Artificial Intelligence for Engineering Design, Analysis and Manufacturing (1998), *12*, 431–445. Printed in the USA.

Evolutionary and adaptive strategies for efficient search across whole system engineering design hierarchies

I.C. PARMEE
Plymouth Engineering Design Centre, University of Plymouth, Drake Circus, Plymouth, PL4 8AA, Devon, U.K.
(Received May 13, 1997; Accepted June 5, 1998)

Abstract

Evolutionary and Adaptive strategies (ES & AS) for diverse multilevel search across a preliminary, whole-system design hierarchy defined by discrete and continuous variable parameters are described. Such strategies provide high-level decision support when integrated with preliminary design software describing the major elements of an engineering system. Initial work involving a Structured Genetic Algorithm (stGA) with appropriate mutation regimes to encourage search diversity is described and preliminary results are presented. The shortcomings of the stGA approach are identified and alternative strategies are introduced. A dual agent strategy (GAANT) involving elements of an ant colony search and an evolutionary search concurrently manipulating the discrete and continuous variable parameter sets is presented. Appropriate communication between the two search agents results in a more efficient search across the hierarchy than that achieved by the stGA, while also simplifying the chromosomal representation. This simplification allows the further development of the preliminary design hierarchy in terms of complexity. The technique therefore represents a significant contribution to configuration design where multilevel, mixed discrete/continuous parameter design problems can be prevalent.

Keywords: Genetic Algorithms; Adaptive Search; Design Hierarchies; Whole System Design

1. INTRODUCTION

During the initial stages of a large-scale, engineering design project the engineer will be presented with the task of identifying initial high-potential system configurations that best satisfy many performance criteria. Such criteria may be qualitative and quantitative in nature and will likely be ill-defined and either incomplete or over emphasized during these higher level stages of the design process. Design at this stage is generally a multidisciplinary exercise requiring extensive cooperation of several groups with varying goals working concurrently within an uncertain and therefore high-risk design environment.

The research described in the paper is primarily concerned with the feasibility/bid stages of large-scale bespoke design projects. These early stages are characterized by a high degree of uncertainty related to limited available data and lack of definition in the initial design brief. A degree of necessary assumption is therefore required and a breadth-first approach is evident. The detail of any system modelling must be commensurate with the degree of confidence in the available data and it therefore follows that caution must be exercised when interpreting the results from such preliminary design tools. Engineering judgement based upon problem-specific knowledge and an understanding of the shortfalls of the software plays a significant role.

The hypothesis underlying the research described is that during these early stages of design, the engineer requires a flexible tool that will provide an efficient search of a high-dimensional design space. The search will result in the identification of high-performance solutions from diverse regions of a design hierarchy described by discrete design decisions and continuous variables.

The genetic algorithm (GA) (Holland, 1975; Goldberg, 1989) generation of a number of high-performance solutions based upon quantitative criteria can provide sufficient information for the engineer to form a qualitative judgement. This benefits the significant requirement for caution related to problem uncertainties and allows conflicting qualitative criteria to be taken into consideration. The selection

Reprint requests to: I.C. Parmee, Plymouth Engineering Design Centre, University of Plymouth, Drake Circus, Plymouth, PL4 8AA, Devon, U.K. E-mail: iparmee@plymouth.ac.uk

of the most appropriate solutions or design region results in a significant reduction in the size of the design space and subsequent concentration of further search in high potential areas. The objective is not to locate the global optimum of the space described by the mathematical model but to select some local optima that best satisfies all current qualitative and quantitative criteria relating to the current design requirements. The assessment, through designer interaction and off-line processing utilizing problem-specific knowledge, of such solutions may lead to the recognition of optimal design direction at that stage in terms of least risk and best performance.

The research therefore concerns the utilization of the GA to achieve an efficient concurrent, multilevel search across hierarchical structures described by discrete and continuous variable parameters. This search should result in the identification of high-performance solutions from diverse regions of such a hierarchy. The continuous variables are directly related to the differing system configurations described by the selected discrete design options. Continuous design sets may therefore differ in terms of number of variables and their combination as illustrated by the simple hierarchical representation of Figure 1. This results in the creation of many continuous design spaces of differing character, size and complexity, each dependent upon specific discrete design configurations. The requirement therefore is for a search strategy that can initially maintain a diverse search across the discrete elements of the hierarchy, while optimally sampling the differing dependent continuous design sets. The initial objective is a rapid decomposition of the hierarchy into a small number of high-performance configurations through the concurrent identification of high potential, diverse design domains.

The overall objective is the development of an exploratory, high-level, decision support tool, which will significantly reduce lead times during this "whole-system" stage of design, while allowing a more extensive search of the available design alternatives. This should result in the identification of competitive solution domains that may have been overlooked during the problem decomposition processes of traditional heuristic design. A high degree of engineer interaction is envisaged in order that qualitative criteria based upon previous design experience can be applied to local, high-performance solutions. Most importantly the strategies will enable the engineer to rapidly survey the potential of diverse regions of a multilevel hierarchy. This offers an alternative to compromising the search space potential by immediately returning to familiar design configurations from previous studies, which may result in premature concentration of search effort and less opportunity for the discovery of innovative solutions.

The GA-based selection of the most appropriate solutions or design regions and subsequent off-line assessment and processing results in a significant reduction in the size of the design space and subsequent concentration of further search in the most appropriate areas. The research therefore concerns the utilisation of the GA to achieve an efficient concurrent, multilevel search across hierarchical structures described by discrete and continuous variable parameters. Much of the initial research in this area involved the manipulation of a design hierarchy for large-scale hydropower systems (Parmee, 1995, 1996) (Fig. 2) and it is this domain that provides the case study for the following paper. Another similar hierarchy currently under investigation concerns the steam cycle of thermal power generation plant, where the discrete elements relate to plant configuration and continuous variables include dimensional parameters and system control settings (Parmee et al., 1996; Chen et al., 1997). The intention is that the strategies and techniques described here should be generic in nature being applicable across a wide range of whole-system engineering design problems.

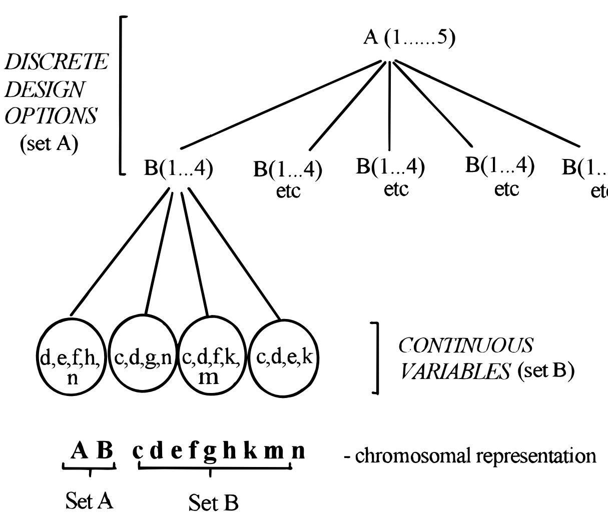

Fig. 1. Simple hierarchy.

2. THE HYDROPOWER SYSTEM

The paper illustrates the development of appropriate adaptive search (AS) strategies by concentrating upon the hier-

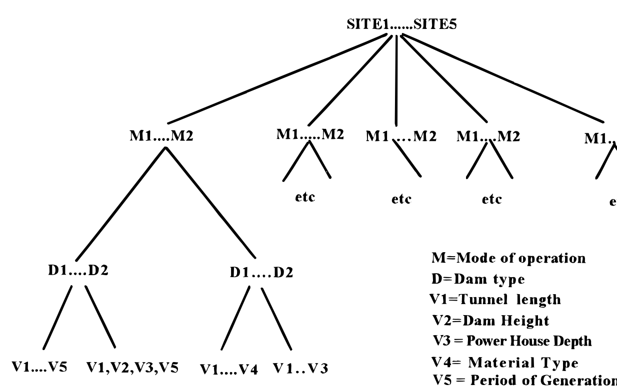

Fig. 2. Basic hydropower design hierarchy.

https://doi.org/10.1017/S0890060498125039 Published online by Cambridge University Press

archy of Figure 2, which simply describes major elements of a large-scale hydropower system. In such a case, the design team will initially be tasked with the identification of high potential hydropower sites either within a particular geographic region or within a specific river network. Although a large number of potential sites may exist, engineering judgement based upon experience, knowledge, and basic analysis will eliminate the majority. However, a small number of sites with differing topographic characteristics will remain and further elimination becomes impossible without a lower-level design and economic analysis to determine best configurations of the main structural elements of the hydropower system and appropriate operational characteristics.

