COMBINING APPROXIMATION CONCEPTS WITH GENETIC ALGORITHM-BASED STRUCTURAL OPTIMIZATION PROCEDURES

P. B. Nair* and A. J. Keane†
University of Southampton, Highfield, Southampton SO17 1BJ, U.K

R. P. Shimpi‡
Indian Institute of Technology, Mumbai 400 076, India

ABSTRACT

This paper presents an approach for combining approximation models with genetic algorithm-based design optimization procedures. An important objective here is to develop an approach which empirically ensures that the GA converges asymptotically to the optima of the original problem using a limited number of exact analysis. It is shown that this problem may be posed as a dynamic optimization problem, wherein the fitness function changes over successive generations. Criteria for selecting the design points where exact analysis should be carried out are proposed based on observations on the steady-state behavior of simple GAs. Guidelines based on trust-region methods are presented for controlling the generation delay before the approximation model is updated. An adaptive selection operator is developed to efficiently navigate through such changing and uncertain fitness landscapes. Results are presented for the optimal design problem of a 10 bar truss structure. It is shown that, using the present approach, the number of exact analysis required to reach the optima of the original problem can be reduced by more than 97 %.

1. INTRODUCTION

Genetic Algorithms (GAs) have shown considerable potential in the solution of optimization problems characterized by non-convex and disjoint solution spaces. In the domain of structural optimization, GAs have been applied with fairly encouraging results. Currently, the advantages and flexibilities offered by genetic search procedures for design optimization problems are well understood. For large design spaces, however, a GA typically requires thousands of function evaluations to locate a near optimal solution. Hence, for many large-scale design problems of practical interest, the GA approach may be computationally prohibitive.

A popular and widely followed practice for optimal design is to make use of a gradient based optimization module linked to an approximate analysis routine, which is continuously updated at each design cycle based on the results of exact model analysis. This practice leads to a computationally efficient search procedure, and hence, the solution of large-scale design problems is made possible in a tractable amount of time. Since line search procedures are utilized in gradient based optimization algorithms, the issue of range of validity of the approximation models or the control of approximation errors can be directly addressed by using *ad hoc* move limits or a trust region framework[1] in the line searches.

In contrast, most of the research work related to GA-based optimization has involved the use of problem specific knowledge to increase computational efficiency (see, for example, references[2-3]). It has been shown that problem specific heuristics can be effectively used to achieve performance improvements. However, there are finite limits to the improvements achievable by using problem specific knowledge alone. The history of theoretical developments and applications of gradient based optimization techniques to design indicates that the most influential factor for their wide spread use has been the ease with which approximation models can be incorporated to achieve substantial savings in the computation cost. The development of faster and more efficient optimization algorithms alone would not have sufficed to make this possible. Taking this cue from the evolution of classical design optimization procedures, the question of how to integrate approxi-

*PhD Student, Computational Engineering and Design Center, Department of Mechanical Engineering.

†Department Head and Director, Computational Engineering and Design Center, Department of Mechanical Engineering.

‡Assistant Professor, Department of Aerospace Engineering.

mation models with genetic search procedures needs to be addressed in order to study their practical applicability for many large-scale design problems, where computational cost is a critical issue. This requirement for studies focusing on the extent to which approximation concepts can be combined with stochastic optimization techniques was also noted in a recent survey of the state of the art in multidisciplinary design optimization (MDO) methodologies by Sobieszczanski-Sobieski and Haftka[4].

A study of the literature reveals that there is a dearth of formal methodologies for using approximation models in design procedures based on genetic search. Since GAs make use of probabilistic recombination operators, the control of step size of the design transitions is not straight forward. This leads to difficulties in the integration of approximation concepts, since most of the design transitions may result in perturbations beyond the range of applicability of conventional approximation techniques, which are typically local in nature (e.g., Taylor series, matrix power series, etc.). Hence, one of the major obstacles toward achieving this goal appears to be the non-availability of approximation techniques which are valid for moderate to large changes in the design variables. However, recent developments in approximate reanalysis techniques for linear structural systems (see, for example, references[5-6]) may remedy this situation to a good extent.

It is of interest to note here that some of the recent studies on application of GAs to large-scale optimization problems have made use of function approximations (see, for example, reference[7]). However, no attempt was made to combine the use of both the exact and approximate model. The use of approximations based on linear regression analysis combined with exact analysis was demonstrated recently by Kodiyalam, *et. al.*[8]. It was pointed out in that paper that better computational efficiency will require more sophisticated strategies and approximation methods.

Eby *et. al*[9] addressed the problem of how to use many analysis models of varying fidelity in the design optimization of a composite flywheel. The problem was approached using a parallel injection island GA, wherein groups of populations are evolved independently, with each population using a different analysis model for their fitness prediction. The conjecture made there was that a GA population evolved using a low-fidelity analysis model would give useful building blocks for the population which made use of a higher-fidelity model. This is indeed a reasonable approach to tackle such problems, and has been shown to give good performance improvements for some cases. However, it is expected that the computational savings offered by this line of approach may not be that substantial. This is primarily due to the fact that, in spite of injecting good building blocks into the population, a GA would typically require of the order of hundreds of exact analyses (for a moderately sized problem), to fine tune the building blocks so as to locate the optima corresponding to the original problem. Furthermore, it should be noted that this approach cannot be used when local approximation models are used in the optimization procedure.

