JOURNAL OF SPACECRAFT AND ROCKETS
Vol. 37, No. 5, September–October 2000

Missile Aerodynamic Shape Optimization Using Genetic Algorithms

M. B. Anderson*
Sverdrup Technology, Inc., Eglin Air Force Base, Florida 32542
and
J. E. Burkhalter† and R. M. Jenkins‡
Auburn University, Alabama 36849-5338

The use of pareto genetic algorithms (GAs) to determine high-efficiency missile geometries is examined, and the capability of these algorithms to determine highly efficient and robust missile aerodynamic designs is demonstrated, given a variety of design goals and constraints. The design study presented documents both the learning capability of GAs and the power of such algorithms for multiobjective optimization. Results indicate that the GA is clearly capable of designing aerodynamic shapes that perform well in either single or multiple goal applications.

Nomenclature

A_{exposed}	=	exposed nozzle area relative to missile body frontal area, in.^{2}
A_{ref}	=	missile body frontal area, in.^{2}
b_t	=	exposed semispan of tail, in.
b_w	=	exposed semispan of wing, in.
C_A	=	axial force coefficient
C_N	=	normal force coefficient
C_{rt}	=	tail root chord, in.
C_{rw}	=	wing root chord, in.
L/D	=	lift/drag ratio
L_{body}	=	total missile body length, in.
L_{nose}	=	nose length, in.
L_{tot}	=	total body length, excluding nozzle, in.
R_{body}	=	body radius, in.
R_{exit}	=	exit radius, in.
TR_t	=	tail taper ratio
TR_w	=	wing taper ratio
X_{let}	=	distance from nose tip to tail leading edge, in.
X_{lew}	=	distance from nose tip to wing leading edge, in.
λ_{tet}	=	tail trailing-edge sweep angle, deg
λ_{tew}	=	wing trailing-edge sweep angle, deg

Introduction

In recent years, researchers have applied gradient-based optimization schemes to aerodynamic design,[1–4] but these methods are subject to undesirable restrictions. For instance, gradient-based optimizers must start with a specified set of initial parameters, which can bias future solutions toward a local optimum in the vicinity of the starting point. Gradient search procedures work most efficiently when there are a small number of design variables and when the variables are essentially independent of each other. As the number of design variables increases and coupling of the variables occurs (most often the case for complex aerodynamic designs), gradient-based algorithms do not have the ability to recombine disparate solutions to form solutions that sample a new portion of the optimization space.

Only recently have attempts been made to couple artificial intelligence, that is, learning techniques to aerodynamic design. The research of Gage and Kroo[5] focused on applying genetic algorithms (GAs) to the topological design of nonplanar wings. In their work the wing was broken into several segments that were allowed to vary in incidence and dihedral angles. The goal of the optimization was to minimize induced drag given a fixed lift. Their approach used both a penalty and repair approach to deal with solutions not achieving the fixed lift value. Bramlette and Cusic[6] have also applied genetic methods to the parametric design of aircraft, and Tong and Gregory[7] have used GAs to conduct the preliminary design of turbine engines after gradient methods had stalled in a local optimum. More recently, Anderson[8] has applied GAs to subsonic wing design with the goal of producing good aerodynamic shapes, with the additional constraint that the structure must not break. Anderson used penalty weights to combine lift, L/D, and structural design margins into one global objective function. Anderson, like Gage and Kroo,[5] pointed out that the achieved solution is strongly dependent on the values of the penalty weights, a very unappealing result. Later, Anderson removed the weighting procedure and instead used a pareto GA on the same problem.[9] The results of this later work showed that pareto GAs are ideally suited for complex optimization problems with diverse goals. Such an approach does mean, however, that the designer must scan the resulting solutions in the pareto optimal set to determine which solution or solutions are most desirable. Unlike single objective problems where there is a clear winner, multiobjective problems require judgement about which solutions are preferable. For cases where there are both aerodynamic and structural goals in the design, there is usually a willingness to trade off some aerodynamic performance to ensure that structural integrity goals are met. Some very recent work[10] shows that a GA can be used to design microdevices also. In this work a viscous micropump is designed to maximize the flow rate between two parallel plates encasing a rotating cylinder. The GA was used to determine the optimum plate spacing, the correct upper plate shape (only the lower plate was flat), and the optimal vertical location of the rotating cylinder to maximize the mass flow of the device at the exit plane of the two plates. Because the governing (Navier–Stokes) equations can be quite time consuming to solve, what should be called a microgenetic algorithm, in combination with a multipoint crossover (parameter-based) scheme and a high mutation rate, was employed in an attempt to reduce the total number of Navier–Stokes solutions attempted. The results indicated that the microgenetic algorithm with relatively few Navier–Stokes solutions (551 total) achieved good pump designs when compared to standard large population GAs. In Ref. 11, a parameter evolution strategy coupled with a full potential method code is used to design transonic wings in a parallel mode. By distributing individual cases to multiple processors, significant turn-around time reductions were achieved. The authors noted that taking advantage of the inherent parallelization capability of population-based search techniques was a strength of the approach. The parameter evolution

Presented as Paper 99-0261 at the AIAA 37th Aerospace Sciences Meeting, Reno, NV, 11–14 January 1999; received 1 May 1999; revision received 5 November 1999; accepted for publication 5 May 2000.