At this prefeasibility stage of the design process, the engineer will be working with limited data relating to ground conditions, flow regimes, etc. Information concerning such aspects will be arriving daily as teams carry out preliminary surveys of the appropriate region. Economic analysis based upon available data will be concurrently identifying the output requirements and costs of the proposed system to establish the project's viability against other power generation development based upon alternative fuel sources. The process is therefore initially characterized by a high degree of uncertainty and much iteration based upon available data and related necessary assumption. Iteration will continue until the level of risk is sufficiently low for firm decisions to be made concerning appropriate design directions. The detail of the system modelling during these preliminary stages must be commensurate with the degree of confidence in the available data. Computational expense must be kept to an acceptable level to allow rapid iteration and comparison of results. The overall objective is to determine the "best" site that, when combined with the related optimal overall system configuration, will provide maximum power output at minimum cost.

The establishment of the basic feasibility of such systems is a time-consuming and therefore costly process. The degree of design iteration generally carried out in a trial-and-error manner under the guidance of domain experts consumes many person-hours. This is unsurprising when one considers a basic representation of the hierarchical structure of the hydropower problem as shown in Figure 2. This simple hierarchy, although only involving a few of the main elements of the system, already describes a complex overall design space of some 10^9 differing site/system configurations. Traditionally the hierarchy would be decomposed down to a level where the dimensionality is sufficiently low to allow a degree of meaningful search. This approach severely compromises the design potential by restricting investigation to local subsets of the design hierarchy that are of a dimension that can be realistically investigated within time and budget constraints. The research described in the following paper addresses this problem by investigating the utilization of AS techniques such as the GA to achieve an efficient concurrent, multilevel search of these complete hierarchical structures.

2.1. The model

At the current state of development, the model of the hydropower system must be regarded as illustrative as opposed to definitive. It is essential, initially, to determine the feasibility of an efficient adaptive search across a highly discontinuous search space described at differing levels by discrete and continuous variables. Research has concentrated upon the recognition of those characteristics that enable the search process to concurrently investigate diverse regions of the design hierarchy. However, although not definitive, the model represents practical design aspects of most of the major subsets of the overall system design. The subsets described by the current model are:

- *The Site*: Five sites are considered each with differing dimensional characteristics relating to idealized valley cross-section and longitudinal valley sections. The longitudinal sections extend up to two kilometers downstream of the dam.
- *Mode of Operation*: Two distinct modes of operation are considered: Baseload supply, where all of the available flow is constantly utilized for power generation and peak power supply, where generation is restricted to those periods providing an increased unit price. The peak power option introduces a further variable parameter—*Period of Generation*—which sets the number of hours of energy production per day. This parameter is only active when the peak power option is selected. Annual energy output is calculated from:

$$\text{Baseload supply: } MWh = 8760 \times Q_i \times Ph_i \times 9.81 \times 10^{-3} \quad (1a)$$

$$\text{Peak power supply: } MWh = 365 \times \frac{Tdf_i}{Gp_i} \times Ph_i \times Pf \times 9.81 \times 10^{-3}, \quad (1b)$$

where

MWh = Megawatt hours produced per annum,

Q = flow rate at site i,

Tdf = total daily flow available at site i,

Gp = total daily generation period at site i,

Ph = potential head at site i, and

Pf = unit price factor (circa 1.5).

Peak power generation incurs cost penalties on pressure tunnel construction, due to the necessary accommodation of significantly increased flow rates. Peak energy output is increased by the unit price factor to allow for the enhanced value of power out.

- *The Dam*: Two dam types are considered—Concrete Gravity and Embankment. A range of materials are

https://doi.org/10.1017/S0890060498125039 Published online by Cambridge University Press

available for the embankment dams resulting in differing slope characteristics and varying dam volume. The embankment dam model takes into account type of material and calculates dam volume (EDv) based upon upstream and downstream slope and dam height (EDh). The width of the dam core (EDw) remains constant for each site. It is assumed that the same material is used both upstream and downstream of the core; that is,

$$Edv = (EDw \times EDh) + \left(\frac{EDh^2}{\sin(Am_t)}\right), \tag{2}$$

where

Am_i = maximum angle of repose for material t.

Up to three grades of fill material are available at each site with varying haulage distance and extraction cost. Type of material directly affects the volume of the embankment dams due to maximum slope requirements. Haul distances for the differing materials are also taken into consideration. The model output consists of total cost (EDc) relating to dam volume and haulage/excavation cost (Hc per cubic metre). Material type is only considered when the embankment dam option is selected; that is,

$$EDc = EDv \times Hc \tag{3}$$

The concrete gravity dam model at this stage is limited to the satisfaction of overturning criteria using the standard "middle third" rule, that is, the resultant force from gravitational and hydrostatic forces must pass through the "middle third" of the concrete dam base. An iterative procedure determines minimum base width dependent upon dam height. Dam volume (CDv) is calculated from the appropriate valley section and dam height (Dh).

- *Pressure Tunnels*: The length of the headrace and tailrace tunnels (PTl) is represented as a single continuous variable at each site within the limits of a 2-kilometer longitudinal valley section, that is, point of exit of the tailrace can be at any point within that 2-kilometer section. Potential head (Ph) is measured between dam crest and tunnel exit.
- *Powerhouse Location*: It is assumed that tunnel length is related to Powerhouse depth (Pd), that is, the deeper the powerhouse the less circuitous the tunnel route to satisfy minimum overburden requirements. Powerhouse depth is therefore represented as a continuous variable between differing upper and lower bounds at each site. Tunnel length is then adjusted by means of a factor (PTf) related to powerhouse depth; that is,

$$PTf = 1 + \left[k_i \times \frac{Pd_i}{(Pd_{\max} - Pd_{\min})}\right], \tag{4}$$

where

k = factor relating to overall topology of each site i and

$Pd_{\max}, Pd_{\min}$ = upper and lower bounds of Powerhouse depth.

The upper and lower bounds of the powerhouse depth can be considered to describe the constraints imposed by the topological characteristics of each site and a trade-off is required between powerhouse construction costs, tunnel costs, and potential head. Overall pressure tunnel costs (PTc) are therefore calculated from:

$$PTc = PTf \times PTl \times Utc, \tag{5}$$

where

Utc = Unit cost (per meter) of tunnel construction. This cost varies dependent upon mode of operation due to the accommodation of increased flow rates.

Powerhouse costs (Phc) are relative to depth, that is:

$$Phc = Phs \times \left(\frac{Pd_i}{(Pd_{\max} - Pd_{\min})}\right), \tag{6}$$

where

Phs = cost of surface powerhouse construction.