The present paper is concerned with the development of an approach for combining approximation models with genetic search procedures. The long term objective is to reduce the computational burden of GA search by reducing the number of exact analysis required to converge to an optimal solution. The approach developed here essentially seeks to create successively improved approximation models based on the results of exact model analysis carried out at key representative points in the design space. The choice of the representative design points at which exact analysis should be carried out is based on observations on the steady-state behavior of a simple GA. This leads to a dynamic optimization problem in which the objective and constraint functions change during the course of the evolution process. Ideas borrowed from the trust region method for nonlinear programming are used to suggest guidelines on when to update the approximation model during the GA search. An adaptive tournament selection procedure is developed for efficiently navigating through such changing and uncertain fitness landscapes. Results are presented for the optimal design problem of a 10 bar truss structure. It is shown that using the proposed approach, the number of exact analysis required to reach the optima of the original problem can be reduced by more than 97 %. Also, it is shown that using the present approach, the GA performance is considerably improved, in terms of both convergence rate and solution accuracy.

2. GENETIC ALGORITHMS

GAs attempt to simulate the phenomena of natural evolution by continuously adapting a population of candidate solutions so as to improve its performance over successive generations. In the context of design optimization, the process of adaptation is carried out by operating on a population of design representations using a selection procedure, and probabilistic operators such as crossover and mutation. The goal of adaptation is to find the design with the highest possible fitness, which is the pseudo-objective function expressed in maximization form. The basic GA used here is fairly standard and the reader is referred to the seminal book by Goldberg[10] for a full description. Some of the main advantages in using a GA for design optimization are outlined below.

1. Gradient evaluations are not required in the search, and hence robustness of the algorithm is relatively unaffected in nonconvex and disjoint solution spaces.
2. Problems with a mix of continuous, discrete and integer variables, and variable complexity design representation can be tackled.
3. There exists the possibility of generating the Pareto optimal frontier for multiobjective design problems in a single run.

The GA used in this study makes use of a tournament selection scheme. This selection scheme has been shown to be robust for a wide variety of problem domains. In this selection scheme, a fitness based tournament is held between s competitors, the winner being the design with highest fitness. Since, only the absolute fitness values are used, this selection scheme is particularly attractive for problems where only the general fitness trends are available. The size of the tournament, i.e., s determines the selection pressure. Selection pressure can be interpreted as the degree to which better designs are favoured in the selection procedure, and hence, the convergence rate of genetic search is in a sense directly controlled by the selection pressure. In GA-based optimization approaches, it is often necessary to tradeoff between robustness and efficiency. Hence, if s is high, even though a good solution will be found quickly, there exists a greater risk of entrapment in a suboptimal solution. In contrast, a low value of s implies a low rate of convergence with increased robustness due to a greater degree of potential allowance for exploration.

3. PROBLEM STATEMENT

Approximation models which are used in engineering optimization can be broadly classified into two categories : (1) Local approximations and (2) Global approximations. It is known that, using a scaling or correction factor (also called the β-correlation method), local characteristics can be imparted to global approximation models (see, for example, reference[11]). This result is of great relevance for the strategies developed here. In other words, any global approximation model can be calibrated to behave like a local approximation model, i.e., there is a tendency for the approximation errors to increase as one moves away from the design points where exact analysis is carried out. Hence, a common strategy for combining approximation models in GA-based optimization becomes applicable to both the classes of models. The design point where exact analysis is carried out is henceforth referred to as an *anchor point.*

The general structural optimization problem can be posed as a nonlinear programming problem of the form :

$$\begin{aligned} &\textbf{Minimize} && \mathbf{F(X)} \\ &\textbf{Subject to :} && \\ & && \mathbf{G(X)} = 0 \\ & && \mathbf{H(X)} \leq 0 \end{aligned}$$

where $\mathbf{F(X)}$ is the design objective function, and $\mathbf{G(X)}$ and $\mathbf{H(X)}$ are the performance constraints. Let $\mathbf{X}^*$ denote the optimal solution to this problem, and $\mathbf{F}^*$, $\mathbf{G}^*$ and $\mathbf{H}^*$ be the values of the objective and constraint functions respectively, at the optimal solution.

Consider the case where an approximation model is available to compute the objective and constraint functions as $\mathbf{F_A(X)}$, $\mathbf{G_A(X)}$ and $\mathbf{H_A(X)}$, respectively. Further, let the approximation model be calibrated at the the anchor design point $\mathbf{X}_{ap}$. As mentioned earlier, the approximation model is exact only at the anchor point, with the prediction error increasing as one moves away from this point. Hence, the anchor point could be viewed as a control parameter which decides the range of validity of the approximate objective and constraint functions.

Approximate analysis can usually be carried out at a fraction of the computational time required for exact analysis. Hence, it is sought to develop an approach which makes minimal use of the exact model, and instead uses a calibrated approximation model to solve the original design problem. The major point of concern in developing such an approach is the guarantee that the modified nonlinear programming problem converges in an asymptotic sense to the optima of the original problem, i.e., $\mathbf{X}^*$.

