Applications of Multi-objective Evolutionary Algorithms to Air Operations Mission Planning

Brad Rosenberg, Marc Richards, John T. Langton, Sofya Tenenbaum, and Daniel W. Stouch
Charles River Analytics, Inc.
625 Mount Auburn St
Cambridge, MA 02138, USA

{brosenberg, mrichards, jlangton, stenenbaum, dws}@cra.com

ABSTRACT

Air operations mission planning is a complex task, growing ever more complex as the number, variety, and interactivity of air assets increases. Mission planners are responsible for generating as close to optimal taskings of air assets to missions under severe time constraints. This function can be aided by decision-support tools to help ease the search process through automation. This paper presents several applications of multi-objective evolutionary algorithms for discovering suitable plans in the air operations domain, including dynamic targeting for air strike assets, intelligence, surveillance, and reconnaissance (ISR) asset mission planning, and unmanned aerial systems (UAS) planning. Lessons learned from these studies are described and an exploration of potential future directions is discussed.

Categories and Subject Descriptors

I.2.8 [**Problem Solving, Control Methods, and Search**]: Heuristic methods

General Terms

Algorithms, Design, Experimentation

Keywords

Evolutionary algorithms, multi-objective evolutionary algorithms, air operations, mission planning

1. INTRODUCTION

Air operations mission planning is a complex process, beginning with translating the commander's intent into a Joint Air Operations Plan (JAOP) through the Joint Air Estimate Process (JAEP) [17]. Planners within the Air Operations Center (AOC) must efficiently improve situational awareness through intelligence preparation of the battlespace (IPB) and generate, analyze, and compare various courses of action (COAs) before one is selected and incorporated into the JAOP. As this moves to the operational environment, the JAOP provides the basis of the daily tactical tasking of air assets under the Air Tasking Order (ATO). Mission planners are responsible for driving the ATO development cycle, tasked with optimally allocating available resources to produce desired effects.

This operational task brings about a significant technical challenge in identifying and implementing the appropriate search technique to generate suitable mission plans. We outline several aspects that characterize the problem requirements and associated search spaces.

First, air operations occur in dynamic environments, with changes to the battlespace occurring rapidly. As a result, mission planners have limited time available for generating, evaluating, and executing a selected plan. Currently, the ATO operates on a 24-hour basis with an overlapping 72-hour development cycle [18]. This process is currently being revised to reduce the kill-chain cycle, transforming it into a continuous or "rolling" ATO. As this process shortens, mission planners will have even less time to generate plans. Automated methods to generate these plans will support mission planners in maintaining this expedited cycle.

Second, the number of assets and tasks is extremely large. Assignment of assets to these tasks results in a combinatorial explosion, classifying it as a combinatorial optimization problem (COP). Exhaustive search of possible plans is, therefore, intractable even for longer planning cycles.

Third, valid mission plans must obey certain hard constraints. An example of a geospatial constraint may include no-fly zones. Temporally, an asset may not be able to reach a target within an associated time window. Other asset capability factors, such as fuel, on-board munitions, speed, etc. can limit the feasibility of potential mission plans. Thus, portions of the search space are discontinuous. As a result, many gradient search methods are rendered insufficient.

Fourth, in addition to hard constraints, there are specific objectives a mission planner is looking to optimize. In targeting, this might be maximizing the probability of success in prosecuting a target. For intelligence, surveillance, and reconnaissance (ISR), a mission planner might desire to maximize sensor coverage. Sometimes, objectives will be in conflict with one another. For example, the asset with proper munitions to successfully prosecute a target might have to route through a hostile portion of the battlespace, increasing its vulnerability. The existence of multiple conflicting objectives produces a complex landscape, one not easily traversed by gradient search algorithms. As a result, naïve search methods will fall victim to becoming stuck in local optima.

This combination of limited decision-making time and a large, complex search space burdens a mission planner in an already stressful environment. Automated systems can assist in this process by generating a selection of mission plans for the decision-maker. We now present our findings in developing automated systems to support air operations mission planning staff and explore directions for future research and development.