*Deputy Director, Development Test Department. Member AIAA.
†Professor, Aerospace Engineering. Associate Fellow AIAA.
‡Associate Professor, Aerospace Engineering. Senior Member AIAA.

strategy is similar to simulated annealing and relies more heavily on random parameter variations than does the survival of the fittest strategy embedded in a GA. Both Refs. 10 and 11 used procedures with essentially high mutation rates when compared to the standard GA. For small populations, however, there continues to be emerging evidence that higher than normal mutation rates are justified. Another emerging trend in optimization using GAs is to couple the GA with more traditional design techniques to speed convergence. The goal of these hybrid methods is to use the global search strength of the GA to provide useful information (starting points) for gradient ascent/descent methods. Reference 12 successfully coupled a GA with a differential corrections method to find flight condition estimates (i.e., Mach number, angle of attack, aerodynamic roll angle, freestream pressure) from static pressures measured around a missile forebody. The GA provided the initial guess for the gradient descent method. Fast and accurate flight condition estimates were consistently obtained with this method. In Ref. 13, a GA was coupled with an inverse design method to design wind turbines (wind mills) that maximize power output at varying wing speeds. The design variables included blade pitch, blade chord, and blade twist distributions with span. In this case, the GA executed the design search and the inverse procedure enforced certain constraints while giving the designer flexibility in choosing which variables to iterate and which to send to the GA. Another hybrid approach[14] coupled a GA with a standard gradient approach to maximize lift coefficient for a multielement airfoil. The design variables included flap angle, gap, overflap, and flap surface geometry. At some prescribed threshold lift coefficient, the GA was halted and the gradient method was begun. This approach is very similar to the method of Anderson et al.[12]; however, note that both of these works determined that selecting the proper threshold for switching algorithms is critical to the final result.

Aerodynamic Code

The aerodynamic prediction methodology used as the basis for this study, AeroDesign, was developed at the U.S. Army Aviation and Missile Command as a design tool for preliminary investigations of cruciform missile configurations. AeroDesign is an empirically based package that was developed using British wind-tunnel test data for a wide variety of configurations, including either a cone or ogive nose, a cylindrical body, up to two sets of four fins (in a + configuration only), and a boat tail. The Mach number regime of the code is from subsonic to high supersonic/low hypersonic, that is, Mach 4.0, flow.

AeroDesign was modified to include two axial force considerations that were not part of the original software. First, the fineness ratio of the nose of the missile is compared to a Sears–Haack (see Ref. 15) body. If the nose of the resulting design is more blunt than the Sears–Haack body, the axial force coefficient is scaled by a ratio of the fineness ratios of the two bodies as

$$\Delta C_{A_{\text{Sears–Haack}}} = \frac{C_{A_{\text{Sears–Haack}}}}{(L_{\text{nose}}/2R_{\text{body}})} \tag{1}$$

where

$$C_{A_{\text{Sears–Haack}}} = \frac{9.67}{(L_{\text{body}}/2R_{\text{body}})^2} \tag{2}$$

As Eq. (1) shows, larger nose fineness ratios have a smaller Sears–Haack axial force penalty. It is recognized that this correction is at best only first order, but the limited axial force prediction capability of the original AeroDesign code necessitates incorporation of such improvements.

The second axial force coefficient correction added to AeroDesign was implemented to correct for cases where the rocket nozzle exit diameter exceeded the diameter of the body. The baseline AeroDesign code has no nozzle aerodynamic prediction capability whatsoever, which in itself is not too bad for most preliminary design studies, particularly when all of the propellant has burned and the plume cannot affect the aerodynamics. Nevertheless, an axial force coefficient correction was added to include a drag increment based on the exposed area of the nozzle that exceeds the body frontal area. Early studies with AeroDesign indicated that without nozzle axial force considerations, the resulting nozzle exit areas could be quite large without an appreciable impact to the overall missile axial force coefficient. The correction is as follows:

$$\Delta C_{A_{\text{nz}}} = C_{A_{\text{flat plate}}}(A_{\text{exposed}}/A_{\text{ref}}) \tag{3}$$

where

$$A_{\text{exposed}} = \pi\left(R^2_{\text{exit}} - R^2_{\text{body}}\right) \tag{4}$$

$$C_{A_{\text{flat plate}}} = 1.0 \tag{5}$$

$$A_{\text{ref}} = \pi R^2_{\text{body}} \tag{6}$$

This correction is used only if R_{exit} is greater than R_{body}.

Aerodynamic Code GA Link

The links to the GA are straightforward. The genetic software used for this study was the IMPROVE code,[16] a general-purpose GA package that can easily be interfaced to a variety of objective functions. The GA is the controlling routine that calls the aerodynamic performance code as needed. The GA passes down the design parameters through the subroutine call statement and the aerodynamics code passes back a measure of how well the design performed in the different goal areas.