The modified optimization problem can then be posed as a bilevel optimization problem of the form :

$$\begin{aligned} &\textbf{Level 1 :} && \\ &\quad\textbf{Find} && \mathbf{X}_{ap} \\ &\quad\textbf{Such that :} && \mathbf{F_A^*} \rightarrow \mathbf{F}^* \\ & && \mathbf{G_A^*} \rightarrow \mathbf{G}^* \\ & && \mathbf{H_A^*} \rightarrow \mathbf{H}^* \\ &\textbf{Level 2 :} && \\ &\quad\textbf{Minimize} && \mathbf{F_A(X)} \\ &\quad\textbf{Subject to :} && \\ & && \mathbf{G_A(X)} = 0 \\ & && \mathbf{H_A(X)} \leq 0 \end{aligned}$$

where $\mathbf{F_A^*}$, $\mathbf{G_A^*}$ and $\mathbf{H_A^*}$ are the optimal values of the objective and constraint functions, respectively, after solving the level 2 optimization problem.

It is to be noted here that $\mathbf{F}^*$, $\mathbf{G}^*$ and $\mathbf{H}^*$ are unknown quantities. Hence, in practice the level 1 problem is terminated when no further improvement in the level 2 search problem is possible, and the objective and constraint function value predicted by the approximation model at this converged state agree with that obtained using the exact model. In gradient based optimization, the updated value of $\mathbf{X}_{ap}$ is chosen as the converged optima of level 2. However, it is not computationally feasible to solve this bilevel problem directly in such a fashion using a GA. Hence, in the present approach, a solution is sought for the modified optimization problem using a single level GA.

It can be seen from the modified problem statement that it is the choice of $\mathbf{X}_{ap}$ which dictates whether the sequence of iterates converge to $\mathbf{X}^*$. Hence, it is necessary to start with an initial choice of the anchor point, say $\mathbf{X}^o_{ap}$, calibrate the approximation model, use it for fitness predictions required by the GA and then continuously adapt the anchor point as the subsequent generations of designs are evolved.

3.1 Some Important Issues

The development of an approach for combining approximation concepts with genetic search procedures involves addressing three fundamental issues

1. Given a set of design vectors in the solution space (typically forming a generation), which points should be chosen as anchor points ?
2. In a given generation, how many anchor points should be chosen ?
3. When should the approximation model be recalibrated ?

The magnitude of error in an approximately evaluated objective/constraint function value at a certain design point depends on its closeness to the anchor point where exact analysis was carried out. When a number of points in the design space are evaluated more exactly than the rest, the simulated design space could be viewed as an uncertain or noisy objective function. For example, if the anchor point is changed, the predicted fitness value for a given design vector would be different for each realization of the calibrated approximation model.

It has been noted in earlier studies on GA function optimization that the GA is robust to certain types and amounts of noise in the objective function. Detailed investigations and results for optimization of noisy functions can be found in the dissertation of Miller[12]. Since the GA makes use of stochastic operators in the search procedure, it can be expected that the uncertainties in the fitness predictions which arise due to the use of approximation models can be ultimately overcome. Also, since the expected value of the fitness function is implicitly used to search for promising regions in the solution space, the GA would tolerate uncertainties arising from approximations. Hence, it is reasonable to expect that the uncertainty in fitness predictions could be directly handled by the GA without an explicit formulation for the expected value and bounds of the uncertainties.

4. PROPOSED APPROACH

This section details the elements of the approach which is developed here to address the three fundamental issues raised in the previous section. For the sake of simplicity only the case involving the use of one approximation model is considered.

4.1 Selecting Representative Anchor points

The most simplistic approach that can be proposed involves a static, random choice of the anchor points. However, initial numerical investigations and the considerations outlined earlier rule out this approach due to potential lack of robustness and poor overall performance. Moreover, such a choice makes it difficult to explain empirically whether the sequence of best solutions will converge to the optima of the original problem.

Consider the convergence trends of a simple GA which is used to optimize a static objective function. Here, as the number of generations increase, the population tends to saturate with designs close to the best value. From this observation, two statements can be made for a GA which is approaching steady-state :

1. The average or mean of all the design vectors in a generation is close to the best design (assuming the effects of genetic drift to be low).
2. The niche count (i.e., the number of designs in an arbitrary neighborhood) of the best design is highest compared to the other designs present in a generation.

The above observations suggest three possibilities for choosing the anchor point at each generation : (1) the average or mean of all the design vectors, (2) the design point with highest niche count, and (3) the design vector with highest fitness in the previous generation. Once the exact model is evaluated at the chosen anchor point, the result could be used to calibrate the approximation model for predicting the fitness of all the designs in that generation.

At a first glance, choices 1 and 2 appear to be better than choice 3 since these choices will result in approximations which are unbiased to potential performance. In order to make use of choice 3, in the initialization phase, the anchor point could be chosen randomly, and in the subsequent generations, the best individual in the previous generation is selected as the anchor point. However, it is considered that this choice could potentially lead to myopic search, with a greater risk of entrapment in suboptimal solutions due to fitness biased approximation error trends.

Interestingly, all the three choices of anchor point would result in the average approximation error in a generation reducing as the GA approaches steady state. Therefore all three choices would result in a similar anchor point as the GA approach steady-state. However, the search trajectory of the GA is expected to be different for each of these choices.