2. FRAMEWORK

To address this issue, we have developed a collection of decision-support tools [11][15][16] to assist mission planners in discovering suitable plans across several domains. The underlying domain-independent framework consists of a method to define an operational scenario, an optimization engine to generate a diverse set of solutions, and a suite of visualization and analysis tools to review, analyze, and visualize generated plans before making the ultimate selection. The optimization engine employs multi-objective evolutionary algorithms (MOEAs) as a base search technology, coupled with a low-fidelity simulation to act as an "abstract wargamer" to evaluate potential plans (see Figure 1).

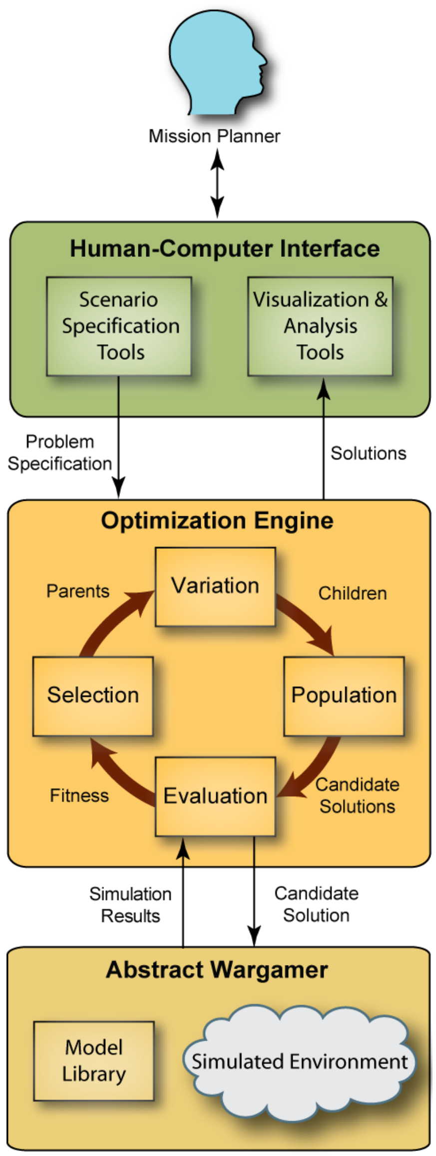

Figure 1: Framework for Air Operations Planning

2.1 Multi-objective Evolutionary Algorithms

Evolutionary algorithms (EAs), and in particular, multi-objective evolutionary algorithms, have been successful as a general search strategy for solving large, complex problem spaces similar to those of air operations mission planning [2]. Several methods are available to handle multi-objective optimization using evolutionary techniques. [5] classifies these into first and second-generation methods. First-generation methods include aggregation, a simple approach where all the objectives are combined into a scalar value (e.g., weighted sum). While being simple and easy to implement, this approach can fail to generate optimal solutions for given search spaces. Much work has been performed since then to extend this work to second-generation methods that overcome these limitations [8][9][20]. One of the more popular algorithms currently used is the Non-dominated Sorting Genetic Algorithm II (NSGA-II) [7].

NSGA-II was selected for our framework due to its speed, ease of implementation, and ability to maintain a diverse set of solutions through non-dominated sorting and crowding distance estimation. NSGA-II was developed by Deb in response to criticisms of the original NSGA algorithm, one of the first multi-objective evolutionary algorithms created. The revised form reduces computational complexity, includes a form of elitism, and removes the need for a sharing parameter.

2.2 Abstract Wargaming

Given limited available time in generating suitable mission plans, rapid evaluation of candidate plans is required. As evolutionary algorithms require a substantial number of evaluations, it is key that we mitigate the bottleneck imposed by the computational complexity of the evaluation process.

Within the framework, rapid evaluation is performed using a low-fidelity simulation environment to act as an abstract wargamer. Within the defense community, wargaming is the process of simulating a potential course of action, alternating back and forth between responses between factions, to determine potential outcomes. The rise of modeling and simulation (M&S) in the defense sector has brought about computer-based simulations to act as wargaming technologies, using either constructive simulations to provide fully automated behaviors for each entity or virtual simulations which require human-in-the-loop (HITL) interaction to provide the necessary behaviors.