Variable parameters within the model are therefore: site, mode of operation, dam type, dam height, material type, period of generation, pressure tunnel length, and powerhouse depth. The model consists of several modules each returning overall cost of a major component from the simple mathematical representations of the factors described above. Pressure tunnel length, dam height, and powerhouse depth are common to all configurations, whereas period of generation is only included when peak power mode is selected and material type is only included upon selection of the embankment dam option as illustrated in Figure 2 and the sample chromosomal representations of Figure 3. Other parameters, the values of that differ between site but do not vary independently, are: Diversion tunnel length; spillway cost and a factor applied to the concrete gravity dam option related to ground stabilization requirements. These parameters are in addition to those already describing the dimensional characteristics of each site, that is, longitudinal and cross-sectional. The model is therefore an illustrative representation of a real-world system. The mathematical simulations of the various elements of the system provide data relating to tunnel length, dam volume, powerhouse depth, and energy output. The system is then costed using data supplied by Knight Piesold and Partners of Ashford, Kent, and a total capital cost is thus generated from those elements included in the analysis. This cost is then increased by 25% to take into account those elements not included such as spillways,

https://doi.org/10.1017/S0890060498125039 Published online by Cambridge University Press

SITE (1 to 5)
MD_1 MD_2 MD_3 MD_4 MD_5 (MD_i = Mode of operation for $SITE_i$)
DT_{11} DT_{12} DT_{21} DT_{22} DT_{31} DT_{32} DT_{41} DT_{42} etc (DT_{ij} = dam type for $SITE_i$, MD_j)
MT_{11} MT_{11} MT_{12} MT_{12} MT_{21} MT_{21} MT_{22} etc (MT_{ij} =material type for $SITE_i$, MD_j)
PD_{111} PD_{112} PD_{121} PD_{122} PD_{211} etc
TL_{111} TL_{112} TL_{121} etc
DH_{111} DH_{112} DH_{121} etc
GP_{111} GP_{112} etc

where:
MD = Mode of operation (1 = peak power supply; 2 = baseload supply)
DT = Dam type (1 = embankment; 2 = concrete gravity)
MT = Material type (only applicable to embankment dams)
PD = Powerhouse depth (relates to all sites and configurations)
TL = Tunnel length (relates to all sites and configurations)
DH= Dam height (relates to all sites and configurations)
GP = Generation period (only applicable to peak power generation mode)

Therefore, a chromosome defining site 2, baseload generation and concrete gravity dam would activate the following parameters:
SITE MD_2 DT_{22} PD_{222} TL_{222} DH_{222} TL_{222}
whereas site 4, peaking generation and embankment dam (rockfill) is represented by:
SITE MD_4 DT_{41} MT_{41} PD_{411} TL_{411} DH_{411} TL_{411} GP_{411}

Fig. 3. Chromosome representation.

generation sets, and transmission. All cost data are based upon average current costs for hydropower development in Central Africa. The Net Present Value (NPV_1) of the total Capital Cost (CC_t) plus an assumed annual operating cost is calculated over a 25-year period at an annual discount rate (m) of 15%. The energy output of the system is similarly discounted (NPV_2) over the same period and a unit cost of energy generation (UC) calculated; that is,

$$NPV_1 = \sum_{n}^{1} \left[\frac{CC_t}{(1+m)^n} \right] \tag{7}$$

$$UC = \left[\frac{NPV_1}{NPV_2} \right] \tag{8}$$

$$NPV_2 = \sum_{n}^{1} \left[\frac{MWh}{(1+m)^n} \right]. \tag{9}$$

The overall objective is to minimize unit cost, UC.

The range of variation of the eight variables are shown in Table 1.

3. THE STRUCTURED GENETIC ALGORITHM

Previous research concerning the manipulation of interrelated discrete and continuous design variables using evolutionary techniques is evident within the engineering structural optimization community. For example, the development of parallel strategies for large mixed-integer structural problems relating to multimember transmission and radio towers (Cai & Thierauf, 1996) utilizing evolution strategies (Rechenburg, 1984). This large mixed-integer problem is decomposed into two subproblems consisting of either discrete or continuous variables and appropriate information exchange between the two domains is introduced. This two-stage approach relating to the geometrical layout and the sizing of truss members has also been investigated utilizing genetic algorithms with applications to trussed-beam roofs and the design of cable-stayed bridges (Jenkins, 1991). Hybrid strategies involving a GA search for optimal discrete geometries and a logic-based approach to determine optimum member sizing were introduced by Koumousis and Georgiou (1994), whereas Hajela et al. (1992) adopt a two-stage, constraint relaxation approach. Later work maintains a two-stage approach (1995) involving an initial topological design for kinematic stability requirements followed by an optimization for response constraints. Leite and Topping have proposed an alternative approach first utilising a simulated annealing algorithm to search the space of possible topological solutions whilst avoiding the generation and development of non-feasible geometries (Leite, 1996).

https://doi.org/10.1017/S0890060498125039 Published online by Cambridge University Press

Table 1. *Variable parameter ranges*

Site	1	2	3	4	5
Modes of operation	2	2	2	2	2
Dam types	2	2	2	2	2
Tunnel length	0–2000 m	0–2000 m	0–2000 m	0–2000 m	0–2000 m
Powerhouse depth	0–100 m	0–50 m	0–150 m	0–125 m	0–100 m
Dam height	0.8–1 Vd*	0.7–0.95 Vd	0.6–0.85 Vd	0.8–1 Vd	0.8–1 Vd
Generation period	4–10 hours	5–10 hours	5–12 hours	4–9 hours	5–10 hours
Material type	3	3	3	3	3

*Vd = Valley depth.

These techniques are primarily sequential in nature, first addressing the discrete space before subsequently searching the continuous space relevant to the identified optimal configuration. The problem domains largely concern routine design utilizing well-defined structural models to provide a measure of relative fitness for each candidate solution. In addition, the continuous design spaces tend to relate to member sizing and, although some difference in the upper and lower bounds may exist, the diversity of the different continuous sets of Figure 1 is not evidenced. To achieve the desired concurrent search of the discrete and continuous spaces, a number of design grammar representations have therefore been investigated (e.g., Antonisse, 1991; Dasgupta, 1992b; Koza, 1992, 1994; Gero, 1994). The fixed-length structure and definition of the hydropower design hierarchy does not require the flexibility and variable length representations of the genetic programming paradigm, thus previous experience within the Center initially led to the incorporation of the Structured Genetic Algorithm (stGA) (Dasgupta & MacGregor, 1992*a*, 1992*b*, 1994) as an appropriate starting point for the development of a suitable global search paradigm. The preliminary design models introduced in the previous section have therefore been integrated with an stGA. The stGA allows parameters defining the characteristics of each site to be encoded in the same chromosome string. Controlling parameters, relating to site, mode of operation, and dam type then "switch on" the relevant parameters within the chromosome string, which relate to the appropriate configuration. The "live" parameters are subsequently passed to the design models. Other parameters remain dormant within the string unless activated in a later population as a result of crossover or random mutation. This process eliminates the possibility of nonfeasible parameter combinations being passed to the mathematical model. The chromosome structure for the hydropower system is shown in Figure 3. The stGA is therefore defining the hierarchical structure consisting of both discrete variables (i.e., site, mode of operation, and dam and material types), which henceforth shall be referred to as set A and continuous variables (dam height, tunnel length, powerhouse depth, flow-rate, etc.), set B, as shown in Figure 1. Traditionally, such a design hierarchy would be decomposed to reduce dimensionality and allow search to take place within an extremely limited, but manageable, design space. The objective of the initial research is to achieve concurrent processing of the hierarchical levels by utilizing the nonlinear search capabilities of the stGA.

The initial choice of the stGA was based upon the results from other research within the Plymouth Engineering Design Centre relating to the algorithm's utilization for the optimization of finite impulse response (FIR) digital filters (Wade et al., 1994; Roberts & Wade, 1994). This work investigated the utility of the stGA for the concurrent manipulation of discrete parameters relating to primitive filter type, delay information, and coefficient addresses within a three-layer representation of the FIR filters. Further work in the area relates to the use of the stGA for geometry-type representation in the application of genetic algorithms to cooling hole geometry design of gas turbine blades (Roy et al., 1996). Some significant success has been achieved in these domains although stGA utilization has so far been restricted to problems requiring few levels of representation.

3.1. Binary mapping

The initial stGA implementation utilized a binary representation of the design parameters (Parmee, 1995). However, this representation immediately presented problems in terms of the probabilities of mutation and crossover being directly related to the length of the binary representation of each parameter. The combination of simple binary switches (i.e., dam type—embankment or mass concrete), multiple-option discrete parameters (i.e., site), and the continuous variables (e.g., tunnel length) results in differing orders of binary representation. The simple design switches can be represented by a single binary digit, whereas discrete multiple options must be represented by an appropriate binary order, the decoded value of which must then be suitably scaled to provide the necessary integers. The order of the binary representation of the continuous variables depends upon the required parameter resolution but is unlikely to be less than five and in the research described here a six-digit binary representation has been used. We therefore have a chromosome string consisting of binary encodings of different order for each parameter type. The probability of parameter

https://doi.org/10.1017/S0890060498125039 Published online by Cambridge University Press

disruption from either random mutation or crossover siting is therefore significantly greater for higher order encodings than for the single binary digit representations.