4.2 Updating the Approximation Model

Due to computational cost considerations, it may not always be feasible to evaluate the exact model at each generation. Hence, it is of interest to investigate whether it is possible to introduce a generation delay before evaluating the exact model at a new anchor point. Henceforth, the number of GA generations over a certain value of generation delay is referred to as a *design cycle.* It is suggested that ideas from trust-region frameworks could be used for developing an efficient procedure for this purpose. Initially one could choose a generation delay arbitrarily, say 4. Then, at the fifth generation, the approximate fitness of the best design found so far or the current average design vector, could be compared with the fitness value predicted by the exact model. This information on how well the approximation model is predicting the trends of the exact model could then be used to control the generation delay. This would allow an adaptive procedure for expanding and contracting the generation delay in accordance with the performance of the approximation model.

Let f and f_a denote the fitness function value predicted by the exact and approximation model, respectively, for the best design or the average design vector at the termination of a particular design cycle. Let f^o denote the corresponding exact fitness value at the beginning of this design cycle. Define $\Delta f = (f - f^o)/(f_a - f^o)$, and let ϵ_1 and ϵ_2 be small positive numbers satisfying the inequality $\epsilon_1 < \epsilon_2$. The algorithm for controlling the value of the generation delay (based on the trust region framework presented in reference[1]) for the next design cycle is given below as :

Case 1 : $\Delta f \geq \epsilon_2$. This case corresponds to the scenario where the approximation model is doing a good job of finding improvement in f over the current value of the generation delay, or there was more improvement in f_a than that in f. Hence, the the generation delay can be increased since the approximation model has proven its utility in finding improvement in the exact fitness f over the current value of the generation delay.

Case 2 : $\Delta f \leq \epsilon_1$. In this case, the approximation model is doing a bad job of predicting the improvement in f, i.e., the improvement in fitness predicted by f_a was far lower compared to that in f or f actually increased . Hence, the generation delay is decreased for the next design cycle.

Case 3 : $\epsilon_1 < \Delta f < \epsilon_2$. Here, the approximation model is doing an acceptable but not especially noteworthy job of predicting the improvement in f. Hence, the generation delay is kept unchanged.

To prevent premature saturation of the population, it is important to ensure that the generation delay never exceeds a certain threshold value which is fixed *a priori*. In the trust region approach, the values of ϵ_1 and ϵ_2 are usually chosen as chosen as 0.10 and 0.75, respectively. A conservative choice of the increment in generation delay could be unity. Note here that there also exists the possibility of increasing the value of ϵ_1 and ϵ_2 as the number of design cycles increase. This would allow fine tuning or reduction of the generation delay as the GA reaches the optimal solution.

It is of interest to note that the same guidelines could also suggest when an approximation model of different fidelity could be used, i.e., one not only controls the generation delay, but also the approximation model used in the design cycle. This issue of model management is important from the standpoint of computational efficiency, and requires further detailed studies.

4.3 Choosing the Number of Anchor Points

Clearly, the number of anchor design points at which exact analysis should be carried out depends on the expected error bounds or range of applicability of the approximation procedure as well as computational cost considerations. For cases in which a low-fidelity approximation procedure is employed, a domain decomposition procedure could be carried out and the anchor points computed for each subdomain. This would serve the purpose of controlling the errors in fitness prediction of the approximation models. Note that in the case of parallel GA implementations, to reduce inter-processor communication overhead due to the requirement of transmitting calibration data to the approximation modules, it would be preferable to carry out the fitness evaluation of each sub-domain independently at different nodes. Hence, load balancing considerations could dictate that the number of design vectors in each sub-domain to be the same. One possible way to do this would be to create clusters into which a design vector is allocated based on the value of the Holder norm ($||\mathbf{X}||_n$) defined below as :

$$||\mathbf{X}||_n = (\sum_{i=1}^{p} \mathbf{x}_i^n)^{1/n}$$

where p is the total number of design variables. The value of n could be chosen so as to define the shape of each domain.

An alternative to choosing more than one anchor point is to create a database of the anchor points generated as the search proceeds. Then the fitness evaluation of a design vector in a generation could be carried out using the approximation model calibrated at the anchor point closest to it, or by averaging the fitness predictions of a family of calibrated models. However, this is feasible only if the approximation model uses a small amount of data for calibration purposes, and further, convergence problems due to the magnitude of approximation errors of the low-fidelity model outweigh the computational burden of database creation, storage and search. The other foreseeable disadvantage of this approach is the possible requirement of interprocessor communication for parallel GA implementations.

4.4 Counteracting the Detrimental Effects of Reduced Selection Pressure

For any choice of the anchor point outlined earlier, the errors in fitness predicted by the approximation model will reduce gradually as the GA approaches steady-state. For example, if the constraints are approximated, the constraint boundary dynamically shifts from generation to generation, with equilibrium being achieved only at the theoretical steady-state. This situation is analogous to genetic adaptation in dynamic environments.

