[3]Yeh, H. H., Ridgely, B., and Banda, S. S., "Nonconservative Evaluation of Uniform Stability Margins of Multivariable Feedback Systems," *Journal of Guidance, Control, and Dynamics*, Vol. 8, No. 2, 1985, pp. 167–174.
[4]McRuer, D., Ashkenas, I., and Graham, D., *Aircraft Dynamics and Automatic Control*, Princeton Univ. Press, Princeton, NJ, 1973, p. 695.

Large-Scale Air Combat Tactics Optimization Using Genetic Algorithms

Sandeep Mulgund*
Arthur D. Little, Inc., Cambridge, Massachusetts 02140
Karen Harper†
Charles River Analytics, Inc., Cambridge, Massachusetts 02138
Kalmanje Krishnakumar‡
University of Alabama, Tuscaloosa, Alabama 35487
and
Greg Zacharias§
Charles River Analytics, Inc., Cambridge, Massachusetts 02138

I. Introduction

THE complete design and specification of air combat tactics for many-vs-many engagements poses a considerable challenge to tactical planners. Although solutions have been developed for one-vs-one or few-vs-few encounters,[1] the results may not generalize to larger engagements where formation tactics become increasingly important. Conventional methods such as trajectory optimization rely on gradient information that may not be available when the free parameters are discrete quantities such as the number of aircraft, weapons type, etc. Previous studies have examined the development of optimized air combat maneuvering via optimal control and machine learning methods.[2] However, the advent of high-speed, low-cost microprocessors has made feasible the use of genetic algorithms (GAs) for solving optimization problems where gradient information is unavailable. GAs have shown considerable potential in solving optimization problems that defy conventional approaches.[3]

In this Note an approach to large-scale air combat tactics optimization using genetic algorithms (ACTOGA) and the results obtained using the approach are presented. A prototype modeling tool has been developed, integrating a tactical engagement simulator with a GA "engine" for performance-based optimization of blue team tactics. The engagement simulator uses point-mass representations of aircraft dynamics, and it allows the simulation of forward-quarter beyond-visual-range (BVR) intercepts.[1] The goal of a BVR intercept is to maneuver one's aircraft to place a threat that is not visible to the naked eye (but visible to onboard sensors such as radar) within the aircraft's weapons envelope and to then launch medium- to long-range missiles against that threat. Only the initial launch of BVR missiles is optimized by the GA; the optimization of close-in "dogfighting" or defensive maneuvering is beyond the scope of this study. The outcome of an engagement (kills, losses, etc.) is used to guide the evolution of the optimized air combat tactics. A novel feature of the approach described here is the use of a "formation hierarchy" in which small, well-known conventional fighting units are aggregated to build large-scale tactical formations. This approach facilitates the design of tactics compatible with existing air combat principles. Excellent blue team performance is demonstrated where both sides are matched in terms of formation size and aircraft capabilities.

II. Air Combat Tactics Optimization

The term *air combat tactics* encompasses several concepts, including 1) the individual maneuvers that pilots use for accomplishing a given objective, 2) formation tactics that specify how small groups of aircraft can work cooperatively, and 3) principles for constructing division tactics that integrate large groups of aircraft.

When seeking to optimize air combat tactics, the objectives must be specified unambiguously. Optimizing individual maneuvers essentially amounts to optimal control design. However, optimal paths for a given situation do not provide any guidance on how to construct effective fighting groups for an arbitrary MvN engagement. Furthermore, there is no way to ensure that a pilot can execute a given optimal maneuver or remember n different optimal maneuvers for n different situations. As such, the approach taken here is to use conventional tactical maneuvers and optimize the manner in which they are used by a large aircraft formation.