Variables Governing Design

The design variables passed by the GA to the aerodynamics code are shown in Fig. 1 and defined in the Nomenclature. For this study, the distance over which the nozzle contracts and expands was assumed to be a constant, proportional (three-to-one) to the radius of the exit. Each variable is constrained with a maximum and a minimum value. The desired resolution of each parameter is also specified. The maximum, minimum, and resolution dictate the size of the optimization space. In GA terms, the number of genes in each chromosome (also known as the number of bits in a base-two system of ones and zeros) is defined as

number of genes

$$= \sum_{n=1}^{\text{number of parameters}} \left\{ \text{Integer}\left[\frac{\text{LN}(\max_n - \min_n)/\text{resolution}_n}{\text{LN}(2)} \right] + 1 \right\} \tag{7}$$

Because it is natural in a design process to specify some desired payload size that must be contained within the external shape, a 16-in.-radius × 300-in.-length cylinder was selected as the space requirement. The payload could, for example, encompass a rocket motor and electronics. Table 1 shows the minimum, maximum, and

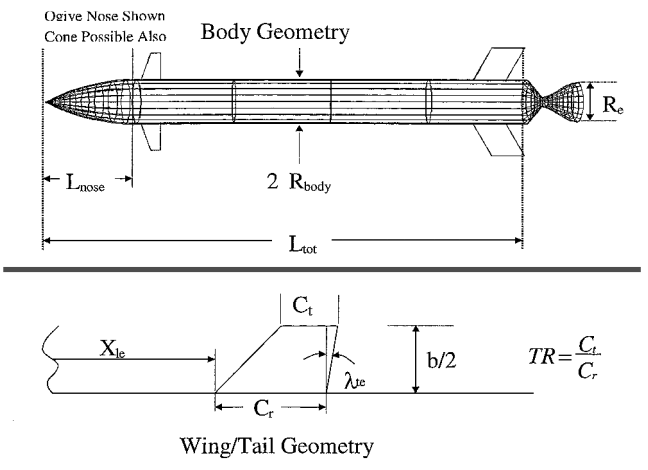

Fig. 1 Basic external aerodynamic shape variables.

Table 1 Maximum, minimum, and resolution of variables governing external aerodynamic design

Parameter	Minimum	Maximum	Resolution	Number of genes
Nose	0	1	1	1
L_{nose}	20	200	5	6
L_{tot}	300	700	5	7
R_{body}	16	20	0.25	5
R_{exit}	4	30	0.25	10
X_{lew}	20	350	2.0	8
b_w	0	80.0	1.0	7
C_{rw}	0	80.0	1.0	7
λ_{tew}	0.0	75.0	5.0	5
TR_w	0	1.0	0.05	4
X_{let}	350	680	2.0	8
b_t	0	80.0	1.0	7
C_{rt}	0	80.0	1.0	7
λ_{tet}	0.0	75.0	5.0	5
TR_t	0	1.0	0.05	4

resolution of each parameter defining the external shape of the missile carrying this payload. The nose parameter denotes either an ogive or a cone.

These parameter definitions mean that 91 genes (or bits) are required to capture the design space. The number of possible designs is, therefore, 2^{91}, or roughly 2500 trillion trillion. The size of the design space immediately leads to doubt about the success of a procedure that does not even start with a reasonable guess at a good design.

Performance Measurement Conditions

Each performance measure was sampled over eight Mach numbers at eight angles of attack to form a sufficient aerodynamic data base. The Mach numbers were 0.3, 0.7, 0.9, 1.1, 1.5, 2.0, 3.0, and 4.0, and the angles of attack were 1.0, 2.0, 4.0, 6.0, 8.0, 12.0, 16.0, and 20.0 deg. The performance at each unique flight condition i contributed to the overall performance of the design through a simple summing procedure. For example, if high L/D is the current performance parameter, the performance measure is calculated from

$$\left(\frac{L}{D}\right)_{\text{total}} = \sum_{i=1}^{\text{no. of conditions}} \left(\frac{L_i}{D_i}\right) \tag{8}$$

The Reynolds number was set to 6×10^6 for each Mach number rather than setting an established pressure altitude.

GA Setup and Mode of Operation

Many variants on the basic GA have been developed to help speed convergence and minimize computational expense. For this study, the two variant algorithms used are creep and elitism. Creep is designed to randomly alter a certain percentage of the parameters by their resolution amount. This technique helps the basic GA to fine tune solutions when an optimum is found. Elitism forces the GA to keep intact the best performer of the current population into the next population. Thus, the best performer cannot be destroyed through crossover or mutation, and can only be replaced by a better performer when one is created. The genetic parameters used to control the probabilities within the algorithm are shown in Table 2.

The population size is roughly three times larger than the number of bits/genes required to represent the parameters. This rule of thumb has been successful for many previous studies but clearly would be very computationally expensive if a higher-order aerodynamic design methodology were used. Coupling computational fluid dynamics (CFD) to a GA for a complex three-dimensional shape would require significant computational resources, whereas with an empirical method like the one used here 250 solutions require only seconds on a Pentium PC.

Results: Goal-Based Designs

The designs shown in this section assume a rigid structure, so that structural considerations play no part in the design process.

Table 2 GA control parameters for external shape design

Parameter	Value
Crossover probability	0.9 (90% of the time 2 survivors are mated they produce offspring)
Mutation probability	0.002 (2 out of every 1000 bits will mutate)
Creep probability	0.05 (5% of the parameters in a population will creep by their resolution)
Population size	250

The cases investigated here incrementally add difficulty to the design problem as additional influences on the performance are activated.

