targets at risk. Note also that although the locations of all units are known, other stochastic factors result in significant standard deviations for the plan with SEAD cost of 50. As location uncertainty increases in case B (Figure 5), a minimum SEAD cost of 70 (one jammer and two weasels) is required to keep attrition below 20 points, but increasing SEAD cost to 90 by adding weasels results in further reducing attrition and increasing targets at risk to nearly 100, while also increasing robustness (reducing the standard deviation indicated by the error bars). Clearly, a mission planner would have to decide whether the additional SEAD cost is worthwhile, but is now armed with the information necessary to justify such a decision. Figure 6 shows that case C — with the highest location uncertainty — results in degraded attrition, targets at risk, and robustness even with higher SEAD costs than for case B. A SEAD cost of 100 is sufficient to achieve results near the best possible, although the highest SEAD cost improves robustness.

We expect that figures like those presented here would guide a planner to the set of plans that correspond to an attractive range of SEAD costs. For example, if case A is expected, then the planner would consider plans from case A with SEAD costs of 60. If case B is expected, then the planner would focus on plans with SEAD costs in the range of 70 to 90.

Figure 7 shows ingressing SEAD forces implementing a plan with SEAD cost of 80 against an example of case B of the Gauntlet scenario. Figure 8 shows the same scenario after the SEAD forces have done their work, clearing the middle of the Gauntlet for the ingressing strikers. The weasels (marked with "W"), jammers ("J") and strikers ("S") approach the Gauntlet from the top of the figure. This plan provides a number of interesting insights and ideas for a mission planner, beginning with the use of a weasel to lead the strike and clear any forward SAM sites so that a jammer can orbit closer to deeper SAM sites. This weasel fires at SAM target tracking radars from close to max ARM range, but cleverly ingresses just deep enough into the SAM envelope to cause the victim radar to emit long enough for the ARM to home to and destroy the target. The second and third weasels are focused on SAM sites more directly in their path (ARM firing range is much shorter than for the first weasel). Their function is to "blow holes" in the defense for the strikers. The second weasel escorts the jammer, leading it slightly as the jammer takes up its orbit position, then strikes SAM sites in the middle of the Gauntlet. The third weasel lags the others, but precedes the strikers, attacking any remaining SAM sites that might remain in their path. This plan provides other insights and ideas to a mission planner, but those discussed here adequately show that the GA is capable of discovering innovative and operationally useful ideas to be implemented in real mission plans.

6.2 Point-defense Scenario

Figures 9–11 show the results from runs of cases A–C of the Point-defense scenario. These results show different trends from the Gauntlet scenario results, consistent with the different approach to protecting targets. First, as with the Gauntlet, the least attrition occurs when air defense locations are known (case A). However, as might be expected, considerably more SEAD resources are required in this scenario than in the Gauntlet. Cases B and C have similar high attrition (around 50), although case C requires fewer SEAD resources to achieve this level of attrition. Also, with higher SEAD resources (above 250 SEAD cost), case C has greater targets at risk than case B. Contrary to the Gauntlet scenario, case B is better for Red and worse for Blue than case C. In case B the location uncertainty leaves the SAM sites close enough to protect their targets, while case C's greater uncertainty would occasionally leave targets unprotected.

Figures 12–14 show the results from runs of cases D–F. Recall that these cases correspond to cases A–C in terms of location uncertainty, but that the SAM sites are more aggressive. Comparing case D with case A shows only minor differences in average scores, although robustness in both targets at risk and attrition are significantly better in case D. Comparing cases E and F with cases B and C shows that higher SAM site aggressiveness signficantly reduces Blue's attrition and improves robustness, apparently due to greater destruction of threatening SAM sites. In fact, further comparing these cases with cases A and D shows that more aggressive emission effectively negates any Red advantage gained by location uncertainty, and improves Blue's robustness as well. These results are very interesting and are consistent with the emission control (EMCON) doctrine, which states that air defense units should not emit unless necessary since emission exposes one's location to the enemy with potentially lethal consequences. They also hint at the potential value of using measures, such as decoys, to encourage Red to emit.

7. CONCLUSIONS AND FUTURE WORK

The results presented here illustrate the potential value of this approach that provides mission planners with a set of non-dominated plans and the quantitative information needed to select particular plans for further consideration. This information includes assessments of both expected performance from the average score, and the robustness of the plans from the standard deviation. We've also illustrated how a selected plan can serve as a model to provide meaningful insights to the mission planner for construction of actionable mission plans.

While our results illustrate the value of this approach, they also lead us to consider the direction of future work. Planned future work includes refinement of the current algorithm, including GA parameters and development of techniques for adaptively varying the number of generations and number of scenario variations for objective function evaluation. Also, while our current approach develops mission plans that consider variations in locations of Red air defense units, we would also like our mission plans to be robust to variations in the numbers and types of air defense units, as well as the tactics they employ. Further, we'd like these variations to reflect points on Red's Pareto front. This suggests the potential value of a coevolutionary approach in which we simultaneously evolve Red's plan with that of Blue. This should lead to robustness across a greater range of variations, and could also provide insight into Red's best configuration and tactics.