The most effective formation tactics employ a basic fighting unit of two aircraft (called a *section* or *element*).[1] Because this is how all fighter pilots learn their craft, it was determined early in this research that it would be most effective for optimized tactics not to deviate from established tactical doctrine. Accordingly, software tactics modules that employ basic fighting units of two aircraft have been developed. Because large numbers of fighter aircraft are difficult to control, formation tactics for large groups can be developed using a hierarchical structure consisting of smaller units or divisions, for example, a four-airplane division called the *fluid four* consists of two elements. Each element consists of two aircraft, but they are treated as a unit. This hierarchical concept is used to develop a GA-based approach to optimized air combat tactics development. Given a palette of air combat maneuvers and standard small-formation tactics as building blocks, GAs are used to determine how they can be integrated to produce large fighting groups that optimize overall combat effectiveness.

Tactics implementation in ACTOGA proceeds as follows:

1) Define a set of commonly used element and division formations (Table 1) as well as the underlying tactical maneuvers and attack tactics.

2) Develop a set of principles for aggregating the small formation tactics for large MvN engagements and implement a method for doing so in the GA software. To illustrate, consider a team consisting of four aircraft. Using only the fighting wing and double attack, the possible team formations are shown in Fig. 1 (assuming both elements use the same two-ship formation). A similar approach can be used to develop large-division formations from smaller two-ship and four-ship groupings.

3) Use the resultant formation tactics to drive the engagement and evaluate the results via the performance metric generator.

4) Optimize the MvN engagement tactics with respect to the performance metrics.

Table 1 Commonly used formations[1]

Name	Symbol	Number of entities
Fighting wing	FW	2
Left double attack	DAL	2
Right double attack	DAR	2
Fighting wing	FW	2
Finger four	FF	4
Left sections in trail	LST	4
Right sections in trail	RST	4
Wall formation	WF	4

Received 9 March 2000; revision received 31 July 2000; accepted for publication 17 August 2000. Copyright © 2000 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved.

*Senior Manager, Technology and Innovation Business, 20 Acorn Park. Senior Member AIAA.
†Senior Scientist, 725 Concord Avenue. Member AIAA.
‡Associate Professor, Department of Aerospace Engineering and Mechanics. Associate Fellow AIAA.
§Principal Scientist. Member AIAA.

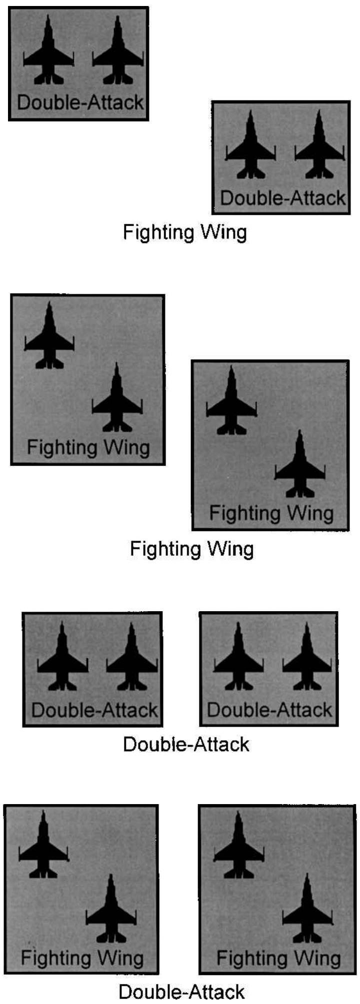

Fig. 1 Potential four-aircraft formations.

A cost function was developed to represent the weighted sum of several terms:

$$J = \sum_{i=1}^{6} k_i J_i, \qquad 0 \le k_i \le 1, \quad 0 \le J_i \le 1 \tag{1}$$

The GA's objective is to minimize J. Each of the k_i coefficients enables weighting the relative importance of each of the J_i terms, which are defined as follows: J_1 = fraction of blue members killed, J_2 = fraction of red team members surviving, J_3 = fraction of blue team members violating separation criteria, J_4 = mean blue team relative advantage assessment, J_5 = standard deviation in blue team relative advantage, and J_6 = mean blue risk assessment.

