
Improving problem definition through interactive evolutionary computation

I.C. PARMEE

Advanced Computation in Design and Decision-Making, Department of Computing, Engineering and Mathematical Science, University of the West of England, Bristol BS 16 1QY, United Kingdom

(RECEIVED October 9, 2001; ACCEPTED March 29, 2002)

Abstract

Poor definition and uncertainty are primary characteristics of conceptual design processes. During the initial stages of these generally human-centric activities, little knowledge pertaining to the problem at hand may be available. The degree of problem definition will depend on information available in terms of appropriate variables, constraints, and both quantitative and qualitative objectives. Typically, the problem space develops with information gained in a dynamical process in which design optimization plays a secondary role, following the establishment of a sufficiently well-defined problem domain. This paper concentrates on background human-computer interaction relating to the machine-based generation of high-quality design information that, when presented in an appropriate manner to the designer, supports a better understanding of a problem domain. Knowledge gained from such information combined with the experiential knowledge of the designer can result in a reformulation of the problem, providing increased definition and greater confidence in the machine-based representation. Conceptual design domains related to gas turbine blade cooling systems and a preliminary air frame configuration are introduced. These are utilized to illustrate the integration of interactive evolutionary strategies that support the extraction of optimal design information, its presentation to the designer, and subsequent human-based modification of the design domain based on knowledge gained from the information received. An experimental iterative designer or evolutionary search process resulting in a better understanding of the problem and improved machine-based representation of the design domain is thus established.

Keywords: Interactive Evolutionary Design; Problem Definition; Information Extraction, Processing, and Presentation

1. INTRODUCTION

Evolutionary engineering design concerns the integration of population-based stochastic search, exploration, and optimization processes with complex, multivariate design problem domains. The research described in this paper particularly relates to the integration of such processes with higher levels of design, where problem conceptualization is a complex, largely human-centered activity, supported by a range of relatively basic computational simulations of the problem domain. The design models typically utilized during this stage of design do not involve computationally ex-

pensive analysis. Evaluation time can be measured in seconds. These models are not intended to provide definitive solutions. They provide an indication of performance that, when combined with experiential knowledge, is sufficient to support initial decision making as to appropriate design direction.

Conceptual design generally consists of a search across an ill-defined space of possible solutions using fuzzy objective functions and vague concepts of the final structure. Solutions and partial solutions are explored and assessed in terms of their feasibility with regard to the constraints and objectives considered relevant at that time. Heuristics, approximation, and experimentation play a major role with a high degree of flexibility, as evident in the establishment of domain bounds, objectives, and constraints (Su, 1990; Navinchandra, 1991). The design environment itself will evolve with the solutions as the designer and design team

Reprint requests to: Prof. I.C. Parmee, Advanced Computation in Design and Decision-Making, Faculty of Computing, Engineering and Mathematical Science, University of the West of England, Frenchay Campus, Coldharbour Lane, Bristol BS 16 1QY, UK. E-mail: iparmee@uwe.ac.uk or iparmee@ad-comtech.co.uk

gain an understanding of the functional requirements and the resulting structures (Gero et al., 1994; Gero & Schnier, 1995; Maher et al., 1995). Simple computer-based models, which may be both qualitative and quantitative in nature, may be utilized in order to establish an initial direction. This decision-making environment is characterized by uncertainty in terms of lack of available data and a poorly defined initial specification. Discovery and the accumulation of knowledge appertaining to problem definition and objective preferences are prevalent in this highly dynamic human- or machine-based process. The following quote (Goel, 1997) relating to creative design captures these aspects:

problem formulation and reformulation are integral parts of creative design. Designers' understanding of a problem typically evolves during creative design processing. This evolution of problem understanding may lead to (possibly radical) changes in the problem and solution representations. [...] in creative design, knowledge needed to address a problem typically is not available in a form directly applicable to the problem. Instead, at least some of the needed knowledge has to be acquired from other knowledge sources, by analogical transfer from a different problem for example. [...] creativity in design may occur in degrees, where the degree of creativity may depend upon the extent of problem and solution reformulation and the transfer of knowledge from different knowledge sources to the design problem.

This paper presents a number of techniques and processes that support human interaction through the machine-based generation of relevant, high-quality solution information appertaining to the design problem at hand. Such information provides problem specific knowledge which, when analyzed and processed by the designer and design team, leads to a better understanding of primary problem factors and a reformulation of the problem to better explore perceived areas of importance.

A brief introduction to the overall concept of interactive evolution is followed by a common decision-making analogy illustrating the manner in which problem definition, in terms of variables, constraints, and objective preferences, can change as problem-specific knowledge is accumulated. An early attempt at capturing such knowledge within an engineering design domain relating to gas turbine blade cooling systems is then introduced, followed by more recent work where a more interactive, continuous search and exploration process was developed.

The objective of this paper is to illustrate the potential of interactive evolutionary processes in supporting the designer during the complex problem formulation aspects of conceptual design. In order to achieve this the interactive system is described as a whole, rather than concentrating upon the finer details of single components. Such detail, however, can be found in the referenced texts.

2. INTERACTIVE EVOLUTIONARY COMPUTATION (EC)

There is a history of research relating to interactive evolutionary computing that, in the main, relates to partial or complete human evaluation of the fitness of solutions generated from evolutionary search. This has generally been introduced where quantitative evaluation is difficult if not impossible to achieve. Examples of applications include graphic arts and animation (Sims, 1991a, 1991b), automotive design (Graf & Banzhaf, 1995), food engineering (Herdy, 1997), and database retrieval (Shiraki & Saito, 1996). Such applications rely upon a human-centered, subjective evaluation of the fitness of a particular design, image, taste, and so forth, as opposed to an evaluation developed from an analytic model.