Case 1: Maximize Normal Force Coefficient over Various Flight Conditions

Case 1 is designed to see whether the GA will maximize the fin areas. Larger fins will, of course, produce higher normal forces than smaller ones. Fin placement does not play a significant part in the design consideration for this case but would certainly be a driver for stability and control (S&C) when S&C becomes a design consideration. Also, because axial force is not considered for this case, the nozzle geometry (exit radius) can essentially be anything and the forebody (nose) can be very blunt without detrimental lift impact. There is significant coupling between the design variables, and this is an important point when contrasting a GA with a gradient optimization scheme. As noted earlier, gradient-based methods have difficulty with variable coupling. This problem is best seen by example for the problem at hand. To increase normal force, an algorithm can do two things for a missile-type shape: increase the wing area or increase the body size. For a gradient method $\text{d}(C_N)/\text{d}(\text{body radius})$ and $\text{d}(C_N)/\text{d}(\text{wing area})$ are always positive, and so a gradient method would be likely to maximize both the body radius and the wing area. That is, even though sensitive derivatives for these terms can be properly established, a standard gradient method would tend to increase both the body radius and the wing area and not trade one for the other. Therefore, a second derivative term, $\text{d}^2(C_N)/\text{d}(\text{body radius})\text{d}(\text{wing area})$, would be required to solve the problem. The problem is that for a fixed maximum span increasing the body size decreases the available wing area. Wings are more efficient at producing normal force than the cylindrical body, so that every increase in the body radius effectively decreases the amount of normal force that can ultimately be generated. This subtle coupling of variables makes gradient methods unattractive for complex design problems because there is no practical way of determining a priori what sensitivity derivatives are needed. Thus, one of the challenges posed to the GA is to produce high normal forces smartly, by effectively maximizing wing surface area through the minimization of the body radius.

Convergence history is important to the understanding of the behavior of a GA. Figure 2 shows the convergence history for this case. The maximum line denotes the performance of the best single member of the population. Average denotes the performance of the average member of the population. This type of convergence behavior is characteristic of GAs with creep: rapid assimilation of good characteristics produces rapid improvements, but then an asymptotic convergence region is reached as the GA attempts to fine tune the solution through creep. Without creep, there are typically fewer changes in the design and the improvements are made in large leaps. Creep certainly does not prevent the large leaps in performance as Fig. 2 shows, but there is admittedly a higher chance that the solution will fall into a local optimum.

Figure 3 graphically shows the design history for the missile. In Fig. 3, and subsequent figures of this nature, only the design that produced the highest aggregate normal force coefficient at each displayed generation is shown. Because GAs are population based, many designs (249 to be exact) within these same generations are not shown. While viewing these quantitative snapshots of the design process, recall that for the high normal force coefficient case neither the nose shape, the nozzle shape, nor wing/tail placement drastically alter the normal force coefficient. Although these particular design traits may make the missile look a little strange at first, one must

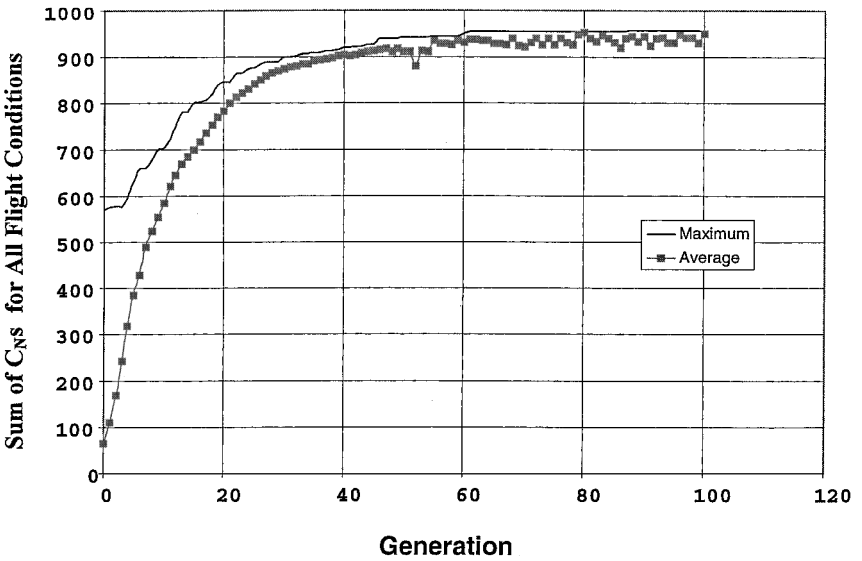

Fig. 2 Convergence history for the maximum normal force coefficient case.

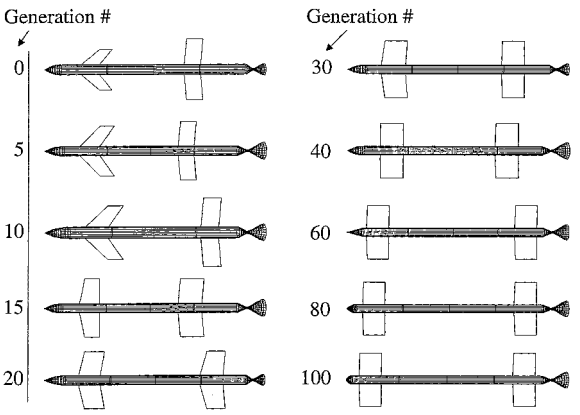

Fig. 3 External shape history for the maximum normal force coefficient case.