This problem is addressed by introducing a weighting to each binary digit: that is,

$$\mathbf{W}_b = ((\mathbf{O}_{\max} - \mathbf{O}_{\min}) + 1)/((\mathbf{O}_b - \mathbf{O}_{\min})) + 1),$$

where

$\mathbf{W}_b$ = the weighting applied to the binary digit,

$\mathbf{O}_b$ = the order of the binary representation of the parameter to which the digit belongs,

$\mathbf{O}_{min}$ = the order of the minimum length binary representation, and

$\mathbf{O}_{max}$ = the order of the maximum length binary representation.

(where "order" defines the number of binary digits describing a parameter).

Having established a weighted population of binary digits Roulette Wheel selection (Goldberg, 1989) can be implemented to select those to be mutated, that is, the larger the weighted value of the gene the greater the probability of selection for mutation. Another factor to be taken into consideration, however, is the binary representation of the multiple-option discrete parameters (i.e., site and material type) and problems associated with the scaling of the decoded values to ensure that each integer describing site/material type has the same probability of selection.

Chromosome length is also a cause for concern. The simple system is described by eight variables: Site, Mode of Operation, Dam Type, Dam Height, Powerhouse Depth, Generation Period, Tunnel Length, and Material Type. However, these parameters must be represented, where appropriate, for each of the branches of the design hierarchy. This results in a chromosome string containing 96 variable parameters. Even by restricting continuous parameter resolution to a five-digit binary representation, a chromosome length of over 400 genes (binary digits) is required to maintain this basic system representation. It is envisaged that a realistic preliminary design model of the system will involve several more layers of discrete/continuous variable parameters with a corresponding significant increase in chromosome length.

All of the above factors indicate that a binary mapping approach may not be the most appropriate. With this in mind, a real-number representation has also been implemented.

3.2. Real number representation

An immediate benefit of real-number representation is a reduction in chromosome length, a significant reduction in stGA coding complexity, and the elimination of problems associated with mutation and crossover site probability plus the elimination of decoding/scaling problems. Single-point crossover is carried out at randomly selected sites located between parameters and mutation is a simple perturbation of randomly selected parameters between their upper and lower bounds (Davis, 1991).

3.3. Preliminary results

Initial runs of the binary stGA and the real-number stGA show a rapid convergence upon relatively high-performance solutions. Knowledge of the discontinuous nature (in terms of discreteness) of the search space caused by the discrete variable parameters introduces a degree of caution when considering these solutions as being globally optimal. Investigation of the number of times the algorithm has visited each of the discrete branches (20 in number) of the hierarchy during preliminary testing reveals that little search across the system has taken place. This is illustrated in Figure 4, which shows the average results from 50 runs of a basic real-number stGA with single-point crossover, 0.02 mutation probability, roulette wheel selection, and an elitist strategy (i.e., the best member of each population is always reproduced). Each line represents a discrete path through the hierarchy while the number of calls relates to the number of times the paths have been visited over 50 generations with a population size of 100 chromosomes.

With the basic representation of Figure 2 we could, of course, decompose the hierarchy to some extent and initiate independent GA searches within each branch. The overall objective, however, is to develop a search methodology that allows concurrent processing of far more complex hierarchies. Such systems will likely involve discrete branching

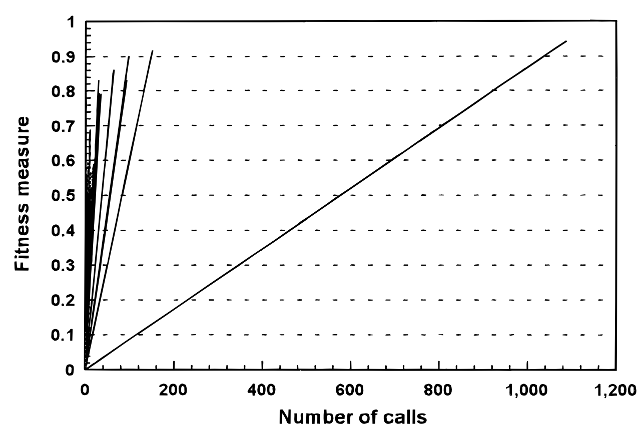

Fig. 4. Basic stGA approach. Each line represents one of the possible 20 paths created by the discrete parameter set A. The number of calls relates to the number of times each discrete set has been passed to the mathematical model. The graph therefore illustrates the degree of search diversity/exploration across the multilevel representation. Fitness values shown here and in Figure 5 are relative to best fitness achieved in the GAANT application of Figure 11. This fitness relates directly to the minimization of unit cost of power output of the system as described in Section 2.1.

below levels described by continuous variables. The introduction of independent GA search along selected branches will rapidly become nonfeasible unless we return to the traditional approach of decomposition and subsequent search of limited regions.

The development of suitable mutation regimes to ensure that the discrete elements of the hierarchy have a sufficient probability of mutation has been necessary to allow an acceptable degree of investigation of the lower-level variables. This has been achieved by assigning independent variable mutation probabilities to the high-level discrete variables (Set A). A uniform mutation probability is applied to the remaining continuous parameters (Set B). Typically a mutation probability of 0.2 that has been applied to set A, while a mutation rate of 0.02 has been uniformly applied to the remaining continuous parameters (set B) (Parmee, 1995). Elitism is not implemented. The effect of this "variable mutation" approach is to promote search across the hierarchy with search effort significantly increased along each of the discrete paths and improved solutions obtained from the majority of the 20 possible system configurations. The optimal solution of Figure 4 is not equalled, however, and so a hybrid approach has been implemented. In this case, varying mutation (no elitism) is implemented for the first 25 generations at which point uniform mutation plus elitism takes over for a further 25 generations. The results from the hybrid approach are shown in Figure 5.

The results shown represent averages from 50 runs of the stGA, each run commencing from a different point. Population size equals 100 chromosomes and single-point crossover with a probability of 0.6 with roulette wheel selection was implemented.

3.4. Discussion of preliminary results

It is apparent from the graph of Figure 4 that although a high-performance solution has been achieved, little search across the design hierarchy has taken place. On average

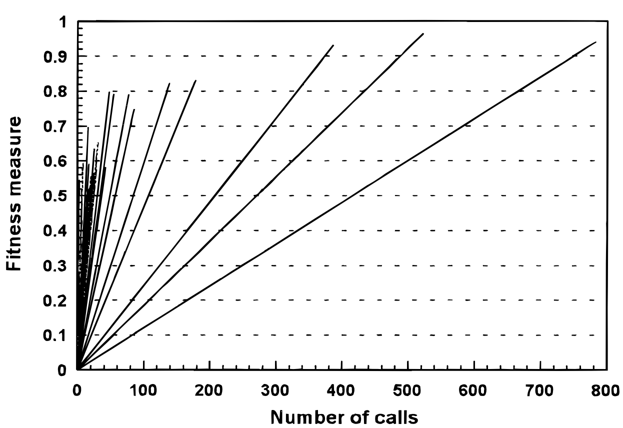

Fig. 5. Hybrid mutation strategy approach.

10% of the paths had not been visited at all during each run of the uniform mutation + elitism stGA. Confidence in this "best" solution actually representing the global optimum must therefore be low especially when one considers the discontinuous nature (in terms of discreteness) of the search space. We must also consider the requirements of the engineer. It is suggested that the identification of a single optimum solution would not be considered satisfactory by the designer during these early stages of the design process. More information would be required concerning the performance of various configurations within other sites to instill confidence and to allow other considerations to be taken into account. It is likely that various objectives exist in addition to the minimization of unit energy cost. Criteria, such as site and design preferences, construction aspects, plant and personnel availability and seasonal site conditions, may have to be considered. The stGA tool must therefore be utilized in a decision support manner and in this respect must provide multiple high-performance solutions from disparate regions of the hierarchy. Another factor here relates to confidence in available data—how valid is a single optimum solution where uncertainty exists concerning the validity of currently available site data?