Partial human evaluation and interaction are also in evidence, for instance, user interaction relating to an evolutionary nurse scheduling system where a schedule model provides a quantitative evaluation of a solution. However, the model may not prove adequate in terms of changing requirements, qualitative aspects, and so forth. In this case the user must add new constraints in order to generate solutions that are fully satisfactory (Inoue et al., 1999). In the pharmaceutical industry, computational biology involves the modeling of biomolecular systems. Genetic algorithms (GA) can provide the search process for the identification of optimal biomolecule combinations. The process can be enhanced, however, by the user introduction of new combinations as an elite solution into selected GA generations (Levine et al., 1997).

All the above applications utilize a major advantage of stochastic population-based search techniques. This relates to their capabilities as powerful search and exploration algorithms that can provide diverse, interesting, and potentially competitive solutions to a wide range of problems. Parmee and Bonham (1999) propose that such solutions can also provide information to the user that supports a better understanding of the problem domain while helping to identify the best direction for future investigation. This perspective relates to human interaction when operating within ill-defined and uncertain decision-making environments in order to improve definition, increase confidence, and identify direction. The role here for EC relates to exploration and the gathering of optimal information from simple conceptual design models. Such information supports model development by the user in an iterative, interactive EC environment, where the first task is to evolve the design space before attempting to solve the problem.

In the development of such highly interactive evolutionary design systems (IEDS), EC may provide the underlying search capability whereas other technologies, some from the computational intelligence domain, can provide the necessary control, information extraction, data-processing, and presentation tasks. Perhaps overall systems could be considered "immersive" in that the user plays an integral, cen-

tral role in receiving optimal or interesting information from the system and analyzing such information off-line before introducing change in terms of the underlying model or evaluation function. In this manner it is possible that user experiential knowledge of a particular problem area can be captured in a further evolutionary search of the redefined design space. It is also likely that the design space from which the final solution is identified is significantly different from the initial space within which the interactive search commenced.

3. MOVING GOAL POSTS

Discovery and knowledge accumulation are aspects of problem solving that are common across decision making as a whole. We can illustrate the manner in which problem spaces change with knowledge gained in a relatively simple manner via decision-making environments familiar to most (Parmee, 2001). Although unrelated to design activities, such environments could be seen as analogous in terms of discovery, knowledge accumulation, problem reformulation, and the eventual identification of a best compromise solution.

For illustrative purposes let us therefore consider a job-related relocation to a new city and the daunting problem of finding a family home. The initial investigation will likely relate to identifying appropriate districts based upon criteria relating to quality of local schools; safety or security issues; proximity to places of work, transport, highway networks, shopping and leisure facilities, and so forth, plus the average price and type of housing and the overall environment. Other criteria relate directly to the ideal property such as maximum cost, number of bedrooms, garden, garage, parking, and so on. Several of the above criteria would be considered hard constraints (i.e., maximum cost) in the first instance.

The decision-making team is the family who would all probably rate the relative importance of the above criteria in a slightly different manner and whose opinions will carry a varying degree of influence. It is also likely that experience relating to former moves will play a role. A pretty clear vision of what the ideal property will look like and consist of will exist initially. A preferred location may also have been identified with an initial ranking of alternative districts.

Initial information gathering will provide quantitative and qualitative data relating to location from a wide variety of sources, some reliable and some only anecdotal. Gradually, an overall picture will be established that will result in the possible elimination of some options and the inclusion of new possibilities. Of significant importance is the possible discovery of other previously unknown locations during explorative trips to those areas already identified. Such locations may be considered interesting and worthy of further investigation.

As details of properties are gathered it will likely become apparent that the ideal solution is hard to find and the concept of compromise becomes a reality. Hard constraints may soften, whereas objective preferences will constantly be discussed and redefined in the light of accumulated knowledge regarding districts and property availability within them. Particular characteristics of areas initially thought to be unsuitable may suddenly appear attractive. The search effort may shift as it is discovered that such areas have suitable properties within the preset price range. Alternatively, the initial hard constraint relating to maximum price may soften as close to ideal properties in favored locations become available. Possible compromises must then be explored in an attempt to accommodate any related increased cost.

The whole decision-making process becomes an uncertain mix of subjective or objective decisions as requirements evolve, objectives rapidly change, and external pressures relating to time constraints begin to take precedence. At the end of the day it is quite probable that the chosen home differs significantly from the one first envisaged. Perhaps a guest bedroom may have been sacrificed and the garden may be far smaller but the location is ideal. Alternatively, the route to work may be longer and more tortuous but a property close to ideal at the right price in an up and coming neighborhood may have been found.

Although seemingly simple, the overall search process is highly complex. Uncertainty, compromise and problem redefinition are inherent features. Analogies relating to design decision-making scenarios are apparent. Such everyday decision-making scenarios can perhaps help us when developing interactive search environments that support more complex decision processes. The concept of unexpected discovery as the decision makers move through the search space is of particular interest.

4. EVOLUTIONARY DESIGN SYSTEMS

4.1. Background

There are many examples of the application of evolutionary computing techniques to specific design problems, the publication of such applications being increasingly evident across the spectrum of engineering journals in recent years. Such applications tend to be domain specific and generally relate to the optimization of a particular component where well-defined mathematical models of the component provide an evaluation function (e.g., Parmee & Vekeria, 1997; Elby et al., 1998; Hajela, 1998; Koza et al., 1998; Olhoffer et al., 2000). More rarely, application has related to whole system design or optimization (e.g., Parmee, 1996*b*; Emmerich et al., 2000; Hillermeier et al., 2000). Significantly less research effort has concentrated on the higher conceptual levels of the design process, where problem representation and uncertainty can cause significant difficulty in terms of the design of an appropriate evaluation function. Gero (1998),

Schnier and Gero (1998), Corney (2000), and O'Reilly and Testa (2000) provide interesting approaches to conceptual design through the utilization of evolutionary and associated adaptive techniques.

It is within uncertain, poorly defined domains that human or evolutionary process interaction has been introduced, as described in Section 2. However, little of this work concerns engineering design. One related area that has received a great deal of attention is that of evolutionary multiobjective satisfaction or optimization, in which there are several examples of applications relating to uncertainty and the development of fuzzy models. The reader is directed to Zitzler et al. (2001) and Deb (2001) for application examples and state of the art theory.