Currently, litle is known on how to guide adaptation in such changing environments. However, it is generally felt that the concepts of diploidy (dominant and recessive genes) could be used to deal with such problems. Alternatively, a measure of the approximation error (e.g., the Eucledian distance of each design vector from the anchor point) could be used along with the phenotype/fitness information in the selection procedure. The present study does not explore these possibilities. Instead, a simple adaptive selection scheme is developed to address the problem of variable selection pressure which arises due to the fitness uncertainties. In the present problem, the dynamic fitness landscape will asymptotically tend to the static (exact) fitness, as the GA approaches steady state. The important domain specific knowledge we have here is that the variance of the design vectors indicates whether the fitness landscape is going through a turbulent phase. If the variance of the design vectors is high, the approximation errors are expected to be high. The converse is true if the design vector variance in a generation is low.

Consider a binary tournament selection scheme in which a tournament is conducted between two competing designs based on their fitness evaluated at a certain generation. If approximation errors are high at that generation, then the resultant selection pressure is low, i.e., the true fitness of the winner may not necessarily be higher than that of the loser. This observation was originally made in the context of GA optimization of noisy functions[13]. This reduction in the selection pressure could potentially lead to a slow rate of convergence to an optimal solution.

To counteract the detrimental effects of variable selection pressure, a modified tournament selection procedure is required in which the tournament size is adaptively selected based on a quantitative estimate of the approximation errors in that generation. This is to ensure that the reduction in selection pressure due to approximation errors do not slow down the GA search considerably. As mentioned earlier, the variance of the design vectors in a generation gives a good idea of the likely magnitude of errors in using the approximation model. Here, three possible values of the tournament size are arbitrarily chosen as 2, 3 and 4. Note that these choices for the tournament size have made based on past experiences, and also to minimize the possibility of premature convergence due to excessively high selection intensity. Let $Avg(\Delta\mathbf{X})$ denote the average of the Eucledian distance of each design vector from the anchor point at a particular generation, i.e.,

$$Avg(\Delta\mathbf{X}) = \frac{1}{m}\sum_{i=1}^{m} ||\mathbf{X}^i - \mathbf{X}_{ap}||$$

where m is the population size.

The tournament selection size s is then chosen using the rules given below.

If $Avg(\Delta\mathbf{X})$ is high

$s = 4$

If $Avg(\Delta\mathbf{X})$ is medium

$s = 3$

If $Avg(\Delta\mathbf{X})$ is low

$s = 2$

The intervals for high, medium and low values of $Avg(\Delta\mathbf{X})$ are problem dependent. Let us consider the case wherein there are p design variables, each of them normalized between 0 and 1. Here, for example, one could choose the intervals for $Avg(\Delta\mathbf{X})$ as - (1) High $\Rightarrow Avg(\Delta\mathbf{X}) > \frac{p^{1/2}}{2}$ (2) Medium $\Rightarrow \frac{p^{1/2}}{5} < Avg(\Delta\mathbf{X}) \leq \frac{p^{1/2}}{2}$ (3) Low $\Rightarrow Avg(\Delta\mathbf{X}) \leq \frac{p^{1/2}}{5}$ It is felt that the robustness of the GA would not be very sensitive to the choice of this interval. In the worst case, if the bounds on medium and high values are widened, the resulting selection pressure (i.e., 3 or 4) is unlikely to lead to premature saturation for a reasonable population size and fixed number of generations. Hence, regardless of the choice made for this interval, the convergence speed and solution accuracy using a adaptive tournament size is expected to better than that achievable by employing a static binary tournament selection operator.

4.5 Empirical Analysis of the Present Approach

This section presents an empirical explanation of how the present approach affects the GA search mechanisms, using analogies with natural evolution. The two main points raised here are survival and adaptation. In the present problem, survival is a function of the fitness as well as the state of the environment. In contrast, for static function optimization, survival depends on the fitness alone. Adaptation in the population occurs when the average fitness increases over successive generations.

In the present approach, the risk of survival of any individual is directly related to the closeness of its genotype to that of the anchor point. Here, the anchor point could be interpreted as a representative individual in a population. In this fitness landscape, it is necessary for all the individuals to stay close enough to the anchor point to ensure survival. If individuals stray away from the the anchor point, the turbulence in environmental payoff due to approximation errors could potentially drive the deviant individuals to extinction. The term potential is used, since the fitness of deviant individuals could be higher or lower than the true value (assuming a sufficiently multimodal landscape). However, when finding the optimum solution or the state with highest fitness, survival instincts have to be balanced by adaptation, so that the average fitness of the population increases over generations. The action of other GA operators is to encourage adaptation. These are two opposing forces, and for the GA to reach the optimum, these forces should co-evolve so as to reach a state of equilibrium. The introduction of a generation delay is expected to help the adaptation process since this procedure keeps the fitness function static for a finite number of generations.

It is known that evolution of a single species over a sufficiently large number of generations will inevitably lead to homogeneity in the population. Hence, it is hoped that in spite of turbulence in the environmental payoff, the population will eventually converge to a homogeneous state.