A second path for further exploration is to reconsider the role of the mission planning system. Due to the dynamic nature of warfare, the mission plan should not be considered a script to followed, but rather a statement of mission goals with an initial plan to be adapted to achieve those goals. In that sense, the mission planning system should seek to produce plans that are both robust and flexible, meaning that we prefer plans that are designed to promote dynamic replanning. The GA-based mission planning system must account for onboard, dynamic replanning when developing mission plans. The GA may also select the most appropriate replanning algorithms (perhaps implemented as intelligent agents) for each platform as part of the mission plan. Such

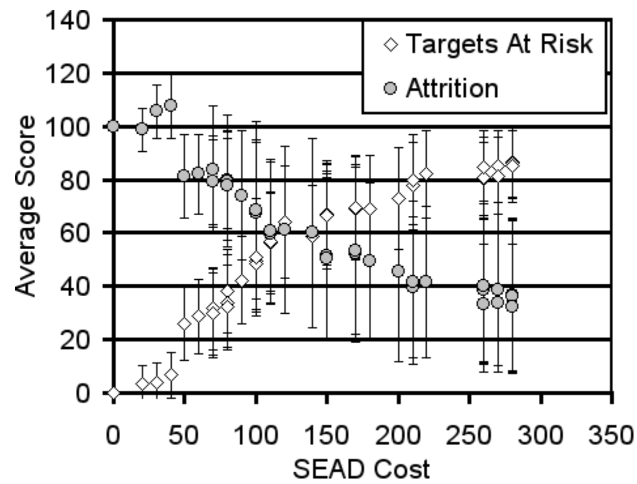

Figure 9: Non-dominated individuals from Point-defense scenario case A.

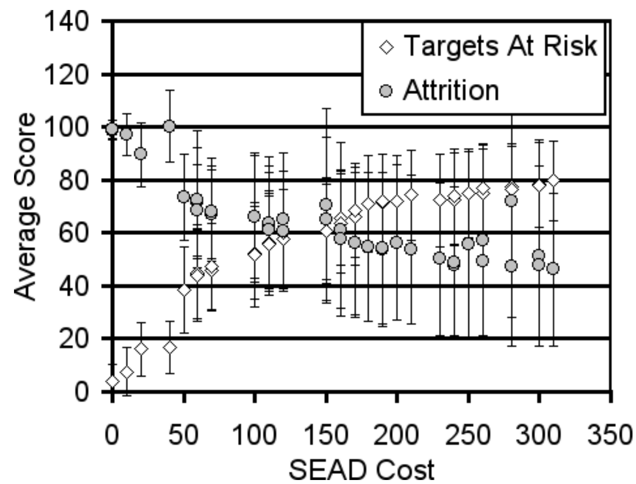

Figure 10: Non-dominated individuals from Point-defense scenario case B.

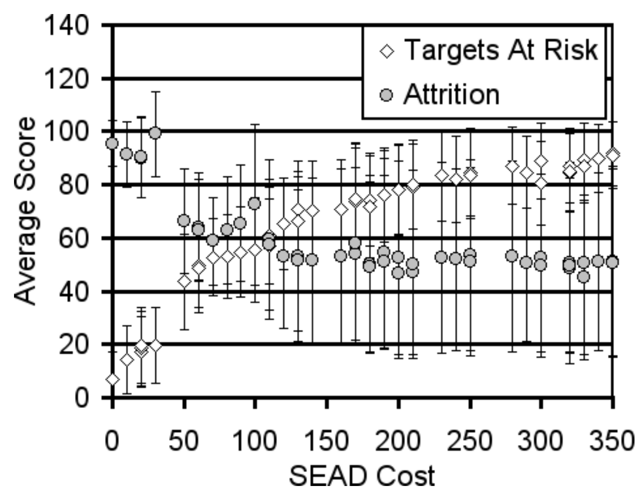

Figure 11: Non-dominated individuals from Point-defense scenario case C.

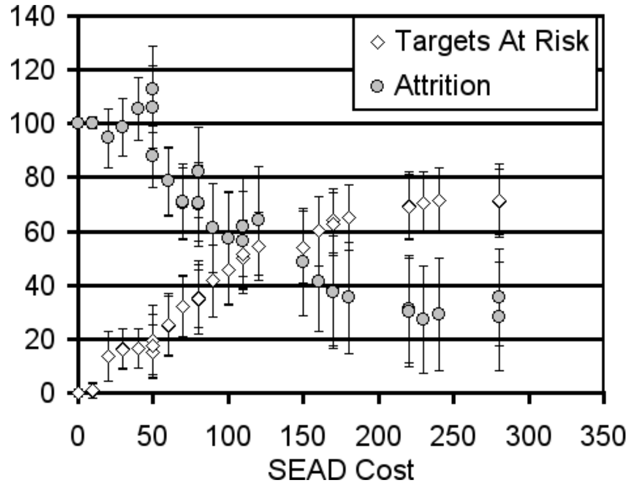

Figure 12: Non-dominated individuals from Point-defense scenario case D.