A prototype interactive evolutionary machine-based search and exploration environment has been proposed that provides relevant problem information to the designer or decision-making team (Parmee et al., 2000; Parmee et al., 2001). The intention is that such information can be processed and subsequent discussion can result in the recognition of similarities with other problem areas and the discovery of possible alternative approaches. Population-based evolutionary search generates a large amount of possibly relevant information, most of which may be discarded through the actions of various operators. Interactive systems support the capture of such information and its utilization in the subsequent reformulation of the problem through the application and integration of experiential knowledge.

It is apparent that a core activity of design decision making relates to the gathering of information concerning diverse aspects of the problem space. It has been proposed (Parmee & Bonham, 1999) that a primary role of evolutionary machine-based search and exploration processes relates to the generation of such information. This moves the utilization of EC away from application over a set number of generations or until some convergence criteria is met. Instead, a more continuous exploratory process, in which changes to objective weightings, variable ranges and constraints based upon the information generated results in a moving, evolving problem space.

Initial related research resulted in the development of cluster-oriented GAs (COGAs). The objective of the COGA approach is the extraction of relevant design information from good solutions over over HP regions of complex multivariate problem spaces. This approach has been improved and developed and now represents an integral component of the IEDS. Other related work concerns evolutionary conceptual design exploration for the identification of optimal alternative system configurations through the utilization of dual-agent strategies for search across mixed discrete or continuous design hierarchies (Parmee, 1996). These approaches led to an initial attempt to concurrently satisfy both quantitative and qualitative criteria through the integration of fuzzy rule bases with evolutionary search. This work, although not interactive, is briefly described in Section 4.2 to illustrate earlier attempts to capture and integrate designer knowl-

edge with evolutionary search. The details of all the above plus associated work relating to evolutionary constraint satisfaction or optimization, structural design, and systems identification can be found in Parmee (2001).

The IEDS supports a relatively continuous, iterative user or evolutionary search process through the utilization of EC, agent-based approaches and a number of other complementary techniques. The manner in which evolutionary computing-based information gathering can support decision making in complex conceptual design environments is illustrated via an overview of the system in Section 5. A more detailed treatment can be found in the individual referenced texts.

4.2. The Qualitative Evaluation System (QES)

The QES strategy is an earlier attempt to provide support to the designer when attempting to determine trade-offs between both quantitative and qualitative criteria. This support utilizes a linguistic rule base (i.e., natural language statements of rules generated in close collaboration with Rolls Royce turbine engineers) that resides within a fuzzy expert system. The rules relate to the comparative effectiveness of a GA-generated design solution in terms of manufacturability, choice of materials, and a number of special preferences relating to in-house capabilities. Quantitative aspects of a design are combined with qualitative ratings to generate a measure of the overall fitness of the solutions. Domain knowledge concerning variable preferences and heuristics is utilized and combined, using a concept of compromise (Roy et al., 1996a, 1996b).

The work concerns the preliminary design of gas turbine engine cooling hole geometries. The primary objective is to minimize the mass flow through the radial cooling hole passage. A GA incorporating adaptive restricted tournament selection (ARTS) identifies a number of single high performance (HP) solutions from the design space. The QES receives all the design variable values of each of these solutions as inputs and develops an overall qualitative rating concerning the effectiveness of the design as a whole.

ARTS is utilized to first identify a number of "good" quantitative (i.e., minimal mass coolant flow) design solutions. These good solutions are next evaluated by the QES, which takes the variable values of each solution as inputs and outputs a qualitative rating for the design. The QES has three components: the fuzzifier, the fuzzy inference engine, and the defuzzifier. Designer knowledge provided by Rolls Royce engineers is stored in a static fuzzy rule base. During the fuzzification stage each variable range is divided into five subranges and expressed using linguistic terms. A crisp value for the effectiveness is obtained through center of gravity type defuzzification.

The knowledge base for the system is developed using fuzzy rules and facts embodying qualitative aspects of the design problem in terms of manufacturability, choice of materials, and the designer's special preferences. The knowl-

edge base is presented in three categories: *intervariable knowledge*, which relates to the relative importance of each variable in terms of the objective function; *intravariabile knowledge*, which relates to preferred subsets of each variable's range (e.g., blade wall thickness needs to be *low* in terms of material cost but *high* in terms of stress considerations; and *heuristics*, which mostly concern specific cases where there is no uncertainty concerning the conclusion.

The inter- and intravariabile knowledge is then integrated using a concept of compromise that is implemented to reduce the severity of qualitative ratings. Intervariabile knowledge determines the degree of compromise possible on every variable (slight compromise, less compromise, compromise, and more compromise).

Having evaluated solutions both *quantitatively* via the ARTS GA process and *qualitatively* via the QES, the results were initially presented in textual form. However, further development combined quantitative and qualitative components in a graphical manner that facilitates overall understanding of the major aspects of the problem. Four solutions for each of three internal cooling-hole geometries (plane, ribbed and pedestal) are presented in Figure 1 with their relative quantitative fitness plainly shown by the major bars of the chart. The qualitative ratings for each solution are then shown as a series of embedded, shaded bars. Such a representation presents much information in a relatively transparent manner. It is apparent, for instance, that, although the plane cast internal geometry provides a low quantitative fitness, the solutions are relatively robust in terms of the qualitative criteria. The qualitative bars relating to the other two geometries show a much greater degree of

fluctuation. This can perhaps provide insight into the problem characteristics that aid the designer in terms of both modeling the system and determining preference rankings for each criterion. For instance, if the priority is for a solution that can be considered low risk in terms of possible problems relating to manufacturing, material, and special preferences aspects, a plane cast internal geometry can be chosen. Losses relating to quantitative performance must then be made up within other areas of the turbine design to achieve a compromise. However, if quantitative performance is paramount, then a pedestal geometry that best satisfies preferences relating to the three qualitative criteria may be considered appropriate, although the robustness of the design may suffer.