Next consider how the adaptive selection operator acts on the evolution process. If there are many deviants from the representative individual in a generation, the selection pressure is kept high. This means that in the selection process, there is a greater tendency to favor the fittest individual. Herein there are two possibilities : either the fitness of this individual is close to its true value or it is undeserved. Now, in the next generation, there will be a tendency for the population to drift close to the genotype of this individual. This implies that the genotype of the representative individual of this generation is closer to the genotype of the fittest individual of the previous generation. Because of this, the turbulence in the environmental payoff for individuals with this genotype will decrease, and hence, a better estimate of the true fitness is achievable. The use of this better estimate would serve the purpose of correcting the search trajectory, if necessary, i.e., this mechanism ensures that the GA cannot be deceived for long if approximation errors guide the search to erroneous regions of the solution space. In contrast, if the selection pressure is kept static and low, it will take a longer number of generations to discover if the fitness of the best individual was undeserved or not. This leads to a slow convergence rate for the search procedure. When the number of deviants decrease, the adaptive selection operator will dictate that a lower selection pressure be used for this generation. This means that exploration is encouraged only when it is safer to do do, i.e., when the genotype of most of the individuals is similar to that of the representative individual.

5. DEMONSTRATION EXAMPLES, RESULTS AND DISCUSSIONS

This section demonstrates some aspects of the proposed approach for the well documented 10-bar planar truss structure. The fixed parameters and design variable bounds for this problem have been directly taken from reference[14], where a GA was employed to locate an optimal solution. Results are presented for three problems. Problems 1 and 2 involve the use of cross-sectional areas as design variables, i.e., a total of ten variables. In problem 1, only stress constraints are considered and in problem 2, both stress and displacement constraints are considered. Problem 3 involves simultaneous optimization of the sizing and geometry variables (a total of 18 variables) to mini-

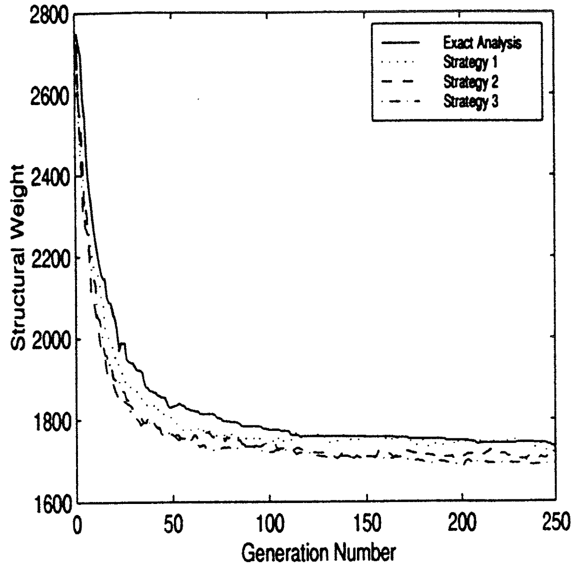

Figure 1: Comparison of Averaged Convergence History of Different Approaches for Problem 1

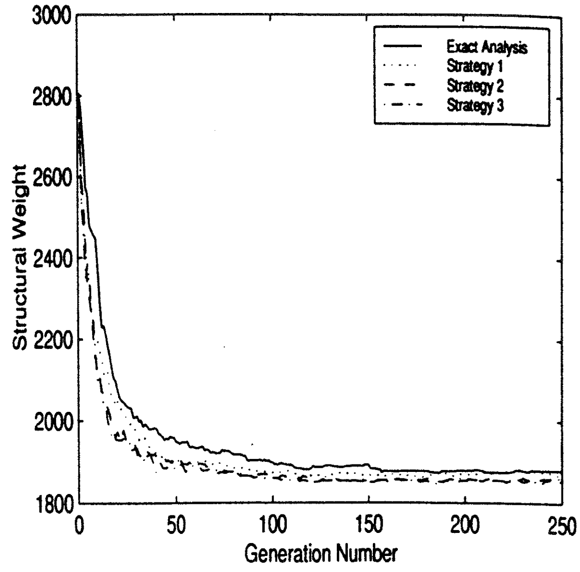

Figure 2: Comparison of Averaged Convergence History of Different Approaches for Problem 2

mize weight subject to stress constraints. The joint coordinates of the truss are allowed to move ± 100″ from the original positions. The stress and displacement constraints for all problems are approximated using a first-order reduced basis method proposed by Kirsch[5]. The pseudo-objective function is formulated using an exterior extended penalty function approach.

This paper presents results for three different variants of the approach outlined earlier :

Strategy 1 : Conventional binary tournament selection.

Strategy 2 : Adaptive tournament selection operator.

Strategy 3 : The guidelines presented in section 4.2 are used for determining the generation delay before the exact model is evaluated. The initial and maximum possible value of the generation delay are chosen as 4 and 8, respectively. An adaptive tournament size is used in the selection procedure.

For all the results presented in this paper, the average design vector was chosen as the anchor point. Further, only one anchor point is evaluated at a given generation.

A population size of 40 was chosen for problems 1 and 2, whereas for problem 3 a larger population size of 50 was used. The probabilities of uniform crossover and bit mutation were kept constant at 0.6 and 0.01, respectively. Creep mutation was applied at a probability of 0.2. An elitist strategy was used to prevent loss of the fittest design. The cross-sectional areas and the joint coordinates were encoded using 6 bits and 10 bits, respectively. The number of GA generations were kept constant for all the runs at 250 and 300 for problems 1,2 and problem 3, respectively. Simulations were carried out for ten different initial populations to compare the performance of the presented approach with a conventional GA which makes use of exact analysis throughout the search. Further, the same initial population was used to run each strategy to ensure a consistent comparison throughout.