Table 3 Final parameter values: case 1

Parameter	Minimum	Maximum	Resolution	Final value
Nose	0	1	1	1
L_{nose}	20	200	5	20
L_{tot}	300	700	5	700
R_{body}	16	20	0.25	16
R_{exit}	4	30	0.25	29.19
X_{lew}	20	350	2.0	44.88
b_w	0	80.0	1.0	80.0
C_{rw}	0	80.0	1.0	80.0
λ_{tew}	0.0	75.0	5.0	0.0
TR_w	0	1.0	0.05	0.967
X_{let}	350	680	2.0	595.88
b_t	0	80.0	1.0	80.0
C_{rt}	0	80.0	1.0	80.0
λ_{tet}	0.0	75.0	5.0	0.0
TR_t	0	1.0	0.05	0.967

focus on the wing/tail planform geometries because they are the most important for now. From these qualitative representations of the design variable histories it is readily apparent that the GA learned to make the wing and tail areas large and to minimize the body radius during the design process. Although early attempts used highly swept wings, they were later replaced by blunt and large wings. The missile body also grew in length during the process, which contributes to higher normal force coefficients. The nose length was also minimized in favor of longer body length, further helping to maximize the normal force coefficient. As these final generations show, there is virtually no root-to-tip taper ratio in the wings or tails. In fact, the final taper ratio is exactly one resolution value away from having no taper at all in the wings/tails. This would appear to show that the GA failed to find the correct taper ratio because logically it would be expected that there should be no taper when maximizing the normal force coefficient. In reality, however, an investigation on the effect of taper on the normal force coefficient showed that, although for subsonic Mach numbers no taper produced better normal force performance, for supersonic Mach numbers slight tapering improved the normal force coefficient. As the taper ratio was manually decreased an additional resolution amount, the aggregate normal force coefficient decreased from what the GA found, meaning that the GA had found an optimum for taper ratio.

According to Fig. 3, there must be a benefit to spatially spreading the wings and fins apart from one another. Numerical studies, in fact, showed just that. The wings and tails were moved toward each other by their resolution amounts and the aggregate normal force coefficient decreased. Moving the tail aft also decreased the normal force coefficient; however, moving the wing forward did marginally increase the aggregate normal force coefficient (0.78% maximum change in C_N). Thus, it would appear that it is possible to numerically prove for this case that the GA did not find the exact optimum in the 100 generations it was allowed to run. However, note that of all of the design variables, only the variable controlling the forward movement of the wing increased the aggregate normal force coefficient (and then only by 0.78%). Table 3 shows the final values for the best design for this case. The nose length was minimized, the body radius was minimized, the body length was maximized, the wing and tail root chords and spans were maximized, and the wing and tail sweep angles were minimized.

Variables such as the nozzle exit radius and the type of nose are not significant until drag is considered.

Case 2: High Normal Force, Low Axial Force

The objective function for this case is simply an accumulation of C_N / C_A over all of the flight conditions:

$$\left(\frac{C_N}{C_A}\right)_{\text{total}} = \sum_{i=1}^{\text{no. of conditions}} \left(\frac{C_{N_i}}{C_{A_i}}\right) \tag{9}$$

The inclusion of a requirement for low axial force coefficients should necessitate two primary design changes over the preceding high normal force coefficient case. First, the nose of the missile should elongate to minimize large wave drag due to a blunt nose [the Sears–Haack correction of Eqs. (1) and (2)]. Second, the nozzle exit radius should be equal to or less than the radius of the body [Eqs. (3–5)]. It might be expected that there would be wing/tail leading-edge

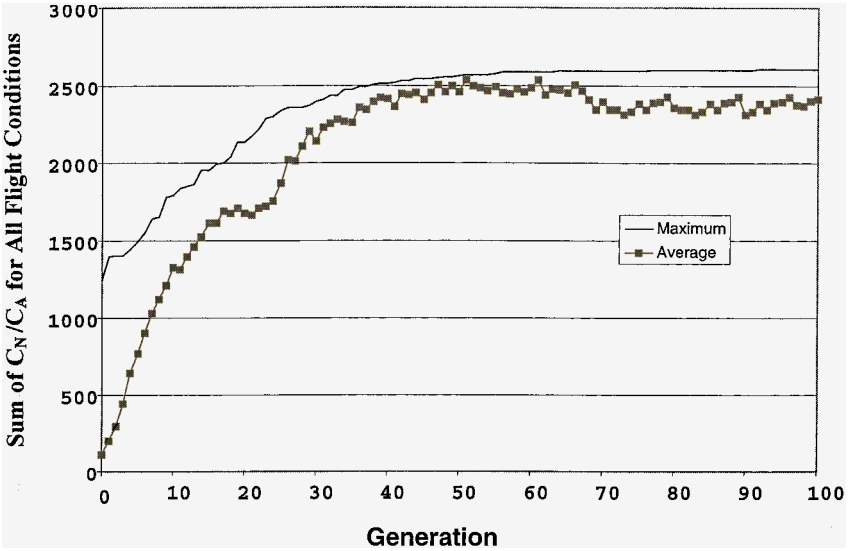

Fig. 4 Convergence history for the maximum C_N/C_A case.