The QES provides a good indication of the relative merits of HP solutions in terms of a number of qualitative criteria. A possible problem area here however is the flexibility of the rule base. There is a requirement for on-line rule changes that should be easy to implement by users with no knowledge of fuzzy inference technologies. A high degree of flexibility in terms of the representation of objectives and their weightings is a major requirement during these higher levels of the design process. It is possible that the major utility offered by the QES relates to more routine design tasks where problem definition is already high.

5. INTRODUCING THE IEDS STRATEGY

The QES illustrates the manner in which qualitative criteria can be combined with quantitative evaluation to provide

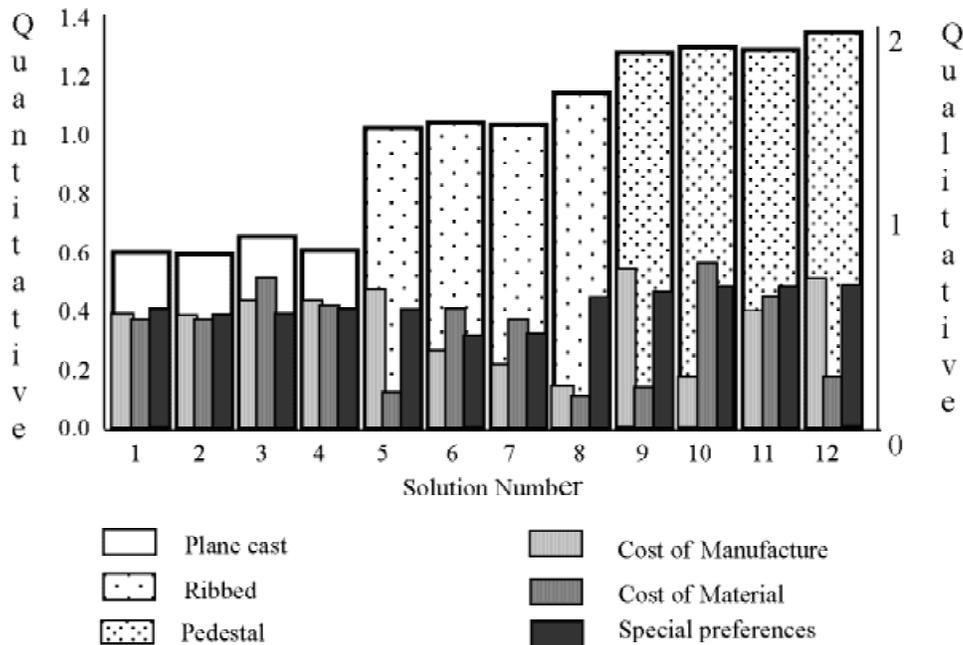


Fig. 1. A graphical representation of ARTS and the QES.

important information to support designer decision making. However, it does not readily support on-line interaction. The requirement for a system that supports the on-line extraction of information that can be presented to the user in a succinct manner, thereby supporting easily implemented change, has led to an investigation of various techniques that can be combined within an overall architecture.

The satisfaction of multiple objectives (i.e., >10) is a major requirement, and such objectives must be entirely flexible in terms of preferences or weightings to allow adequate exploration of the problem domain and to support a better understanding of the complex interactions between variable space and objective space.

The developed system involves a number of machine-based processes that can communicate, as shown simply in Figure 2. The user is an integral part of the system in accepting, analyzing, and processing information from the system, introducing objective change via the preferences component and variable parameter and constraint changes directly into the evolutionary component. The evolutionary component can operate in several modes:

1. single evolutionary process,
2. multiple individual evolutionary processes with no co-operation, and
3. multiple coevolving processes.

5.1. HP region identification and information extraction

Mode 1 relates to regional identification aspects where the COGAs of the information gathering component extract information relating to variable interaction, variable redun-

dancy, and the setting of appropriate variable parameter ranges. The basic structures of COGAs and the associated adaptive filter have been described in a number of papers (Parmee, 1996a; Parmee & Bonham, 1999). Their function relates to the rapid identification of HP regions of complex, multivariate design space. At the COGA core is an explorative GA. High levels of search space exploration are required and were initially promoted through variable mutation COGA regimes. More recently, the integration of various sampling techniques (Bonham & Parmee, 1999b) with the COGA strategy has improved exploratory performance. The adaptive filter extracts and scales populations (in terms of fitness) with a continuous GA evolutionary process by only allowing solutions that lie above a filter threshold to pass into a final clustering set. The design exploration capabilities are well described in Parmee and Bonham (1999), along with extensive discussion relating to possible interactive utilization.

COGA operation within the information gathering component is illustrated via the preliminary design of military aircraft. This is a complex design domain characterized by uncertain requirements and fuzzy objectives relating to the long gestation periods between the initial design brief and realization of the product. Changes in operational requirements in addition to technological advances cause a demand for a responsive, highly flexible strategy in which design change and compromise are inherent features for much of the design period. Design exploration leading to innovative and creative activity must be supported. The ability to introduce rapid change to satisfy the many operational, engineering, and marketing considerations as they themselves change is essential. In this case, the COGA software is manipulating the BAE Systems MINICAPS model (Webb, 1997). This model is a much condensed version of the Computer-Aided Project Studies (CAPS) suite of soft-