Where possible, our abstract wargamer uses an agent-based modeling (ABM) approach. Agent-based modeling has gained wider acceptance in recent years given increases in computational processing power. In addition it has the distinct advantage over classic mathematical models (e.g., Lanchester equations) [10] of being able to directly model the entities in a system. This allows for the emergence of the system-level effects from the low-level interactions of distinct entities. This can be advantageous when attempting to model a complex system such as air operations.

For our framework, we have leveraged the MASON multiagent simulation toolkit [13] to allow us to capture the key variables for coarse-grained simulation. MASON was selected given its strong performance characteristics, versatility, and friendly licensing options for embedding in governmental and commercial systems. Using an air operations model built on top of MASON, we are able to run thousands of simulations within minutes.

3. APPLICATIONS

Using this framework, we have developed several decision-support tools for the defense community for air operations mission planning. Our first application addressed dynamic targeting for air strike assets. We extended this work for the ISR planning domain and are generalizing this approach to support UAS mission planning. We now present our efforts in developing these applications.

3.1 Dynamic Targeting for Air Strike Assets

The first application developed using this framework was in assisting the dynamic targeting cell within an AOC in retasking air strike assets when pop-up targets emerge that need to be immediately addressed. When these targets emerge, planners within the dynamic targeting cell must rapidly generate and evaluate retasking alternatives with a given set of constraints. Questions asked are:

- What assets are available to prosecute the target?
- What assets can get to the target within the window of opportunity?
- What assets have munitions on-board that can successfully prosecute the target?
- What are the repercussions of this retasking?
- How do I address the cascading effects?

As this process occurs during the real-time execution of an ATO in an operational environment, decisions must be made within minutes. As a result, planners can be greatly assisted by the use of automated systems to aid in determining options.

To address dynamic targeting, a mission planner wishes to achieve several objectives. A subset of those included in our application were: (1) maximize the probability of success for each mission, (2) minimize the vulnerability incurred by friendly assets from hostile threats, (3) minimize disruption to the original ATO (i.e., the total number of retasked assets), and (4) minimize the overall impact to the commander's intent (e.g., abandon low priority targets over high priority targets).

A candidate retasking solution given this problem formulation consists of an allocation of available assets at the injection time of a pop-up target to available targets (see Figure 2). A pre-filtering process is performed to determine feasible taskings for each target. For example, an asset that has already expended its ordnance cannot be tasked to strike another target. Additionally, an asset that cannot reach the target before the end of the target's window of opportunity cannot be tasked.

Asset	ALPHA	BRAVO	…	ZULU
Task	TGT1	TGT2	…	TST

Figure 2: Dynamic Targeting Mission Plan Genetic Representation

Using this direct encoding, standard genetic variation operators of 2-point crossover and uniform mutation are used. Invalid solutions may occur where two assets are assigned to a single target. In this case, a repair operation randomly selects one of those allocations to discard and substitutes it with an available allocation. If no allocation is available, the asset is not assigned to any target.

This approach has shown promise for the dynamic targeting domain. It can also be extended for supporting the Combat Plans Division within the AOC for analyzing the effects of potential emerging targets during ATO development. It is currently being used to support modeling and simulation-based experimentation.

3.2 ISR Asset Mission Planning

While dynamic targeting and deliberate targeting are a vital part of air power, it certainly does not operate in isolation. The ISR division within the AOC is responsible for providing accurate, relevant, and timely intelligence to decision-makers, increasing situational awareness and predictive intelligence. Similar to targeting, ISR planners must develop collection plans, tasking assets with specific platform and sensor capabilities to meet intelligence requirements. Extending from our work in the dynamic targeting domain, we applied our framework to assist ISR planners working on the ISR Operations Team in the AOC.

Under the targeting domain, objectives were focused in successful prosecution of targets and minimizing impact and disruption to the original plan. In the ISR domain, our objectives shift. A subset of those included in our application were: (1) minimize the total distance travelled by each ISR asset, (2) maximize the obedience to time window requirements, and (3) maximize the coverage of intelligence collected given requirements and available assets.