Considering first J_1 and J_2, a desirable solution would be one where none of the blue players are killed (i.e., $J_1 = 0$) and none of the red players survive (i.e., $J_2 = 0$). These are the basic performances desired from the optimization. However, by themselves, these considerations are insufficient to generate an effective tactical solution. The other terms address some key issues that characterize an effective tactical formation. The J_3 term captures the idea that team members should not come too close to each other because of the potential collision risk. Thus, a penalty is applied when any two aircraft violate a separation constraint (set to 2500 ft for the results that follow).

The J_4, J_5, and J_6 terms measure relative tactical advantage and situational risk, based on an assessment of the tactical situation. This assessment, in turn, is completed on the basis of belief network situation models.[4] The relative advantage belief network (not shown in the interest of brevity) quantifies a pilot's belief as to who possesses the instantaneous tactical advantage (blue, neutral, or red).

III. System Evaluation

To demonstrate GA performance, two different tactical scenarios were considered, each with a given red team formation.

Scenario 1a: 16v16; blue formation not optimized by GA

Scenario 1b: 16v16; red formation same as scenario 1; blue formation optimized by GA

Scenario 2: 12v16; examine blue performance when at a numerical disadvantage

A. Scenario 1a

The first scenario allowed the examination of the blue team performance in a simulation where both blue and red tactics were specified manually: the blue team's tactics were deliberately not optimized. Table 2 shows several performance measures for this scenario. Clearly, the blue team fared quite poorly. The relative advantage and risk measures refer to the belief network metrics described earlier. The mean relative advantage and risk measures are scaled between 0 and 1. Overall, the nonoptimized blue tactics provided a poor solution.

B. Scenario 1b

For the next scenario the red team formation and tactics used in scenario 1a are used. However, now the GA optimized the blue force tactics. This resulted in the configuration shown in Fig. 2. The GA selected a lead-trail (LT4) formation for the blue team, consisting of (from front to back) a finger four (FF) followed by three sections in trail (ST).

Table 2 shows the scenario's outcome. It is clear that the GA had a profound impact on the blue team's tactical effectiveness: blue suffered zero casualties, while red suffered 16 casualties. The blue team had better relative advantage and risk measures than the red team. The margin was determined to be statistically significant by the T-test. Furthermore, the standard deviation in both risk and relative advantage was zero, meaning that every blue player obtained the same performance measures. These results indicate that even with perfectly even blue/red aircraft performance, proper selection of tactical formation and intercept geometry can make a substantial difference.

C. Scenario 2

With a set of successful evaluations of ACTOGA's performance in 16v16 engagements, its performance was assessed in a scenario where the blue team was outnumbered. Figure 2 shows the 12v16 engagement used for scenario 2, where the red team was given the same configuration as in scenario 2. The GA fitness function weights were the same as those used in scenario 1b.

Table 2 shows the outcome of the second scenario, and it indicates that the GA once again obtained excellent results. Although there was no statistically significant difference in relative advantage and risk between the red and blue forces, the GA achieved a perfect result (in terms of kills and losses). Some blue players entered 1v2 engagements, which is unavoidable when the blue team is outnumbered.

Note the structural similarities in the blue team formation between scenarios 1b and 2. In both cases the root formation is a four-group lead trail. This suggests that it may be possible to generalize the specific results observed across a range of engagements (e.g., lead-trail formations can be effective against sections in trail).

Table 2 Outcome of scenarios

	Scenario 1a		Scenario 1b		Scenario 2	
Measure	Blue	Red	Blue	Red	Blue	Red
Casualties	9	7	0	16	0	16
Mean relative advantage (lower is better)	0.49	0.49	0.40	0.59	0.48	0.44
Standard deviation in relative advantage	0.03	0.03	0.00	0.03	0.12	0.12
Mean risk (higher is better)	0.44	0.36	0.85	0.66	0.85	0.71
Standard deviation in risk	0.28	0.28	0.00	0.19	0.01	0.17