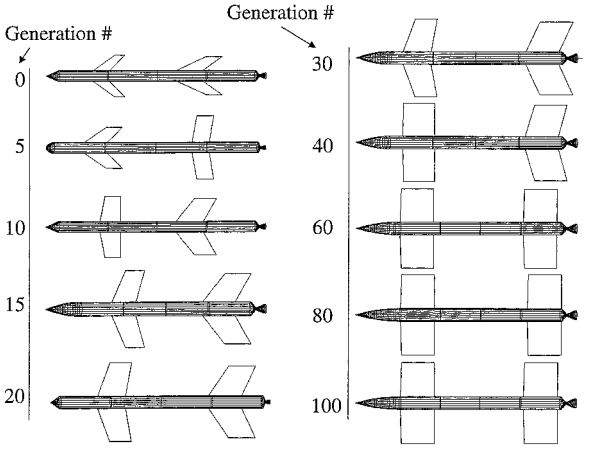

Fig. 5 External shape history for the maximum C_N/C_A case.

sweeps to forestall the sharp transonic drag rise that accompanies the transition from subsonic to supersonic flight. However, given the weighting toward supersonic Mach numbers in the preplanned flight conditions (five of eight Mach numbers are supersonic), the GA might find that high normal force coefficients due to unswept wings might outweigh the delay in axial force coefficient rise that sweep would certainly produce.

Figure 4 shows the convergence history for this case. As before, there is rapid assimilation of good characteristics (the power of GAs with good diversity in the population), but then an asymptotic convergence region is reached as the GA attempts to fine tune the solution through creep. After generation 70 (roughly), there is no continued improvement in the best member of the population, and the average member actually starts to decline (which looks a little odd). An investigation into the cause of the declining average stems from geometry conflicts that the GA produces as it begins to fine tune the solution. A geometry-conflicts subroutine checks to make sure that the wings are not overhanging the missile nose and that the tails are not overhanging the nozzle portion of the missile. Through time, as the GA began to maximize the wing and tail areas (Fig. 5), it moved these surfaces toward the extremes of the nose and tail to increase normal force coefficient as it did before. Because the nose of the missile began to elongate to reduce axial force coefficient, the forward placement of the wing eventually began to conflict with the length of the nose. Designs with conflicts return a zero value for the objective function. As more conflicts were encountered during the fine-tuning process, the average for the population did indeed decrease. Nevertheless, the GA design history (Fig. 5) shows that the algorithm produced good high C_N/C_A designs. As expected, 1) the wings and tails became exactly those that were designed under the high C_N case (taper ratio included), 2) the body radius was exactly minimized, 3) the nozzle exit radius became less than the body radius, and 4) the nose grew in length to minimize bluntness (to reduce the axial force coefficient). Unlike the preceding case, the wing leading-edge location was exactly placed where the nose ended, thereby attempting to maximize the possible separation distance between the wing and the tail. The tail location was within one resolution increment of being at the aftmost position possible. A test was performed where the tail was moved to its aftmost possible position and the aggregate C_N/C_A only changed by 1.5 (roughly 0.058% better than the design produced by the GA).

Case 3: High Normal Force, Low Axial Force, Minimal Static Margin

This case includes a measure of the stability of the airframe in addition to the aerodynamic performance. AeroDesign can calculate the static margin, which is the distance between the center of pressure and the center of gravity, once the location of the center of gravity is input to the code. The center of gravity is in turn calculated by the components that make up the missile. One of the most significant contributors to the weight and center of gravity is the rocket motor propellant. For this study, the propellant was not burned, so that the center of gravity remained fixed at one location. The fuel mass was fixed at 6500 lbm in a length of 300 in. and a radius of 15.7 in. A 1000-lbm payload was also included. The payload was assumed to begin at the end of the nose section (beginning of the body) and its volume and center of gravity location was calculated by assuming a payload density of 0.1 lbm/in.^{3} The outer body of the missile was assumed to be made of a titanium alloy (density 0.23 lbm/in.^{3}), as were the wings, tails, and nozzle. The rocket motor case was assumed to be a steel alloy with a density of 0.29 lbm/in.^{3}, an ultimate strength of 195,000 psi, and a design pressure of 3000 psi (which defines its thickness). Because the GA controls the external design of the missile, a mass-properties subroutine must be called for each attempted design. This subroutine calculates the total missile mass, center of gravity location, and inertial properties. After the mass properties for the design have been determined, the AeroDesign code can then calculate the normal force coefficient, axial force coefficient, center of pressure, and static margin for each Mach number and angle of attack.

Equations (1–5) are used again to measure the aerodynamic performance, but static margin is such a different design consideration that it is natural to question how a measure of this goal can be included into the design process. Certainly a weighting procedure could be used to combine the aerodynamic performance goal and the static margin goal into one overall performance measure, but,

just as Anderson[8] has shown, there is no clear way to know a priori the correct weighting values. For this study a strict pareto GA was employed to avoid weights. The strict pareto GA considers each design and compares how each design works in each goal area. Superior designs are those that perform better than their competitors in each goal area. Designs that are easily dominated disappear and, at the end of the design process, the final population is called the *p*-optimal set because it contains a mixture of designs that work reasonably well in each goal area. Certainly there will be a clearly superior performer in each goal area; however, the best multiobjective performer will probably not also be the best performer in a given specific goal. The *p*-optimal set is made up of hybrid designs that have survived based on their ability to perform across multiple goals. The aerodynamic performance goal is calculated as before, whereas the static margin goal, which is to minimize the static margin, is calculated as

$$\bar{X}_{\text{static}} = \sum_{i=1}^{\text{no. of conditions}} \left| X_{\text{cp}_i} - X_{\text{cg}_i} \right| \tag{10}$$

where the X locations are measured positive aft of the nose tip. This construction will not guarantee a stable missile, but will simply minimize either the stability or the instability of the missile. The pareto scheme will attempt to find solutions that will perform reasonably well in each of the two goal areas, but it can be expected that some of the earlier aerodynamic performance exhibited will be lost to the desire to have a minimal static margin. Note also that the addition of a new design goal further enhances the desirability of a nongradient-basedoptimization method. Gradient methods are difficult to correctly implement with multiple goals, and there is always the question of the frequency/method of switching between the different goals during the design process.