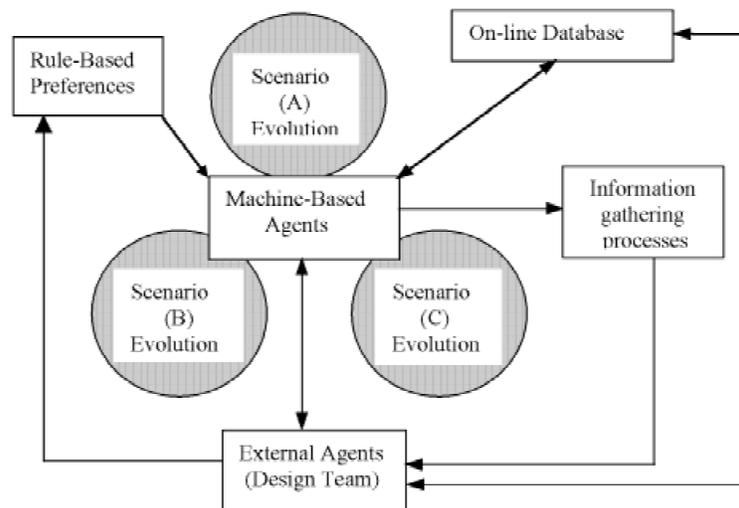


Fig. 2. A schematic of the interactive evolutionary design station.

ware for conceptual and preliminary airframe design. MINICAPS maintains many of the characteristics of the overall suite, especially in terms of multiple objectives. The 9 input variables that define the problem space can generate up to 13 outputs.

Developed software allows identified HP solution clusters to be projected onto any 2-dimensional (2-D) hyperplane of the included variables. Figure 3 illustrates the manner in which COGAs can be utilized to better understand the distribution of HP solutions within a complex high-dimensional fitness landscape (i.e., a landscape described by multiple variable ranges and a measure of performance of all possible variable value combinations). The results of consecutive independent COGA runs are shown on 2-D hyperplanes relating to wing thickness to chord ratio and wing leading edge sweep angle variables. Initially, low adaptive filtering thresholds allow high levels of lower

fitness solutions to pass into the final clustering set. This is illustrated by the wide ridge of solutions arcing from the top left to the bottom right of the 2-D hyperplane [Fig. 3(a)].

As filtering becomes increasingly severe in subsequent runs, this region becomes a narrow ridge of solutions and we gradually see the emergence of two HP regions in the top left and lower right quarters [Fig. 3(b,c)]. At higher filtering thresholds, the two specific regions of HP, A and B, become apparent [Fig. 3(d)].

These examples clearly illustrate the effect of varying the filtering threshold. Low filtering provides the designer with maximum information relating to the general nature of the search space. Conversely, high filtering greatly reduces the set cover and produces a limited number of near optimal solutions.

The perceived utility of this approach is that information relating to a wide range of solutions is available during

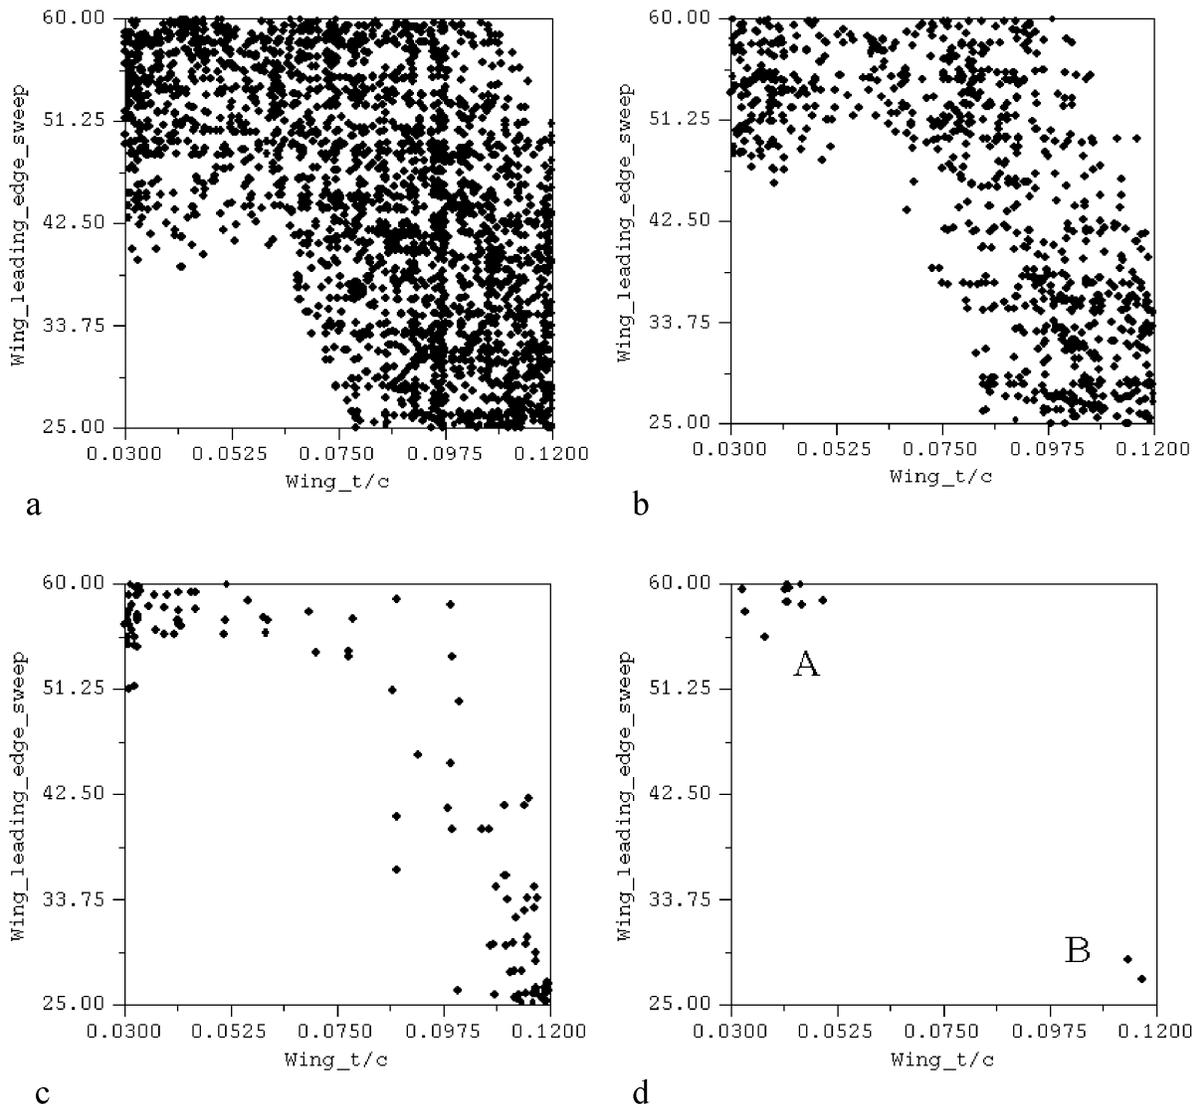


Fig. 3. The application of variable mutation COGA to preliminary airframe design.