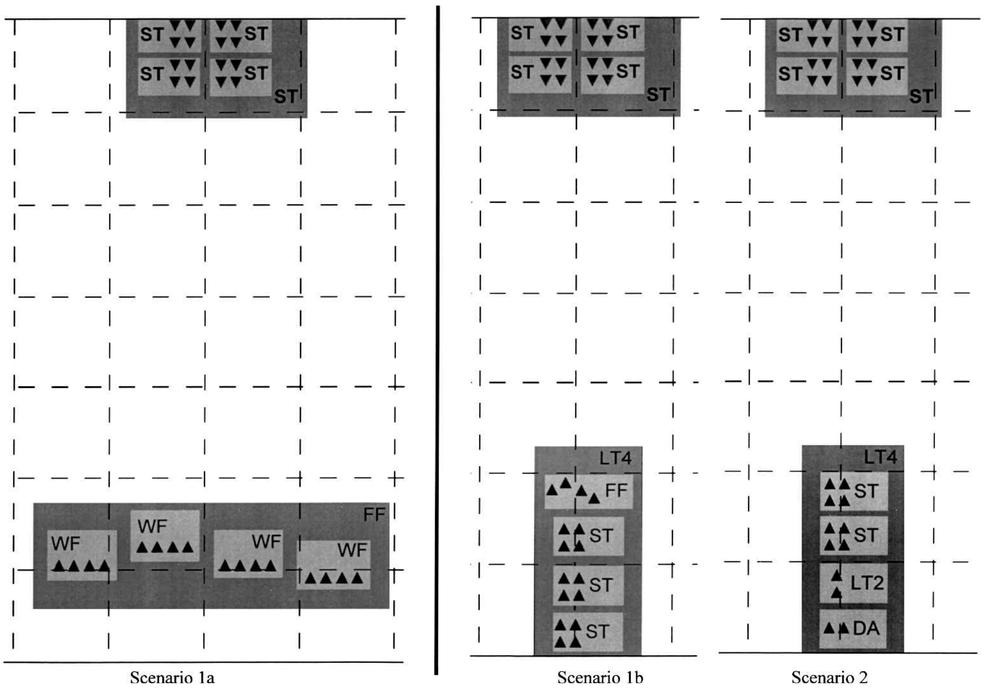

Fig. 2 Scenarios: 1a, no optimization, and 1b and 2, blue team optimized by GA (DA, double attack; LT2, LT4, lead-trail; FF, finger four; ST, sections-in-trail; WF, wall formation).

This could lead to rule extraction from GA solution databases to identify tactical rules of thumb for characterizing effective tactics.

IV. Conclusions

Air combat tactics become an ideal candidate for GA optimization when the problem is formulated in a manner conducive to genetic coding. The large numbers of elements (aircraft and groups) in many-aircraft combat formations and their interrelated behavior are defined effectively by the application of a hierarchical coding scheme. By altering the structure of the performance function to vary the significance of several tactical attributes, the GA-optimized solutions have been tailored to make them sensible from an air combat point of view. The GA provides solutions with generalized characteristics when faced with similar enemy formations.

Acknowledgments

This project was performed under U.S. Air Force Contract F33657-97-C-2037. The authors would like to thank the Technical Monitor, Tim Menke of Wright–Patterson Air Force Base, Ohio, for his continued support of this research.

References

[1]Shaw, R., *Fighter Combat: Tactics and Maneuvering*, Naval Inst. Press, Annapolis, MD, 1988, pp. 11–151.

[2]Smith, R. E., and Dike, B. A., "Learning Novel Fighter Combat Maneuver Rules via Genetic Algorithms," *International Journal of Expert Systems*, Vol. 8, No. 3, 1995, pp. 247–276.

[3]Krishnakumar, K., and Goldberg, D., "Genetic Algorithms in Control System Optimization," *Journal of Guidance, Control, and Dynamics*, Vol. 15, No. 3, 1992, pp. 735–740.

[4]Pearl, J., *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, San Mateo, CA, 1988, pp. 46–58.