The "variable mutation" approach results in a diverse search process where a much higher proportion of the design hierarchy has been visited and a number of high-performance solutions have been located but the optimal solution of Figure 4 is not identified. The introduction of the hybrid approach achieves a diverse search across the hierarchy, while also resulting in the identification of better performing solutions than those of Figure 4. Although the results show that it is possible to establish a diverse search across a complex design hierarchy using an stGA and appropriate mutation regimes, the performance of this approach for a more complex structure is not known. It is apparent that an alternative approach with the basic design hierarchy presented here would be to carry out 20 independent searches along each of the discrete paths. Although time-consuming, this would provide optimal solutions across the entire structure. However, it is envisaged that to develop a meaningful system model, a more complex hierarchy will be required involving an increased number of layers and significantly more variable parameters. This will eliminate the opportunity for decomposition and independent search.

These initial results show that the appropriate utilization of the Structured Genetic Algorithm can:

- provide a single high-performance solution from the design hierarchy described in the text within 50 generations.
- provide multiple high-performance solutions from the same hierarchy within 50 generations.

In all cases a real-number representation of the parameters appears to be adequate and is far less complex than a binary mapping.

https://doi.org/10.1017/S0890060498125039 Published online by Cambridge University Press

3.5. Deficiencies of the STGA and mutation strategies

For the strategies so far discussed, crossover can be potentially disruptive to the desired formation of high-performance parameter sets particularly in the earlier generations. A high diversity of system configuration described by the three main discrete parameters (set A) in early generations leads to a high level of crossover of the continuous variables (set B) between differing configurations. This leaves little opportunity for the evolution of better continuous variable sets for any particular discrete configuration. Crossover merely perturbs either A or B depending upon the location of the crossover site with little relevant improvement in the continuous variables when related to a particular site, mode of operation, or dam type. The situation improves as a degree of convergence becomes apparent. The probability of crossover between "like" configurations then increases and there is therefore an increased opportunity for an overall relevant improvement in parameter set B. This represents a serious problem in that the opportunities for premature convergence are prevalent. The initial hydropower hierarchy is simple and a significant increase in complexity will result from further development. This is bound to involve additional sets of discrete parameters at lower levels of the hierarchy, which will likely exacerbate the problem and increase the probability of premature convergence.

Another area for concern is the inherent parameter redundancy within each of the stGA's parameter strings. To prevent the generation of nonfeasible (lethal) parameter combinations, it is necessary to include parameter representations for each of the discrete design options. This requires a string containing 96 parameters of which only 8 (i.e., those relevant to the discrete configuration) will be passed to the design model as shown in Figure 2. A high proportion of crossover and mutation could therefore be ineffective unless restricted to those areas of the string currently active. Added to this are the problems associated with the excessive number of parameters required in an stGA representation, as the hierarchy is developed and problem complexity increases.

4. SIMPLIFYING THE PARAMETER REPRESENTATION

The ideal representation would be a strategy that ensures the avoidance of lethal parameter sets, while allowing the information exchange evident during traditional crossover. Crossover could be restricted to parameter strings that are identical in terms of the discrete parameter set (i.e., set A). This would allow an exchange of information relevant to a particular design configuration and the subsequent evolution of that configuration in terms of the continuous variables (set B), while also allowing a straightforward eight-parameter chromosome representation.

However, this crossing of "like" configurations in terms of the discrete parameters does not allow their perturbation and their subsequent improvement. This problem could be addressed by introducing the variable mutation probability approach described earlier, that is, introduce a high mutation probability in set A and a lower probability to set B. The overall strategy would therefore involve two individual search agents that are operating autonomously, that is, a simple hill climber manipulating the discrete set A and a genetic algorithm manipulating the continuous set B. Communication, however, is inherent and exists at a lower level between the individual chromosomes of the continuous set B via information exchange during crossover and at a higher level between the two search agents via the evaluation of each string and subsequent selection using roulette wheel (Goldberg, 1989). An improvement upon this approach can be achieved by introducing lower-level information exchange between the discrete sets, A, of each generation by introducing elements of an ant colony metaphor for the manipulation of the discrete set. This introduction forms the basis of the GAANT algorithm.

4.1. The ant colony algorithm

The ant colony algorithm (ANT) is analogous to the foraging strategies of ant colonies (Colorni et al., 1991, 1992; Bilchev & Parmee, 1995, 1996). The technique relies upon multiagent cooperation to concentrate the search in those areas where "food" is abundant. The amount of food available in this instance being related to the relative fitness of the design solution as defined by the engineering system model. The initial ant resource is initialized by generating uniformly random starting directions from a virtual "nest," as shown in Figure 6. The parameter sets defining these points are sent to the design model and the allocation of further ant resource along each trail is proportional to the relative fitness of the returned values from the design model (Fig. 7). A search radius, *R*, is defined, which determines the maximum extent of the subspace to be considered in

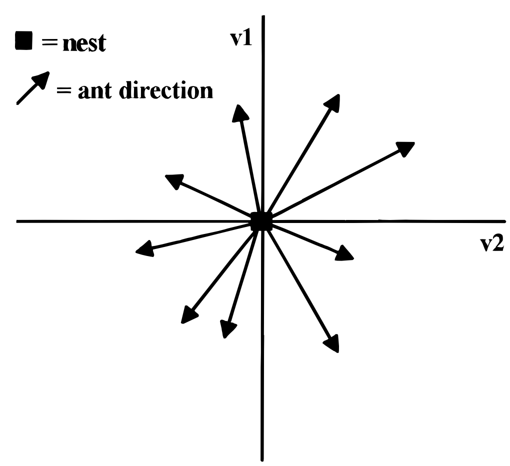

Fig. 6. Random selection of initial search directions.

https://doi.org/10.1017/S0890060498125039 Published online by Cambridge University Press

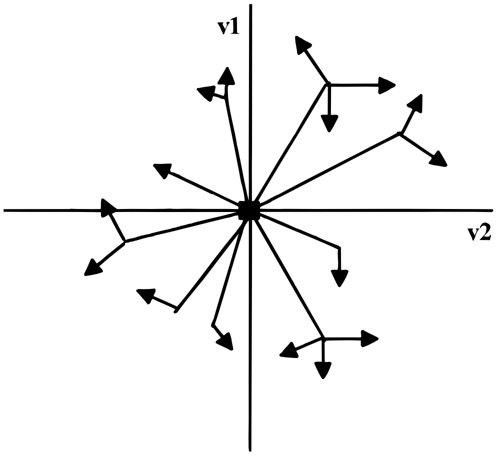

Fig. 7. Fitness proportionate distribution of nest resource.

each cycle. Directions are then randomly selected from each initial point and further ant resource follows these directions in steps not greater than the radius *R*. Each ant thus defines a new point in the design space. Poorly performing trails are evaporated after some preset number of cycles (Fig. 8) and the ant resource from these trails is reallocated around the better performing trails. In its simplest form, ANT could be considered to represent a multi-start hillclimber with communication between each element. It is this communication aspect that prevents rapid convergence upon local optima and provides the basis for the algorithm's search capabilities, which result in the identification of high-performance solutions.

4.2. Gaant

Two particular Ant Colony "operators" are of interest:

- **fitness proportionate distribution**: similar to fitness proportionate reproduction—in this case, the number of software "ants" distributed down each "trail" is proportionate to the relative strength (i.e., fitness) of that trail. Each discrete path of the hierarchy represents a possible "trail."
- **evaporation**: if the strength of a particular "trail" does not improve over a preset number of iterations, then that "trail" is evaporated and the released "ant" resource is redistributed around the better "trails."

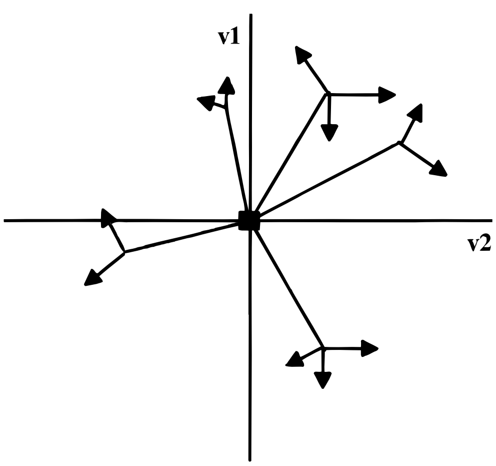

Fig. 8. Evaporation of poor performance trails.