Note that for optimization using strategy 1 and 2, exact analysis is carried out only once per generation, i.e, for the average design vector. For all the design vectors in the population, the stress and displacement constraints are approximated using the results of this precise analysis. In contrast, strategy 3 uses a variable number of exact analysis which is equal to the total number of design cycles.

The convergence trends of the GA search for problem 1 averaged over 10 runs for the three strategies are compared with GA search using exact analysis only in Figure 1. A statistical analysis of the GA performance for this problem is summarized in Table 1. The corresponding results for problem 2 are shown in Figure 2 and Table 2. The convergence trends and optimal solution obtained by the different solution approaches for problem 3 are summarized in Figure 3 and Table 3, respectively.

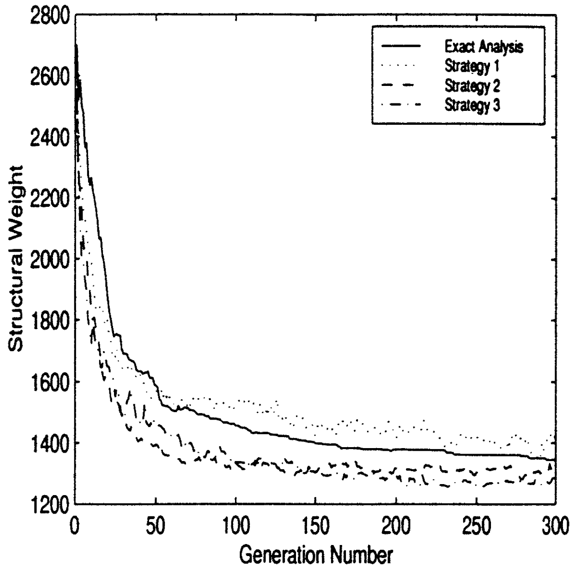

Figure 3: Comparison of Averaged Convergence History of Different Approaches for Problem 3

Based on these results, the following observations can be made :

1. Rather surprisingly it appears that the use of approximation concepts seems to improve the GA search accuracy instead of inhibiting convergence to a near optimal solution. The main reason for this lies in the fact that the GA search procedure exploits the variable selection pressure mechanism induced by the approximate fitness predictions to improve its search efficiency and accuracy. This could also be attributed to the nature of the genetic operators of selection, crossover and mutation, which allow the individuals to adapt in varying fitness landscapes.

2. Strategy 2 and strategy 3 give the best results for all the problems. The difference between these two strategies do not appear to be very significant for problem 2. Further, it can be seen from Figures 1-3 that these strategies converge to the vicinity of the optima faster than strategy 1. This demonstrates clearly how the use of an adaptive tournament selection size helps the GA procedure to tradeoff between exploration and convergence speed in accordance with the uncertainties in fitness predictions. Also, the standard deviation in the optimal solution found by the GA is generally found to decrease when approximate fitness predictions are used.

3. For problems 1 and 2, it was observed that approximation errors encountered during the initial GA generations were usually less than $\pm$ 20%. These errors reduce rapidly to the order of $\pm$ 5%, which leads to the generation delay in strategy 3 always increasing to 8, after which it is kept fixed. These trends could be attributed to the global characteristics of Kirsch's approximation method and the choice of the anchor point. However, even such low approximations errors were observed to lead to oscillation of the best solution in the vicinity of the constraint boundary, especially for problem 2, where a displacement constraint is active at the optima. This oscillatory behavior is to be expected when approximation models are used in the search, and has been reported earlier for gradient based optimization algorithms. Two possible ways to alleviate this undesirable behavior are to either increase the number of anchor points in the final stages of the GA search or use a fuzzy constraint handling strategy.

4. In problem 3, the use of geometric design variables lead to approximation errors as high as 100%. Hence, this problem is representative of many practical applications where high-fidelity approximation models are seldom available. This high magnitude of approximation errors leads to difficulties in finding a feasible solution for strategy 1, which primarily relies on the steady-state hypothesis to converge to the optima of the original problem. This is mainly because the population size and number of generations used for this problem are not high enough for the GA to reach steady-state. In general, constraint violations of the order of 5-6% were observed for strategy 1. In contrast, strategy 2 and strategy 3 had no difficulty in finding a feasible solution. This can be primarily attributed to the adaptive tournament selection operator which enforces a pseudo steady-state. Also, strategy 3 performs better than strategy 2 for this problem because of the generation delay which allow individuals a reasonable number of generations to adapt to changing fitness environments.

5. Substantial savings in the computational cost have been achieved by using the proposed approach. For all the problems considered, the number of exact analysis required to converge to an optimal solution could be reduced by more than 97%. In particular, for strategy 3 the reduction in number of exact analysis were of the order of 99.6%.

6. CONCLUDING REMARKS

An approach for combining approximation models with genetic search procedures has been developed. The important issues which arise in the development of a computationally efficient framework for tackling this problem are addressed. It is shown that this problem can be posed as a bilevel optimization problem, the solution of which involves adaptation in a dynamic environment. Numerical results obtained for a simple demonstration example provide a good idea of the computational savings offered by the presented approach. In general, the number of exact analysis may be reduced by more than 97%. Also, it is shown that both the GA search efficiency and accuracy are improved by using approximation models for fitness predictions.