preliminary runs utilizing low filter settings. Regions of medium performance in terms of primary objectives may be discovered that may be considered HP if objective preferences change. Objectives initially considered of lesser importance may become more significant if a certain payoff can be achieved by adopting an alternative design approach based on solutions discovered during these exploratory runs. Such regions can therefore be presented for design team discussion and further off-line investigation and processing. This may result in a redefinition of the problem and of the variables, constraints, and objectives of the defining model.

During collaborative work, this approach has consistently been supported by practicing engineers who are very aware of the possible existence of interesting solutions that require further investigation, in addition to a concentration of effort on “HP solutions” in terms of major objectives. The COGA approach supports the identification of interesting regions or solutions that the designer may wish to revisit and investigate in finer detail. This relates back to the house-hunting analogy, where interesting districts may be discovered while traveling to a favored location. In the house-hunting case, however, search is across a 2-D plane, whereas in terms of design, the negotiation of a complex high-dimensional surface is necessary. This involves identifying interesting solutions while concentrating on those objectives perceived at that time to be most significant.

Upon successful identification of HP regions, local perturbations can be applied to the solutions defining each region to generate further solutions, thereby improving the regional set cover. Such perturbed solutions do not pass through the adaptive filter, and therefore may be of relatively low fitness. The standard deviation of the fitness of all the final solutions within the region will now give an indication of solution sensitivity within that region. In this manner, information relating to solution robustness can be accumulated to further assist the designer in the selection of appropriate design direction.

Another aspect of this approach relates to designer confidence in the preliminary design models being utilized. The return of single HP solutions may be considered high risk due to uncertainties relating to the relatively simple analysis adopted by the model. The identification of similarly medium- or high-performing solutions in the immediate neighborhood could prove reassuring. Alternatively, a lack of such solutions may cast doubt upon the (possibly erroneous) single solution, leading to analysis and further development or amendment of the preliminary design model.

One aspect of variable interaction is shown in Figure 4. The left-hand column shows HP regions relating to the turn rate and excess power, plotted in the gross wing plan area or wing aspect ratio hyperplane. These graphs plainly indicate to the engineering designer the settings for the upper and lower bounds of the two variables for further search effort. Clicking on any solution will cause its value and the variable vector associated with it to appear on the screen.

The right-hand column, however, shows the corresponding distribution of HP solutions in the climb mach number cruise height hyperplane. In this case a more uniform distribution of such solutions across this hyperplane is evident. The engineer must tread carefully when assessing appropriate bounds for these two variables. It was initially assumed that these hyperplanes indicate low sensitivity of the respective objectives to these variables because HP solutions are available throughout the variables’ ranges (Parmee & Bonham, 1999). However, further analysis of the solution distribution indicates that it is possible that the higher performing solutions may occupy isolated peaks within the fitness landscape and could therefore be considered highly sensitive to slight perturbation of the variable values. The designer in this case has the following possible options:

- Maintain the existing variable ranges of that hyperplane and search across the whole space described by them in order to identify individual optimal solutions.
- Concentrate search in major areas containing the very HP solutions.
- Select a very HP solution that lies within a preferred subset of the variables’ ranges and accept that this solution offers “the best” values for the variables. This, in effect, transforms the two variables to fixed parameters, therefore reducing the dimensionality of the overall design space.
- Identify a small grouping of very HP solutions and radically reduce the space so that subsequent search, in terms of the two variables, is concentrated within this succinct region. It is possible that such a region offers greater solution robustness due to the density of very HP solutions.

Further research involving software agent analysis of the variable vectors relating to the solutions of HP regions is continuing. Objectives of this analysis include agent-generated advice to the user relating to possible variable range reduction while directing the designer to those hyperplanes containing high-value information, that is, information appertaining to the setting of appropriate variable bounds or relating to highly sensitive variable combinations. The overall objective of such agent-based support is to lessen the load on the designer relating to the processing of information concerning multidimensional aspects of the problem at hand.

5.2. COGAs as multiobjective information gatherers

Mode 2 concerns the concurrent identification of HP regions of the problem space relating to individual objectives and the subsequent definition of common HP regions where best-compromise solutions satisfying several objectives may be found. Again, the COGA techniques of the information gathering component are involved. Identified HP regions can be overlaid upon selected 2-D hyperplanes described