The concepts supporting these two operations have been adapted slightly and integrated with the manipulation of the discrete parameter set. The flow chart of Figure 9 illustrates this integration. To further explain the process, the values of set A are randomly selected in generation 1 and combined with a randomly selected population of set B. The initial population of discrete parameters survives for a preset number of generations (n), while the as-

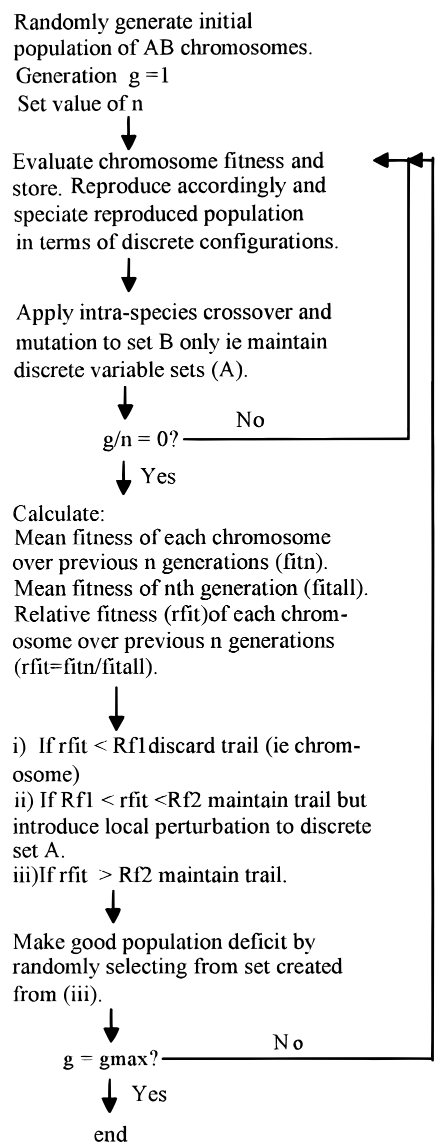

Fig. 9. GAANT flow chart.

https://doi.org/10.1017/S0890060498125039 Published online by Cambridge University Press

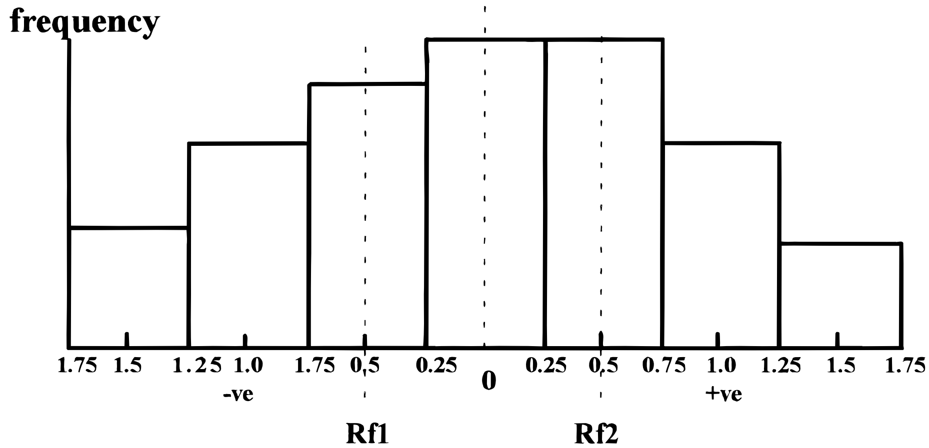

Fig. 10. Scaling of rfit and introduction of Rf thresholds.

sociated continuous parameters are manipulated by a simple GA. A combination of like parameter sets during crossover and reproduction is controlled by means of the speciation of each chromosome in terms of like configurations of the discrete parameters. Crossover then only occurs between members of the same species. Evolution of set B continues over each generation, whereas improvements in set A are achieved as follows: The average fitness of *each* chromosome is calculated over n generations then compared to the average fitness of the chromosomes of the nth generation. Evaporation, duplication and perturbation of the discrete parameter sets is then established in accordance with their relative fitness ($rfit$), which is represented in terms of their average fitness ($fitn$) over n generations and the average fitness of members of the nth generation ($fitall$), that is, $rfit = fitn/fitall$. This allows the following communication:

- low-level communication between the chromosomes of set B resulting in the evolution of the continuous parameters within the bounds imposed by their discrete system configuration;
- low-level communication every nth generation between the chromosome sets representing the discrete parameters, which results in their gradual improvement; and
- high-level communication between the two agents in the form of relative fitness of an entire string over n generations.

Evaporation, duplication, and perturbation are controlled at the nth generations by introducing two thresholds Rf_1 and Rf_2 (Fig. 10). If $rfit$ is less than Rf_1, then the trail is evaporated, that is, the chromosome is not reproduced. If $rfit$ is greater than Rf_2, then the trail is maintained (i.e., the chromosome is reproduced) and further resource is allocated from the evaporated trails (i.e., the population deficit created by trail evaporation is made good by randomly selecting chromosomes from those with a fitness higher than Rf_2). Finally, if $rfit$ lies between Rf_1 and Rf_2, the discrete parameters are randomly perturbed to create a new trail.

5. RESULTS AND DISCUSSION

Initial results displayed in the same format as Figures 4 and 5 are shown in Figure 11 and should be compared to the hybrid stGA results of Figure 5. All results are based upon a population size of 100 chromosomes with $n = 5$ generations; $Rf_1 = -0.75$ and $Rf_2 = 0.75$ (in terms of standard deviations from the mean). Table 2 shows a more detailed comparison between the dual mutation regime approach and GAANT. Initial results illustrate the change in performance related to the overall number of calls to the fitness function and the setting of n. Fitness is shown relative to that of the initial GAANT implementation with $n = 5$ and number of calls = 2500. The results are averaged from 100 runs of the various algorithms. Standard deviation (SD) of the number of calls along the best path and of the fitness along that path over the 100 runs is shown. Table 2 also shows the number of times individual runs have failed to visit paths and the

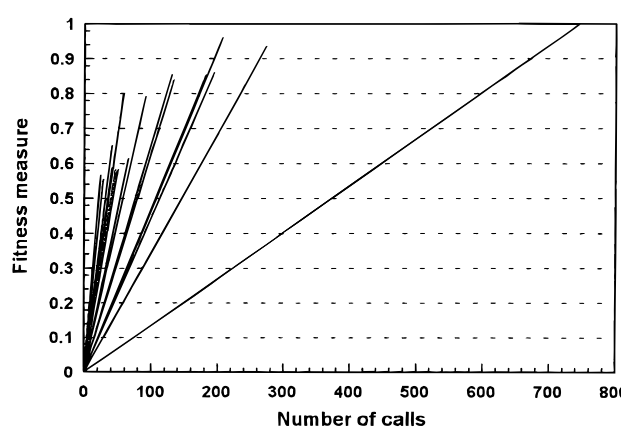

Fig. 11. GAANT implementation.

https://doi.org/10.1017/S0890060498125039 Published online by Cambridge University Press

Table 2. *Comparison of stGA and GAANT approaches*

	Hybrid stGA Approach		GAANT Implementation			
			$n = 5$		$n = 10$	
	2500 Calls	5000 Calls	2500 Calls	5000 Calls	5000 Calls	7500 Calls
Max solution	0.95	0.97	1	1.01	1	1.01
No. of solutions > 0.9	3	3	3	3	3	3
No. of solutions > 0.8	5	5	7	9	9	10
No. of calls along best path	611	1727	711	1277	1287	2545
SD of No. of calls (Best path)	580	1706	207	460	413	738
SD of Fitness (Best path)	0.21	0.26	0.05	0.05	0.05	0.04
No. of missed paths	93	78	0	0	0	0
Max No. of misses of any path	20	26	0	0	0	0

worst case showing the max number of "misses" related to a particular path over the 100 runs.