A candidate collection plan consists of an allocation of available ISR assets to an ordered collection of observable sites. Each ISR asset available is outfitted with a set of sensors, which enable it to gather intelligence collection types (e.g., IMINT, RADINT, COMINT). Each observable site has a set of related intelligence requirements, matched to an intelligence collection type.

Our initial approach to representing a collection plan was to use a multi-tiered genetic representation. In this representation, an individual is made up of one or more "asset-route" genes, where each asset-route gene contains a collection of waypoints for observable sites for an associated ISR asset. An example of this representation is illustrated in Figure 3. Lowercase letters are used as unique identifiers for observable sites. "D" refers to the depot, the initial location of the ISR asset. It is used to emulate refueling and has the added effect of producing multiple routes. "N" is a no-operation designator to emulate a variable length genetic representation.

ISR Asset	wp1	wp2	wp3	wp4	wp5	wp6	wp7	wp8	wp9	wp10	wp11	wp12	wp13	wp14	wp15
Predator	D	D	D	D	D	D	c	a	b	e	f	d	D	D	D
Global Hawk	c	a	b	e	f	d	D	D	D	D	D	D	D	D	D
U2	D	D	h	i	j	o	n	m	l	k	D	D	D	D	D
Rivet Joint	a	c	b	d	e	f	D	D	D	D	D	D	D	D	D
JSTARS	b	c	a	f	e	d	D	D	D	D	D	D	D	D	D

Figure 3: ISR Mission Plan Genetic Representation

Genetic variation operators function on both the individual level and on the gene level (see Figure 4). Crossover is performed using single-point crossover, selecting a subset of full asset-route genes from one parent and combining it with the complement from the other parent. At the gene level, an additional single-point crossover is performed per ISR asset between two candidate plans, swapping the respective parts of partial asset routes. Uniform mutation is then performed, selecting a viable collection site and replacing the current waypoint.

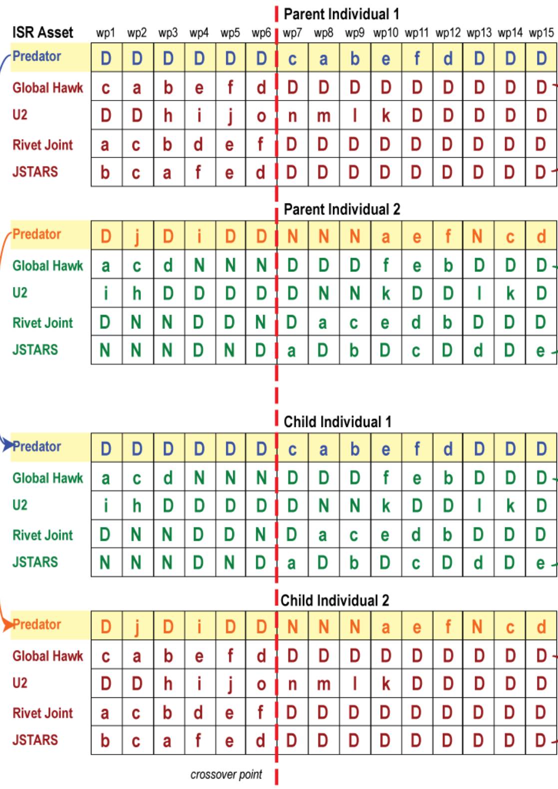

Figure 4: ISR Mission Plan Multi-tiered Genetic Variation

During our experimentation using this genetic representation, we recognized that our approach was similar to that of cooperative coevolutionary algorithms (CCEAs) [14]. In CCEAs, each population contains individuals representing a component of a larger solution. Evolution of these populations occurs in parallel, interacting only to obtain fitness.

Our most recent research has included revising our initial approach to leverage CCEAs. In this approach, a subpopulation is formed for each available ISR asset, the individuals of which constitute an individual collection plan for that asset. A complete collection plan is then formed by selecting a single member from each subpopulation. Crossover and mutation is performed on the subpopulation level as in a standard evolutionary algorithm. For evaluation, a member of a subpopulation collaborates with the best of the previous generation from each other subpopulation to form a complete candidate solution. This individual is then assigned the fitness of the complete solution. (see Figure 5). We are currently exploring the performance of this approach in comparison to our initial approach of using a multi-tiered representation.