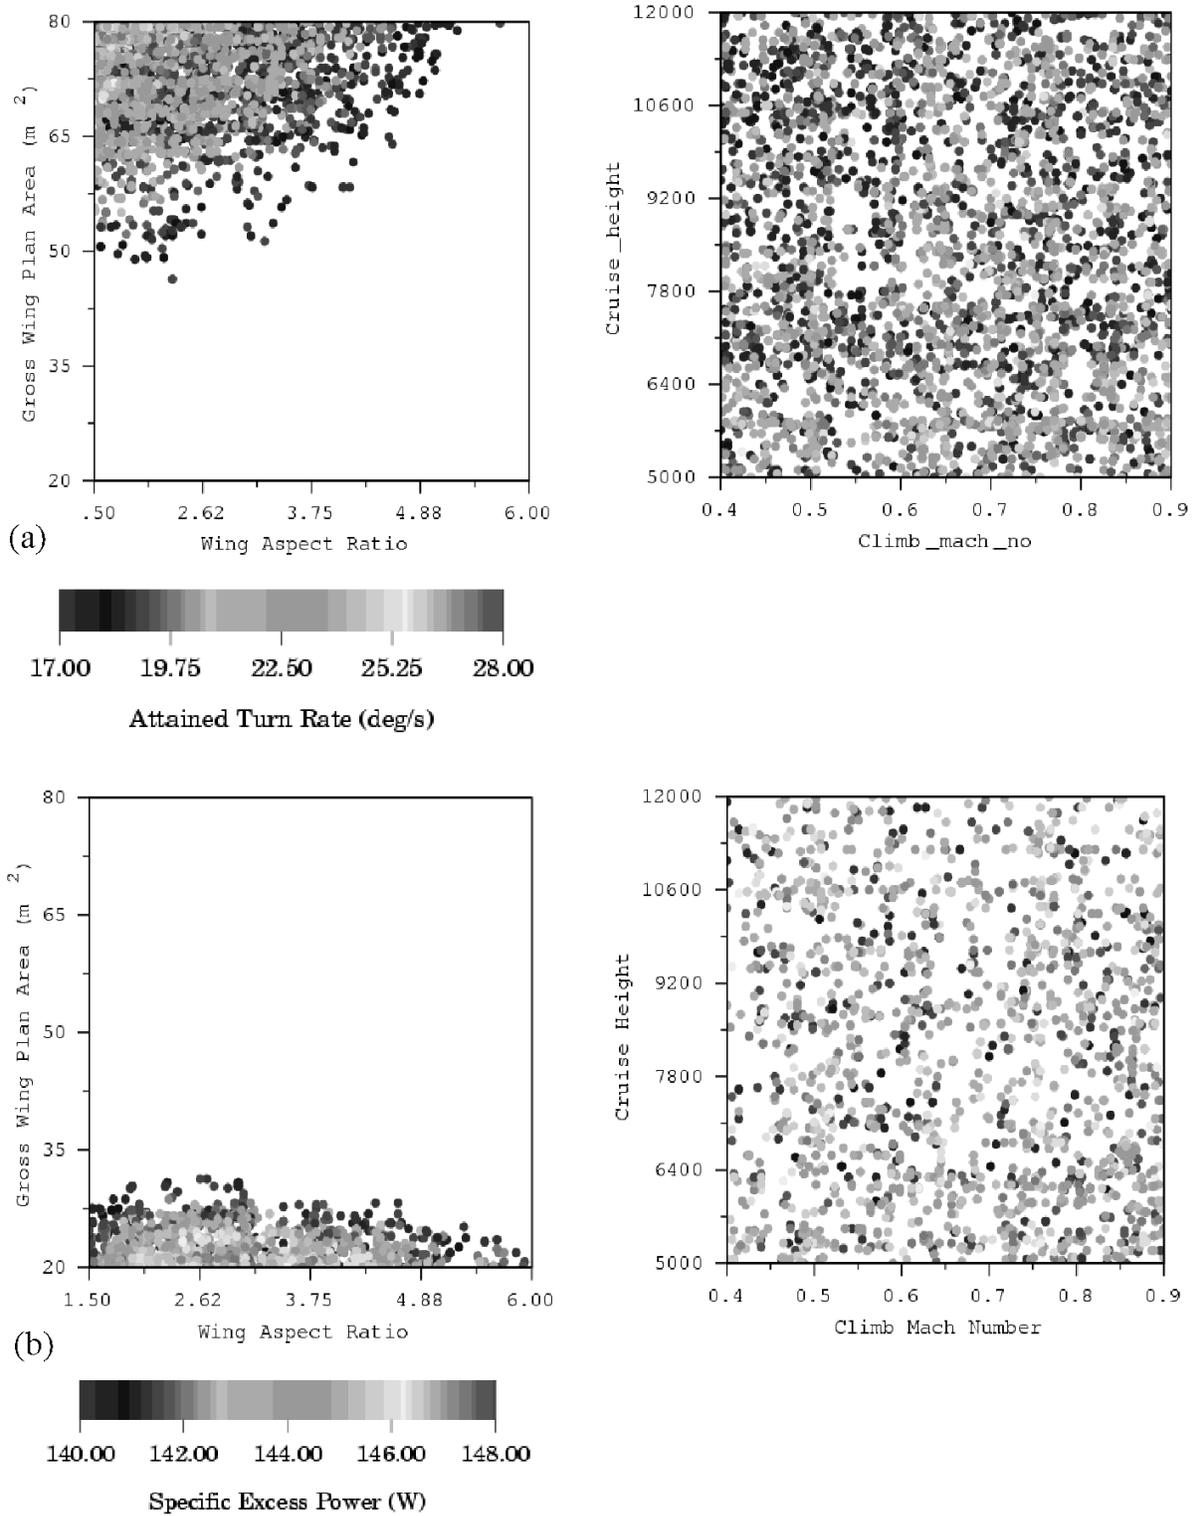


Fig. 4. A comparison of a projection of results on two differing hyperplanes for (a) the attained turn rate objective and (b) the specific excess power objective.

by pairs of variables selected from the set of variable parameters that describe the problem space.