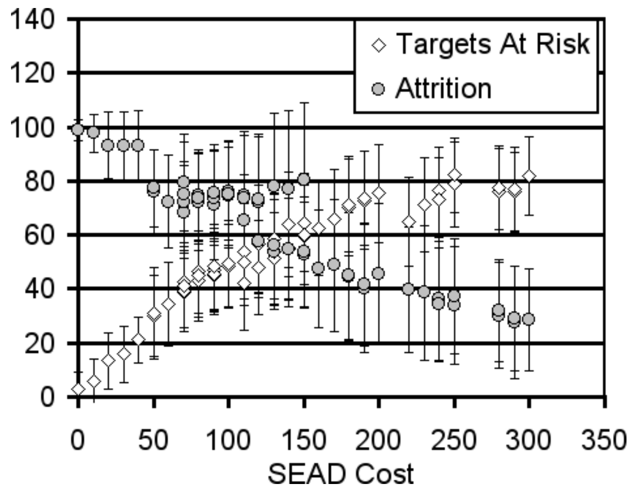

Figure 13: Non-dominated individuals from Point-defense scenario case E.

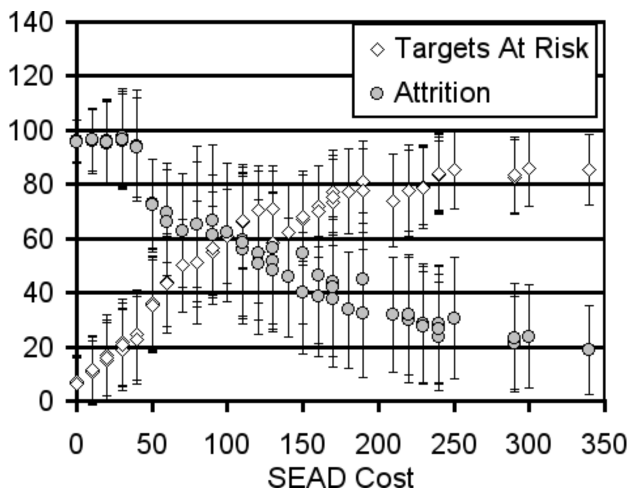

Figure 14: Non-dominated individuals from Point-defense scenario case F.

a capability could significantly improve the performance of the mission planning system and the units that implement and adapt the mission plan.

A third path for exploration is to expand this capability beyond mission planning and into the domain of campaign planning and control. Many of the results presented here, especially for the challenging Point-defense scenario, suggest unacceptable risks for some missions. A GA-based campaign planning tool would reduce risk by creating missions that set the stage for later missions. For example, such a tool might bias early missions toward pursuit and destruction of air defenses so that later strike operations would be more likely to succeed. As the environment changes, the campaign plan would automatically be adapted by the GA. Such a campaign planner/controller would be a significant extension of the capability presented here.

8. ACKNOWLEDGMENTS

This work was funded by System of Systems Analytics Corporation internal research and development.

9. REFERENCES

[1] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A Fast and Elitist Multi-Objective Genetic Algorithm: NSGA-II. Technical Report 200001, Kanpur Genetic Algorithms Laboratory, Indian Institute of Technology Kanpur, 2001.

[2] M. Grace, P. Stieber, S. Minarik, B. Tanju, and F. Escobar. An Automated Tool for Mission Planning in GPS-Denied Areas. In *Proceedings of Institute of Navigation's National Technical Meeting 2002 Conference*, pages 176–181, San Diego, CA, 2002.

[3] S. J. Louis, J. McDonnel, and N. Gizzi. Dynamic Strike Force Asset Allocation Using Genetic Algorithms and Case-Based Reasoning. In *Proceedings of the Sixth Conference on Systemics, Cybernetics, and Informatics*, pages 855–861, 2002.

[4] S. Luke, C. Cioffi-Revilla, L. Panait, and K. Sullivan. MASON: A New Multi-Agent Simulation Toolkit. In *Proceedings of the 2004 SwarmFest Workshop*, 2004.

[5] P. Pham, J. Ridder, and P. Pace. User's Guide: Advanced Reactive Electronic Warfare Simulation (ARES), Version 1.12. Technical Report NPS-EC-01-002, Naval Postgraduate School, March 2001.

[6] S. W. Soliday. A genetic algorithm model for mission planning and dynamic resource allocation of airborne sensors. In *Proceedings, 1999 IRIS National Symposium on Sensor and Data Fusion*, Laurel, MD, May 1999. Applied Physics Lab.

[7] A. S. Yilmaz, B. N. McQuay, H. Yu, A. S. Wu, and J. C. Sciortino. Evolving Sensor Suites for Enemy Radar Detection. In *Genetic and Evolutionary Computation Conference Proceedings, Part II*, pages 2384–2395. Springer, July 2003.