3.3 UAS Mission Planning

While manned ISR assets have received much attention in years past, UAS have more recently come into the limelight for providing ISR. The number and variety of these platforms continues to increase. Currently, the generation of routes for UAS assets is human-mediated. In addition, the use of multiple unmanned aerial vehicles (UAVs) increases the complexity for a human decision-maker. To help reduce this burden, we have extended our work from the ISR asset mission planning domain to that of constellations of UAVs.

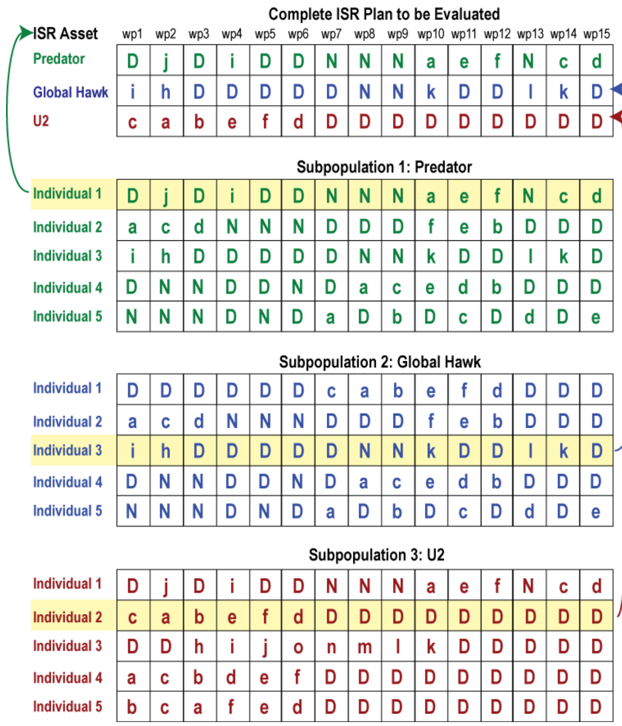

Figure 5: Cooperative Coevolution Approach

Using a genetic representation similar to that for ISR asset mission planning, we use the concept of an Air Maneuver Network (AMN), developed by the Battlespace Terrain Reasoning and Awareness (BTRA) program within the Army's Engineering Research and Development Center (ERDC), as a topological representation of a target area of interest (AOI) in place of direct observable sites. An AMN (see Figure 6) is a planar network, the nodes of which represent air control points and ground control stations connected by edges specifying bi-directional transitions between nodes. These edges represent flight segments with battlespace information (e.g., military value, weather, sensor performance, etc.) derived from BTRA spatial products [19].

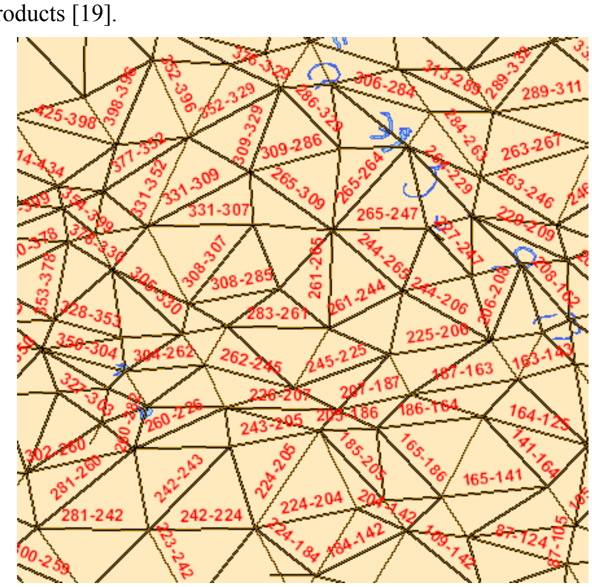

Figure 6: Air Maneuver Network Representation

For our effort to provide decision aids for UAS mission planning, we optimized for several objectives, a subset of which included: (1) minimize the latency between visits, (2) maximize the total coverage of the AOI, (3) maximize the time utilization, (4) minimize conflicts between multiple UAVs, and (5) maximize the cumulative military value.