Figure 6 shows the convergence history for the average among the members of the population in each goal area. Figure 6 shows the classic behavior expected during a tradeoff-type optimization approach. The static margin and aerodynamic performance continued to improve for the average member, at least until generation 90–95 for this case. There is slight improvement beyond this point, but not significant enough improvement to warrant additional generations. Figure 7 shows the performance of each member of the *p*-optimal set in each goal area. As Fig. 7 shows, there are a great many solutions that perform well in each goal, and there are members of the population that are clearly superior in each goal. There is an abundance of solutions with accumulated static margins between 10 and 20 in. with an accumulated C_N / C_A between 2050 and 2100. There are superior aerodynamic performers around 2200–2250 C_N / C_A, but their static margins are between 100 and 200 in. Clearly, some aerodynamic performance was sacrificed for good static margin performance. The pareto approach produces an abundance of solutions with an accumulated static margin near 10 in., allowing the designer to choose from a variety of similar performing

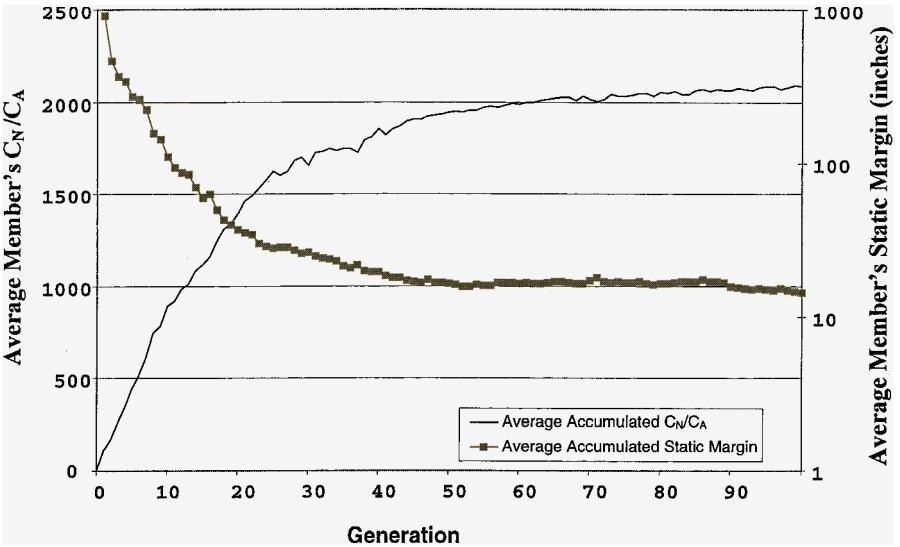

Fig. 6 Convergence history for the maximized C_N/C_A, minimized static margin case.

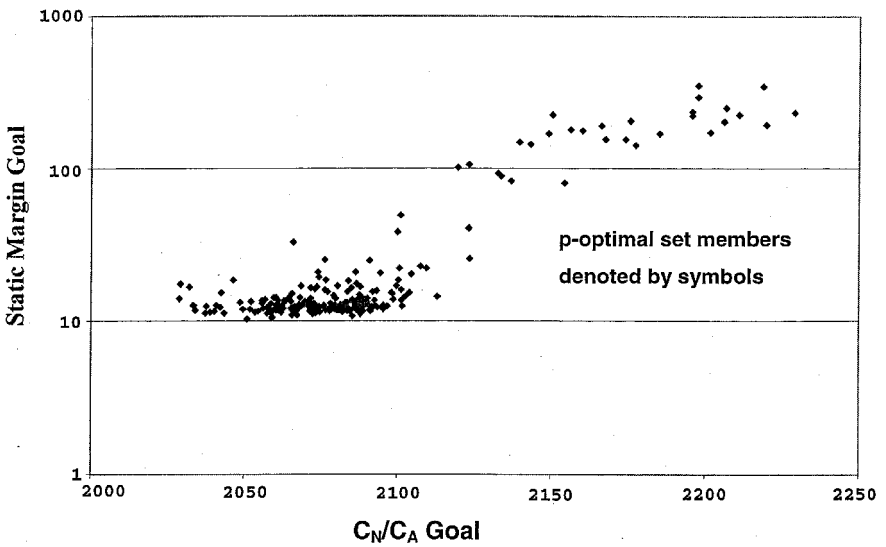

Fig. 7 Performance variation for *p*-optimal set.