It is apparent from Figures 5 and 11 and from the number of solutions greater than 0.8 in Table 2 that the GAANT approach can provide improved performance in terms of maximum fitness across a larger number of paths than the dual mutation/stGA approach. The standard deviation (*SD*) of the calls and of the fitness along the best path has been calculated to give an indication of the robustness of the two approaches. The very high *SD* of the hybrid stGA approach prompted further investigation of the results that showed a far from normal distribution, which indicates a lack of robustness. This is further supported by the number of paths that are not visited over the 100 generations of the test runs. The GAANT approach, however, shows far greater stability from the standard deviations and the complete cover of the 20 discrete paths. Although both of the techniques finally converge upon the same configuration (i.e., best discrete path), the GAANT solution from that path is significantly better than that found by the stGA.

Other aspects currently under investigation include the values of Rf_1 and Rf_2 and the redistribution and evaporation strategies of the ant colony manipulation of the discrete parameters. Extensive experimentation is planned based upon a number of test hierarchies of varying complexity.

6. FURTHER WORK

The research has led to the application of GAANT to more complex structures relating to the optimization of thermal power system configuration. In this case the feed-heater layout of the steam cycle of nuclear power plant has been addressed (Chen & Parmee, 1997). Discrete variables relate to steam tapping points and the number and configuration of the feed heaters. Continuous variables include control settings and system component dimensions. Preliminary results show a significant reduction in design lead time in addition to significant increases in predicted power output (circa 0.3%). The integration of GAANT and GAANT variants has resulted in totally machine-based design processing that has replaced an engineer/machine based process. Overall design time has been reduced by approximately 75%. This further work also indicates the generic applicability of the strategies. Research is now continuing by addressing more fundamental issues relating to the structure of the GAANT strategy, while also improving the generic aspects through the introduction of the technique and variations to a wide range of engineering problem domains.

The concepts behind the GAANT representation have also been applied to the manipulation of variable-length multi-level mathematical function representations. The objective here has been to improve the calibration of preliminary design models to empiric data or to results from a more in-depth analysis (FEA or CFD). This is achieved by identifying those areas of coding where insufficient knowledge or the requirement of keeping computational expense to a minimum has resulted in unavoidable function approximation. A contributing factor may be the inclusion of empirically derived coefficients (i.e., discharge, drag, etc.). The objective is to evolve improved coding within these areas to achieve a better calibration with existing empiric data or results generated from a more in-depth, computationally expensive analysis. If this is possible, then the element of associated risk would be correspondingly lessened, while rapid design iteration can still be achieved utilizing these simple, but more representative models. Initial research indicated that adaptive techniques and genetic programming (GP) (Koza, 1992,1994) when utilised for system identification can achieve these objectives to a limited extent. It soon became apparent that the problems associated with the crossover of continuous coefficients between differing discrete functional structures causes similar problems to those identified in Section 3.5 relating to the successful crossover of useful information within the fixed-length design hierarchies. As previously noted, the exchange of information from continuous design spaces to unrelated discrete design configurations does not promote the formation of high-performance variable parameter combinations. This semantic disruption problem has been recognized within the

https://doi.org/10.1017/S0890060498125039 Published online by Cambridge University Press

GP community (Iba et al., 1996). The success of the GAANT strategy has stimulated experimentation to assess the benefits of a similar approach to the manipulation of the variable-length hierarchies describing the mathematical functions and engineering representations. Some significant preliminary results have been achieved (Watson & Parmee, 1997). The indication is that further development of the GAANT concepts and their integration with GP techniques for the manipulation of variable-length representations has some considerable potential.

7. CONCLUSION

A comparison of the GAANT approach with the previous work involving the stGA shows that the strategy can find better solutions for a greater number of discrete configurations. Diversity of search is thus maintained and with fewer calls to the model in some cases (Fig. 11). The GAANT technique is also exhibiting a much greater degree of robustness than the stGA approach, which is of significant importance in terms of generic applicability of the strategy.

Of equal if not more significance during these early stages of development is that improvement has been achieved using a basic chromosomal representation of the eight variable parameters that describe the system. The extensive chromosomes of the stGA are no longer required, thus problems concerning redundancy and complexity of implementation can be avoided. This allows us to further develop the system hierarchy to achieve a more realistic preliminary design model involving further levels comprising of discrete and continuous variable sets. The complexity of the stGA representation would have seriously restricted such a development. The research has therefore established an alternative strategy that can maintain sufficient search diversity and improve results, while also allowing the continuation of overall system development. In this respect it can be considered to represent a significant contribution to the successful integration of adaptive search with whole-system design.

The results represent an initial basic implementation of the GAANT algorithm. The implementation of an experimental strategy to investigate the main characteristics of the dual-agent approach is likely to result in further improvements in performance. The generic aspects of the work must be considered. The GAANT approach is proving to be applicable to a range of whole-system design problems and its potential could therefore be considerable within the early stages of design. Current work investigating the integration of these strategies with thermal system plant configuration (Section 6) is indicating that application of the algorithm is generic. Significant improvement in predicted power output and design lead time is being achieved by manipulating the whole plant models utilizing basic variations of the GAANT algorithm. The indication is that the current dual agent philosophy will provide a basis for further development. Initial application to the thermal system problem would not have been possible without the simplification of the chromosomal representation offered by GAANT.

With regard to the hydropower system application, significant improvements and additions to the model are required to provide an acceptable design tool. Areas that must be addressed include: hydrological and geological considerations; river diversion; hydraulic characteristics; storage; spillway requirements; turbomachinery, generator sets, and transmission; and more definitive modelling of the water-retaining structures. However, it must be understood that the resulting design tool is for use during the preliminary stages of the design process (i.e., prefeasibility/bid stage). The resulting mathematical representations need not be overly complex for reasons explained previously. In its present form, the GA-driven model may provide some rapid insight relating to site potential in the early days of the project, but it must be considered illustrative only in its present form.

The computational expense associated with the GAANT manipulation of the hydropower model is low. Runs involving 2500 calls to the evaluation function are completed, on average, within a 3-min period on a two-processor, SPARC 10 workstation. Even considering the basic hierarchy used here, the GAANT processing of the problem would far outperform a human designer using either computer-supported heuristic techniques or more deterministic gradient-based optimization tools to independently search each of the continuous variable domains related to the 20 possible discrete design configurations of this illustrative hydropower hierarchy. It could be argued that, in this case, an experienced engineer could eliminate many of the discrete paths with high confidence. However, the purpose of this initial work has been to establish the strategy using a relatively simplistic test case. The results show that concurrent processing of the different levels of the hierarchy can be achieved and that diverse high-performance solutions can be identified. The rapid, exploratory aspects of the strategy have therefore been illustrated. These aspects have been further proven from application in the thermal system domain introduced in Section 6.

The intention has been to develop a strategy that supports exploration during the early stages of the design of a large-scale system, where poorly defined initial data and uncertainty are major factors. The research has established such a strategy and therefore offers an exploratory, decision support environment that could be of major benefit to such design domains. It is suggested that large-scale civil engineering certainly falls into this category due to the interface with natural systems often remote from the design office. However, recent research involving preliminary airframe design is illustrating similar characteristics in this case in terms of long-range forecasting of military airframe requirements. Work is now proceeding in this area and the GAANT approach will likely play a fundamental role. With respect to the far bettter defined thermal system design problem, GAANT has offered the means to allow extensive experimentation using the system simulation software, which has resulted in the discovery of novel feed heater layouts. Such experimentation/exploration would not have taken place if extensive man–machine involvement had been required.

A criticism of the GAANT approach could relate to the complexity of the disjoint continuous spaces. Currently little more than a simple genetic algorithm (SGA) is being utilized to search these domains and future work must address the utilization of higher performance GAs such as the CHC (Eshelman, 1991) or Fast Breeder (Muhlenbeim & Schlierkamp-Voosen, 1993) to better handle problems associated with deceptive fitness landscapes, high modality, and heavy constraint. This is being addressed in current PEDC research in this area. However, the utilization of the SGA is justified in this preliminary investigation and initial establishment of the strategy. Consideration must be given to the validity of the design model, however. Extensive search to locate an elusive peak or trough could be considered wasteful if the probability of that solution, being erroneous under more detailed analysis, is high.