Preliminary experiments using this approach show promise, and as the number and variety of these assets grow, we envision that automated mission planning tools will become increasingly important. We are currently exploring methods to manage heterogeneous constellations and are moving towards ways to incorporate the dynamic nature of the battlespace into the optimization process.

4. FUTURE DIRECTIONS

In the process of developing these decision-support tools to support air operations mission planning, we have come across a number of desirable features, many of which strike to the core of evolutionary computation research. These include incorporating user preferences (i.e., commander's intent) into evolutionary multi-objective optimization, constraint-handling, leveraging the knowledge of human experts in search, supporting dynamic replanning, and developing applications that can be used in real-time.

4.1 Incorporating User Preferences

Approximation of the Pareto front is a desirable goal in multi-objective optimization, attempting to provide a decision-maker with the tradeoff space among several objectives to make an ultimate selection based on desirable preferences. In COA analysis, this is desirable as well, as distinguishability will be a filtering criterion in evaluating a COA. Unfortunately, however, we have found that decision-makers have difficulty in accepting a solution set with multiple tradeoffs. Often, decision-makers want to weigh the importance of specific objectives over others. However, this leads the decision-maker to have to derive a complex equation essentially comparing apples and oranges (e.g., 3% decline in average probability of success is worth a 1% decrease in incurred vulnerability). This approach, using linear aggregation, has often shown to prematurely converge, bias the direction of search, and fail to locate optima in nonconvex portions of the solution space. Pareto dominance-based approaches in MOEAs (such as NSGA-II in our framework), however, treat each objective independently, and do not provide the end user with a means to give preference to certain objectives. What is needed is a compromise approach that approximates the Pareto front and incorporates user preferences into the optimization process.

Handling preferences in evolutionary multi-objective optimization has been surveyed by Coello Coello in [3] where it is described that preferences can be expressed a priori, a posteriori, or in an interactive manner. Linear aggregation is one method to define preferences in advance. Another method may be to modify the dominance definition, relaxing it from strict Pareto-dominance to one based on the preferences of the decision-maker. A posteriori methods enable a decision-maker to choose after search takes place. In our traditional approach, using MOEAs, the spread of solutions approximating the Pareto front is down-selected based on user preferences. Finally, we can investigate using interactive genetic algorithms as a method to bring the decision-maker into the process as the optimization process is performing. In this situation, the algorithm may query the decision-maker for comparisons between multiple solutions, allowing the human to guide the search.

4.2 Handling Constraints

As identified above, air operations mission planning requires obedience to hard constraints as imposed by the environment, geography, temporal conditions, and asset capability. It is insufficient for an automated decision-support tool to return an infeasible plan as a candidate solution. A method to handle these constraints properly and efficiently is needed.

Methods for handling constraints in EAs has been explored in detail over the last decade and best summarized in [6] and further in [4]. Traditionally, the use of penalty functions is used to decrease the overall fitness of an infeasible solution. Unfortunately, however, this can introduce a multitude of new parameters, produce premature convergence, or in some instances converge to infeasible solutions. Another method is to maintain feasibility through the use of special representations and genetic operators. This is the approach we take in applications to air operations mission planning, either by creating a closed system using specialized variation operators or by adding a repair operation to change an infeasible solution to a close feasible solution. Unfortunately, this takes away from the generalizability of EAs and can also lead to biasing the search, potentially limiting proper exploration of the search space. Another method is to treat constraint violation as additional objectives in a MOEA to be minimized. This approach, however, can easily lead to devoting precious computation time searching infeasible regions of the search space.

4.3 Leveraging Expertise

A decision-support tool is just that, a tool to support decision-making. As a result, it is desirable to somehow incorporate the knowledge of a human expert into the optimization process. Extracting and encoding this knowledge is an open area of research within the field of evolutionary computation.