Further work is required to study the performance of the present approach, particularly when other approximation models of lower fidelity are used in the search procedure. It is expected that the use of the domain decomposition strategies outlined here would be useful for controlling approximation errors during the GA search. The possibility of using a measure of the expected approximation errors (i.e., the Eucledian distance of the design vector from an anchor point) along with the fitness information in the selection procedure also merits consideration. Development of efficient architectures for parallelizing the present approach is another research area worthy of future investigations.

This research has several long-term ramifications. Most importantly, it shows that approximation concepts can indeed be combined with a great degree of success in GA-based optimization procedures. Further, this study also serves as a demonstration of the inherent power of the approximation concepts being employed, and the robustness of GAs in dynamic and uncertain environments.

ACKNOWLEDGMENTS

This work was supported by a grant from the Faculty of Engineering and Applied Sciences at the University of Southampton.

REFERENCES

[1]Alexandrov, N., Dennis, J. E., Lewis, R. M., and Torczon, V., "A Trust Region Framework for Managing the Use of Approximation Models in Optimization," *ICASE Report No. 97-50*, 1997.

[2]Furuya, H., and Haftka, R. T., "Combining Genetic and Deterministic Algorithms for Locating Actuators on Space Structures," *Journal of Spacecraft and Rockets*, Vol. 33, No. 3, 1996, pp. 422-427.

[3]Nagendra, S., Haftka, R. T., Gurdal, Z., and Watson, L. T., " Improved Genetic Algorithms for the Design of Stiffened Composite Panels, " *Computers and Structures*, Vol. 58, No. 3, 1996, pp. 543-555.

[4]Sobieszczanski-Sobieski, J., and Haftka, R. T., "Multidisciplinary Aerospace Design Optimization : Survey of Recent Developments," *34th AIAA Aerospace Sciences Meeting and Exhibit*, AIAA Paper No. 96-0711, Reno, Nevada, 1996.

[5]Kirsch, U., "Improved Stiffness-Based First-Order Approximations for Structural Optimization", *AIAA Journal*, Vol. 33, No. 1, 1995, pp.143-150.

[6]Nair, P. B., Keane, A. J., and Langley, R. S., "Improved First-Order Approximation of Eigenvalues and Eigenvectors," *AIAA Journal*, to appear.

[7]Lee, J., and Hajela, P., "Parallel Genetic Algorithm Implementation for Multidisciplinary Rotor Blade Design," *Journal of Aircraft*, Vol. 33, No. 5, 1996, pp. 962-969.

[8]Kodiyalam, S., Nagendra, S., and DeStefano, J., "Composite Sandwich Structural Optimization with Applications to Satellite Components," *AIAA Journal*, Vol. 34, No. 3, 1996.

[9]Eby, D., Averill, R. C., Gelfand, B., Punch, W. F., Mathews, O., and Goodman, E. D., "An Injection Island GA for Flywheel Design Optimization" Invited Paper, *Proceedings of the 5th European Congress on Intelligent Techniques and Soft Computing*, Sept., 97.

[10]Goldberg, D. E., "*Genetic Algorithms in Search, Optimization and Machine Learning*," Cambridge MA : Addison-Wesley, 1989.

[11]Chang, K. J., Haftka, R. T., Giles, G. L., and Kao, P. -J., "Sensitivity-Based Scaling for Approximating Structural Response," *Journal of Aircraft*, 30, 1993, pp. 283-288.

[12]Miller, B. L., "Noise, Sampling and Efficient Genetic Algorithms," *PhD Dissertation*, University of Illinois at Urbana-Champaign, 1997.

[13]Miller, B. L., and Goldberg, D. E., "Genetic Algorithms, Tournament Selection and the Effects of Noise," *IlliGAL Report No. 95006*, University of Illinois at Urbana-Champaign, 1995.

[14] Hajela, P., and Yoo, J., "Constraint Handling in Genetic Search Using Expression Strategies," *AIAA Journal*, Vo. 34, No. 11, 1996, pp. 2414-2420.

Approach	Exact Analysis	Strategy 1	Strategy 2	Strategy 3
Average Weight	1729	1718	1714	1694
Standard Deviation	22	12	14	10
Minimum Weight	1698	1661	1655	1625
No. of Exact Analysis	10000	250	250	33 (Avg.)

Table 1: Comparison of Optimization Results for Problem 1

Approach	Exact Analysis	Strategy 1	Strategy 2	Strategy 3
Average Weight	1880	1862	1854	1861
Standard Deviation	9	7	2	2
Minimum Weight	1855	1823	1842	1854
No. of Exact Analysis	10000	250	250	33 (Avg.)

Table 2: Comparison of Optimization Results for Problem 2

Approach	Exact Analysis	Strategy 1	Strategy 2	Strategy 3
Average Weight	1345	1407	1301	1279
Standard Deviation	21	15	19	16
Minimum Weight	1219	1342	1236	1185
No. of Exact Analysis	15000	300	300	39 (Avg.)

Table 3: Comparison of Optimization Results for Problem 3