It should be stressed that the utilization of the exploratory capabilities of GAANT, when uncertainty is prevalent, should involve a high degree of engineer interaction. In this manner, designer intuition and problem-specific knowledge can be applied to locally optimal solutions that well-reasoned decisions concerning future design direction can be made. Again, confidence factors relating to model resolution must be considered.

There is now little reason to believe that further modification plus the introduction of other, perhaps more appropriate, adaptive search paradigms will not result in satisfactory diverse search across more complex hierarchies. The exploratory aspects of the approach are now well-founded and it has been shown that these will allow the engineer to explore regions of the design space that are currently beyond reach within the constraints imposed relating to budget and available design time.

ACKNOWLEDGMENTS

The Plymouth Engineering Design Centre is one of six EPSRC Engineering Design Centres that have been established in recent years at various U.K. academic institutions. Knight Piesold and Partners, an international civil and environmental consultancy based in Ashford, Kent has provided support in the form of discussion, advice, and cost data for the research described in the paper. The authors thank Knight Piesold and EPSRC for their support of our research activities.

REFERENCES

Antonisse, HJ. (1991). A grammar based genetic algorithm. In *Foundations of Genetic Algorithms* (Rawlins G.J.E., Ed.), pp. 193–204. Morgan Kaufman, Los Altos, California.

Bilchev, G., & Parmee, I.C. (1995). The ant colony metaphor for searching continuous design spaces. *Proc. AISB Workshop on Evolutionary Computing; Lecture Notes in Computer Science 993*, pp. 25–39. Springer-Verlag, London.

Bilchev, G., & Parmee, I.C. (1996). Constrained optimisation with an ant colony search model. *Proc. 2nd Int. Conf. on Adaptive Computing in Eng. Design and Control*, 145–151.

Cai, J., & Thierauf, G. (1996). Structural optimisation of a steel transmission tower using parallel evolution strategy. *Proc. Adaptive Computing in Eng. Design and Control '96*. 18–25.

Chen, K., Parmee, I.C., & Gane, C.R. (1997). Dual mutation strategies for mixed-integer optimisation in power station design. *Proc. IEEE Int. Conf. Evolutionary Computation* 385–390.

Colorni, A., Dorigo, M., & Manniezzo, V. (1991). Distributed optimisation by ant colonies. *Proc. Eur. Conf. Artificial Life*, 134–142.

Colorni, A., Dorigo, M., & Maniezzo, V. (1992). An investigation of some properties of the ant algorithm. *Proc. PPSN '92*, pp. 509–520. Elsevier Publishing, New York.

Dasgupta, D., & MacGregor, D. (1992*a*). Designing application-specific neural networks using the structured genetic algorithm. *Proc. Int. Workshop on Combination of Genetic Algorithms and Neural Networks*, pp. 87–96. IEEE Computer Society Press.

Dasgupta, D., & MacGregor, D. (1992*b*). Nonstationary function optimisation using the structured genetic algorithm. *Proc. Second Int. Conf. Parallel Problem Solving From Nature (PPSN)*, pp. 145–154. Springer Verlag, Brussels.

Dasgupta, D., & MacGregor, D. (1994). A more biologically motivated genetic algorithm: The model and some results. *Cybern. Sys.: An Int. J. 25(3)*, 447–469.

Davis, L. (1991). *Handbook of genetic algorithms*. Van Nostrand Reinhold, New York.

Eshelman, L.J. (1991). The CHC adaptive search algorithm: How to have safe search when engaging in nontraditional genetic recombination. In *Foundations of Genetic Algorithms and Classifier Systems*, (Rawlins, G.J.E., Ed), pp. 265–283. Morgan Kaufmann, San Mateo, California.

Gero, J.S. (1994). Evolutionary learning of novel grammars for design improvement. *AIEDAM 8*, 83–94.

Goldberg, D.E. (1989). *Genetic algorithms in search, optimisation & machine learning*. Addison-Wesley Publishing Co., Reading, Massachusetts.

Hajela, P., Lee, E., & Lin C.Y. (1992). Genetic search strategies in multicriterion optimal design. *Structural Optimisation 4*, 99–107.

Hajela, P., & Lee E. (1995). Genetic algorithms in truss topological optimisation. *J. Solids Structures 32(22)*, 3341–3357.

Holland, J. (1975). *Adaptation in natural and artificial systems*. The University of Michigan Press, Ann Arbor.

Iba, H., Degaris, H., & Sato, T. (1996). A numerical approach to genetic programming for systems identification. *J. Evolutionary Computation 3(4)*, 417–452.

Jenkins, W.M. (1991). Towards structural optimisation via the genetic algorithm. *Computers and Structures 40(5)*, 1321–1327.

Koumousis, V.K., & Georgiou P.G. (1994). Genetic algorithms in discrete optimisation of steel truss roofs. *ASCE J. Computing in Civil Eng. 8(3)*.

Koza, J. (1992). *Genetic programming*. MIT Press, Cambridge.

Koza, J. (1994). *Genetic programming II*. MIT Press, Cambridge.

Leite, J.P.B. (1996). Parallel adaptive search techniques for structural optimisation. PhD Thesis. Dept. of Mechanical and Chemical Engineering, Heriot-Watt University, Edinburgh, UK.

Muhlenbeim, H., & Schlierkamp-Voosen, D. (1993). Predictive models for the breeder genetic algorithm. *Evolutionary Computation 1(1)*, 25–49.

Parmee, I.C. (1995). Diverse evolutionary search for preliminary whole system design. *Proc. 4th Int. Conf. AI in Civil and Structural Eng*, 199–204.

Parmee, I.C. (1996). The development of a dual-agent strategy for efficient search across whole system engineering design hierarchies. *Lecture Notes in Computer Science 1141*, Springer-Verlag, ISBN 3 540 61723 X, pp. 523–532.

Parmee, I.C., Gane, C., Donne, M., & Chen, K. (1996). Genetic strategies for the design and optimal operation of thermal systems. *Proc. Fourth Eur. Congress on Intelligent Techniques and Soft Computing*, 419–423.

Rechenburg, I. (1984). The evolution strategy: A mathematical model of Darwinian evolution. In *Synergetics: From Microscopic to Macroscopic Order*, (Frehlend E., Ed.), Vol. 22, pp. 122–132. Springer Series in Synergetics, New York.

Roberts, A., & Wade, G. (1994). Optimisation of finite wordlength filters using a genetic algorithm. *Proc. 2nd Int. Conf. Adaptive Computing in Eng. Design and Control*.

Roy, R., Parmee, I.C., & Purchase, G. (1996). Integrating the genetic algorithm with the preliminary design of gas turbine cooling systems. *Proc. 2nd Int. Conf. Adaptive Computing in Eng. Design and Control*, 228–235.

Wade, G., Roberts, A., & Williams G. (1994). Multiplier-less FIR filter design using a genetic algorithm. *IEE Proc.—Vision and Signal Processing 141(3)*, 175–180.

Watson, A., & Parmee, I.C. (1997). Steady-state genetic programming with constrained complexity crossover using species sub-populations. *Proc. Seventh Int. Conf. Genetic Algorithms*, 315–321.

I.C. Parmee has been working in the field of evolutionary and adaptive computing since 1988. His particular interests involve the integration of such techniques with current design practice and the development of cooperative evolutionary/adaptive strategies that support and enhance engineering design abilities. He is Director of the Plymouth Engineering Design Centre at the University of Plymouth and coordinates a research team, all of whom are involved in evolutionary/adaptive design research. The team collaborates closely with industry in fields ranging from large-scale civil engineering to aerospace and electronic design. Dr. Parmee has approximately 80 journal, conference, and book publications in this evolutionary design field and is a Reader in evolutionary/adaptive computing in design and manufacture within the School of Computing.

https://doi.org/10.1017/S0890060498125039 Published online by Cambridge University Press