One simple method to include expert knowledge into the optimization process is to seed the initial population with good solutions in addition to solutions to provide genetic diversity. Unfortunately, this requires that potential solutions to the specific or a similar scenario have been previously defined, which is rarely the case. Another approach is augmenting the evolutionary algorithm with a case-based memory of past problem solving attempts, periodically injecting appropriate intermediate solutions to the population [12]. This requires, however, the ability to identify similarity between problems and solutions. Another approach is to use interactive evolutionary algorithms, allowing the decision-maker to act as the evaluation function for candidate solutions. Unfortunately, this negates the key advantage of using an automated approach.

4.4 Dynamic Replanning

Air operations mission planning is situated within a dynamic environment where changes to the battlespace occur rapidly. It does little good to a decision-maker if the results returned by a decision-support tool are already obsolete. As the desire of the air operations community is to reduce the kill-chain cycle time, planning becomes a continuous process. While evolutionary algorithms have seen success in solving static optimization problems, more research is needed to explore how they can be used for dynamic optimization problems.

A simple method to provide the ability to dynamically replan is to re-start the optimization process when a change occurs in the battlespace. This is the approach we take with our decision-support tool for dynamic targeting. A monitoring process observes the current battlespace for significant changes, requesting an optimization when reality diverges from the original plan. While this is satisfactory for our purposes, if significant changes occur often or the time to run the algorithm is long, this is insufficient.

The field of evolutionary approaches to dynamic optimization problems provides insight to other methods to support dynamic replanning. A survey of these approaches is provided in [1]. In these cases, the evolutionary algorithm attempts to satisfy one or more of the following goals: (1) finding a tradeoff between convergence to the Pareto front and the cost of adaptation, (2) discovering solutions that are robust to changes in the environment, and (3) discovering solutions that are flexible. One method to do this is simply increasing diversity (e.g., increasing probability of mutation) when a change is detected. Another is to maintain diversity in the population, avoiding convergence altogether to keep robust solutions. Memory-based approaches can also be used, similar to using case-based reasoning to keep partial solutions that can be used later. Finally, using a multi-population approach can also explored, having a small population to scout and maintain promising areas of the search space.

4.5 Supporting Real-Time Applications

As these decision-support tools are needed to assist mission planners in an operational environment, their utility is dependent upon their ability to return a set of candidate solutions as quickly as allowable. Since evolutionary algorithms require a substantial number of evaluations, this is not always possible given our current capabilities. Methods to reduce the computation time of the optimization process without sacrificing usefulness could assist in helping provide applications that can operate in real-time.

Our preferred method for reducing computation time is to use low-fidelity simulation as a form of abstract wargaming. Capturing only the key variables necessary to provide a coarse-grained analysis aids in rapidly exploring the large search space of possible mission plans. Solutions returned from this process can then be analyzed in more fine-grained detail by a human decision-maker. When designing an abstract wargamer, however, one must be careful that the level of abstraction complies with the expectation of the user. As with any modeling approach, failure to capture the necessary variables can lead to solutions of little use.

Complementing low-fidelity simulation is the rise of modeling and simulation (M&S) within the defense community for training, experimentation, and analysis. Since evolutionary algorithms can take advantage of parallel processing, combining EAs, M&S, and grid computing can reduce the overall time required to optimize a collection of mission plans.

5. CONCLUSIONS

Air operations mission planning is a complex process, growing ever more complex given the increase in the number, variety, and interactivity of air assets. As we attempt to reduce the kill-chain cycle time, defense planners will need to rely on automated decision-support tools to help identify suitable mission plans. Our framework of combining multi-objective evolutionary algorithms with abstract wargaming has showed considerable promise, and we have applied it to several portions of the air operations domain, including dynamic targeting, ISR asset mission planning, and UAS route generation. However, there are significant areas where we can improve, coinciding with major areas of research within the evolutionary computation community. Leveraging the state of the art in these areas and contributing to their advancement can aid our mission planners dramatically.

6. ACKNOWLEDGEMENTS

The work presented here has been supported by Government Contract Nos. FA8750-04-C-0174, FA8750-07-C-0052, FA9550-07-C-0068, and W9132T-06-C-0040. The authors would like to thank Dawn Trevisani and Joseph Caroli of the Air Force Research Laboratory Information Directorate (AFRL/RI), John Tangney of the Air Force Office of Scientific Research (AFOSR), and Kirk McGraw of the U.S. Army Engineer Research and Development Center (ERDC) Construction Engineering Research Laboratory (CERL) for their assistance and vision.