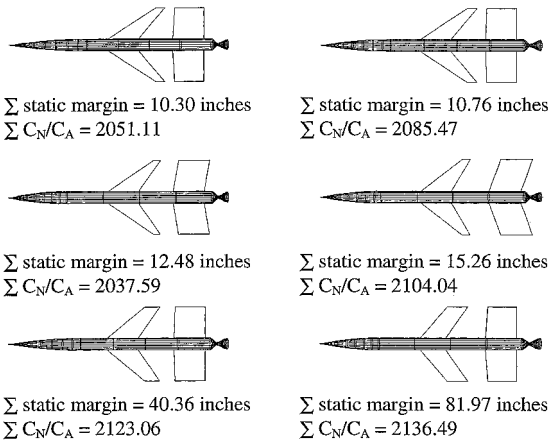

Fig. 8 Some available designs at generation 100: maximized C_N/C_A, minimized static margin case.

solutions, based perhaps on some other criteria such as manufacturing simplicity or weight. Several of the better solutions are shown in Fig. 8. These solutions have different levels of static margin and aerodynamic performance. All of the solutions have large tails to increase C_N/C_A. This result should have been expected because large tails would serve also to help the static margin performance. From these solutions, it is apparent that the better static margin solutions have more highly swept wings and more taper than the higher C_N/C_A solutions. More generally, note that each of these solutions has swept wings. Clearly, the GA learned to keep the wing areas fairly large while at the same time keeping the center of pressure far enough aft to keep the static margin reasonable. An increased wing area improved C_N/C_A performance; however, the static margin goal seemed to suffer at the same time.

Conclusions

The GA is clearly capable of designing aerodynamic shapes that perform well in either single or multiple goal applications. For multiple-goal/multiple-objective problems, the strict pareto algorithm provides hybrid designs that seek to perform well in each goal area, while providing significant diversity in the p-optimal set. The p-optimal set allows designers to choose among several candidate solutions.

It is also clear from these results that complex problems are not a deterrent to the use of GAs in solving a problem. Any of the models used in the present solutions could be replaced by more accurate routines because the interface between the GA and the various prediction codes involves only a simple, one-line subroutine call. For example, the aerodynamic prediction code used in the present analysis, AeroDesign, could be replaced by a more robust code such as Missile Datcom. The only restriction that is placed on these codes is that they must be capable of providing answers to any design that is developed by the GA code. For codes in which some designs are off limits, penalties must be artificially programmed into the code to reject those designs. Consequently, the final solution resulting from the GA design process is only as good as the codes involved. For the present case, no attempt has been made to quantify the accuracy of the final configuration because the goal was to develop the techniques and demonstrate the capability of the GA approach. This has clearly been done, but it is recognized that more accurate codes could improve the final results but not likely change the final general configurations very much.

Finally, it is clear that the methodology has been established in which GAs can be successfully used to circumvent months and perhaps years of labor in pursuing designs that work well individually but not collectively as a system. The use of GAs can optimize system approaches and will most certainly alter the manner in which system designs are finally extracted.

References

[1]Kroo, I., and Takia, M., "A Quasi-Procedure, Knowledge-Based System for Aircraft Design," AIAA Paper 88-4428, Sept. 1988.

[2]Kroo, I., Altus, S., Braun, R., Gage, P., and Sobieski, I., "Multidisciplinary Optimization Methods for Aircraft Preliminary Design," AIAA Paper 94-4325, Sept. 1994.

[3]Gage, P., and Kroo, I., "Development of the Quasi-Procedural Method for Use in Aircraft Configuration Optimization," AIAA Paper 92-4693, Sept. 1992.

[4]Consentino, G., and Holst, T., "Numerical Optimization Design of Advanced Transonic Wing Configurations," AIAA Paper 85-0424, Jan. 1985.

[5]Gage, P., and Kroo, I., "A Role of Genetic Algorithms in a Preliminary Design Environment," AIAA Paper 93-3933, Aug. 1993.

[6]Bramlette, M., and Cusic, R., "A Comparative Evaluation of Search Methods Applied to the Parametric Design of Aircraft," *Proceedings of the Third International Conference on Genetic Algorithms*, edited by J. D. Schaffer, Morgan Kaufmann, San Mateo, CA, 1989, pp. 213–218.

[7]Tong, S. S., and Gregory, B. A., "Turbine Preliminary Design Using Artificial Intelligence and Numerical Optimization Techniques," American Society of Mechanical Engineers, Paper 90-GT-148, June 1990.

[8]Anderson, M. B., "Potential of Genetic Algorithms for Subsonic Wing Design," AIAA Paper 95-3925, Sept. 1995.

[9]Anderson, M. B., "Using Pareto Genetic Algorithms for Preliminary Subsonic Wing Design," AIAA Paper 96-4023, Sept. 1996.

[10]Sharatchandra, M. C., Sen, M., and Gad-el-Hak, M., "New Approach to Constrained Shape Optimization Using Genetic Algorithms," *AIAA Journal*, Vol. 36, No. 1, 1998, pp. 35–42.

[11]Gregg, R. D., and Misegades, K. P., "Transonic Wing Optimization Using Evolution Theory," AIAA Paper 87-0520, Jan. 1987.

[12]Anderson, M. B., Lawrence, W. R., and Lopez, J. L., "Supplementing Gradient Search with Genetic Algorithm in Air Data Estimation System," AIAA Paper 94-1931, 1994.

[13]Selig, M. S., and Coverstone-Carroll, V. L., "Applications of a Genetic Algorithm to Wind Turbine Design," *Journal of Energy Resources Technology*, Vol. 118, March 1996, pp. 22–28.

[14]Cao, H. V., and Blom, G. A., "Navier–Stokes/Genetic Optimization of Multielement Airfoils," AIAA Paper 96-2487, 1996.

[15]Sears, W., "On Projectiles of Minimum Wave Drag," *Quarterly of Applied Mathematics*, Vol. 4, No. 4, 1947, pp. 361–366.

[16]Anderson, M. B., "Software and Users Manual for the IMPROVE Code," Auburn Univ., AL, 1998.

M. S. Miller
Associate Editor