7. REFERENCES

[1] Branke, J. 2002. Evolutionary Optimization in Dynamic Environments. Springer.

[2] Coello, C., Veldhuizen, D. V., & Lamont, G. 2002. Evolutionary Algorithms for Solving Multi-Objective Problems. New York: Kluwer.

[3] Coello, C. A. 2000. Handling Preferences in Evolutionary Multiobjective Optimization: A Survey. In Proceedings of Proc. of the 2000 Congress on Evolutionary Computation. Piscataway, NJ: IEEE Service Center. 30-37.

[4] Coello, C. A. 2002. Theoretical and Numerical Constraint Handling Techniques Used With Evolutionary Algorithms: A Survey of the State of the Art. Computer Methods in Applied Mechanics and Engineering, 191(11-12), 1245-1287.

[5] Coello, C. A. 1999. A Comprehensive Survey of Evolutionary-Based Multiobjective Optimization Techniques. Knowledge and Information Systems, 1(3), 269-308.

[6] Coello, C. A. & Carlos, A. 1999. A Survey of Constraint Handling Techniques Used With Evolutionary Algorithms. (Rep. No. Lania-RI-9904). Laboratorio Nacional de Informatica Avanzada.

[7] Deb, K., Agrawal, S., & Pratap, A. 2002. A Fast Elitist Non-Dominated Sorting Genetic Algorithm for Multi-Objective Optimization: NSGA-II. In Proceedings of Parallel Problem Solving From Nature VI Conference.

[8] Fonseca, C. M. & Fleming, P. J. 1993. Multiobjective Genetic Algorithms. In Proceedings of IEE Colloquium on Genetic Algorithms for Control Systems Engineering. London, UK: IEE.

[9] Knowles, J. D. & Corne, D. W. 1999. The Pareto Archived Evolution Strategy: A New Baseline Algorithm for Multiobjective Optimisation. In Proceedings of Congress on Evolutionary Computation. Washington, D.C..

[10] Lanchester, F. W. 1956. Mathematics in Warfare. In J. R. Newman (Ed.), The World of Mathematics (pp. 214-246). New York: Simon and Schuster.

[11] Langton, J. T., Caroli, J. A., & Rosenberg, B. 2008. An Evolutionary Algorithm Technique for Intelligence, Surveillance, and Reconnaissance Plan Optimization. In Proceedings of SPIE Defense and Security Symposium.

[12] Louis, S. J. & McDonnell, J. 2004. Learning With Case-Injected Genetic Algorithms. IEEE Transactions on Evolutionary Computation, 8(4), 316-328.

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

[14] Potter, M. A. & De Jong, K. A. 2000. Cooperative Coevolution: An Architecture for Evolving Coadapted Subcomponents. Evolutionary Computation, 8(1), 1-29.

[15] Rosenberg, B., Burge, J., & Gonsalves, P. 2005. Applying Evolutionary Multi-Objective Optimization to Mission Planning for Time-Sensitive Targets. In Proceedings of Genetic and Evolutionary Computation Conference (GECCO) 2005.

[16] Tenenbaum, S., Stouch, D., McGraw, K., & Fichtl, T. 2008. Multi-Objective Optimization to Support Mission Planning for Constellations of Unmanned Aerial Systems. In Proceedings of SPIE Defense and Security Symposium.

[17] United States Air Force. 1-22-2000. Air Force Doctrine Document 2-1: Air Warfare.

[18] United States Air Force. 6-8-2006. Air Force Doctrine Document 2-1.9: Targeting.

[19] Visone, D. L. 2005. Battlespace Terrain Reasoning and Awareness. In Proceedings of ESRI User Conference.

[20] Zitzler, E., Laumann, M., & Thiele, L. 2001. SPEA2: Improving the Strength Pareto Evolutionary Algorithm. (Rep. No. Technical Report 103). Zurich, Switzerland: Computer Engineering and Networks Laboratory (TIK), Swiss Federal Institute of Technology (ETH) Zurich.