Figure 5(a) shows HP regions relating to three MINICAPS objectives: attained turn rate (ATR), specific excess

power (SEP), and ferry range (FR), which are plotted on the gross wing plan area or wing aspect ratio, variable parameter hyperplane. All objectives are considered to be of equal importance. No objective preferences or weightings

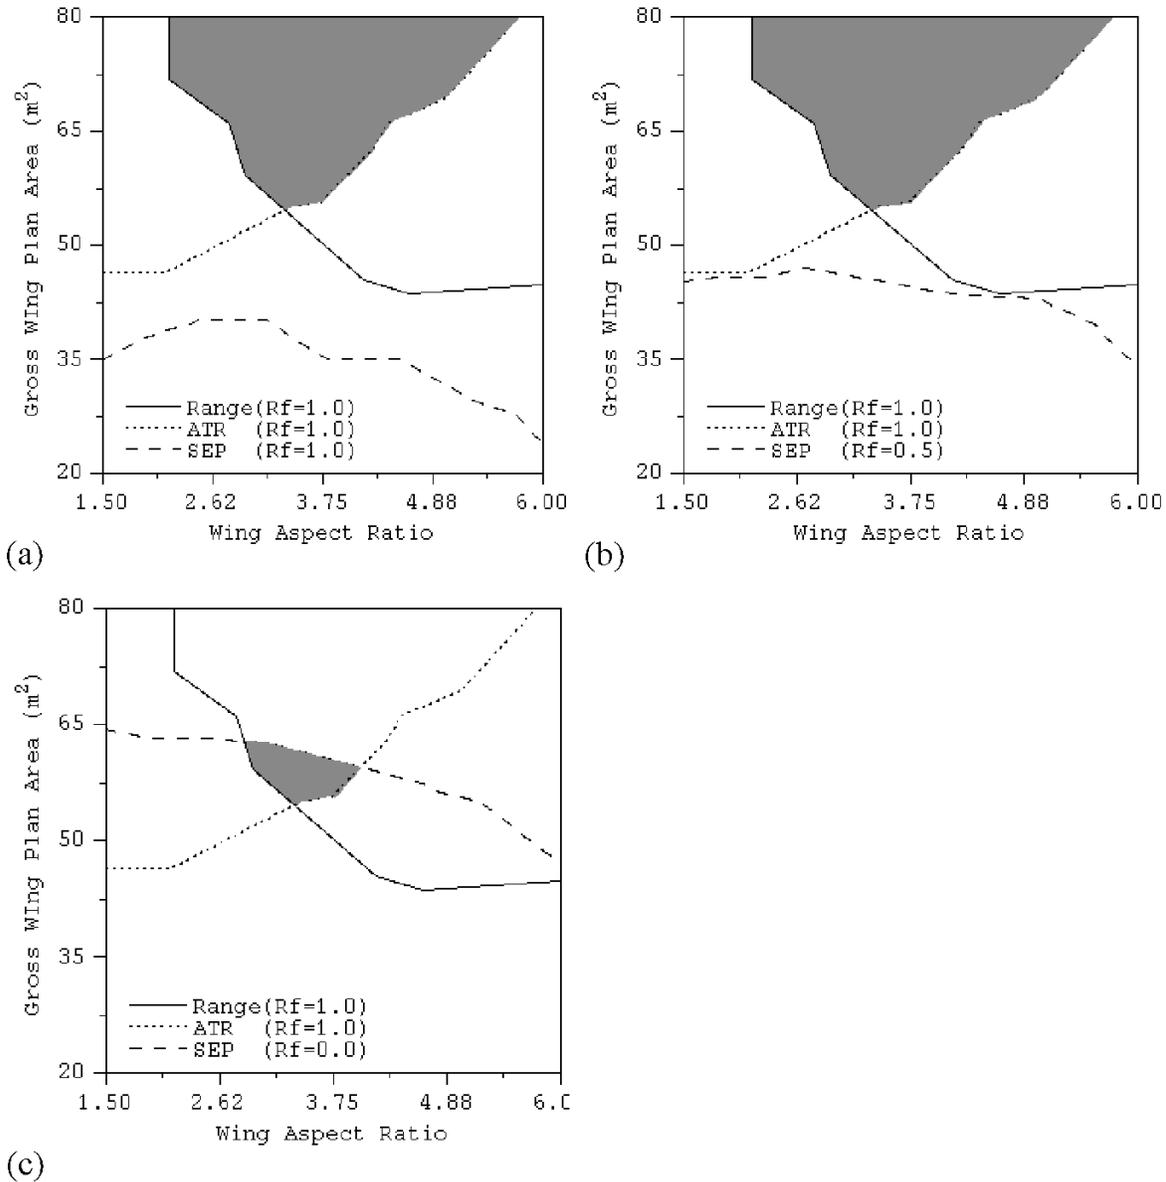


Fig. 5. The identification of compromise HP regions relating through filter threshold relaxation. (a) A common region for ferry range and attained turn rate (ATR) has been identified but specific excess power (SEP) objectives cannot be satisfied. (b) Relaxing the SEP filter threshold allows lower fitness solutions through and boundary moves. (c) Further relaxation results in the identification of a common region for all objectives.

are introduced. As can be seen, regions relating to ATR and FR overlap, forming a region containing HP compromise solutions. There is no mutually inclusive HP region relating to SEP, however. By returning to COGA and reducing the severity of the adaptive filter in relation to SEP solutions, the SEP region can be expanded, as shown in Figure 5(b,c) until a mutually inclusive region involving all objectives is identified. This relaxing of the adaptive filter threshold allows lower performance SEP solutions through to the final clustering set. This could be considered equivalent to a lessening of the relative importance of this objective (i.e., redefining preference ranking of the objectives).

The technique allows the projection of objective space onto variable space. This gives a visual appreciation of the interaction between the various objectives and supports the user in the determination of initial preferences concerning their relative performance. A clear indication of the degree of conflict between objectives is apparent, and experimental changes to the filtering factor relating to each objective can indicate the degree of difficulty likely to be encountered upon the introduction of more definitive multiobjective optimization processes.

Current work is investigating the relationship of the solutions within the mutually inclusive regions to the non-

dominated solutions of the Pareto frontier. This work indicates that a good approximation to the Pareto front is contained within these regions. This approximation can be realized by identifying the nondominated solutions contained within the HP regions relating to all of the objectives.

It is stressed that this visual representation provides an indication only. Even if mutually inclusive compromise regions are apparent in all variable parameter hyperplanes, a possibility still exists that such compromise regions do not exist to the extent suggested in the graphical representations. This could be due to highly convoluted HP regions. However, solution vectors describing the regions are available, and agent-based systems can utilize this data to check the validity of compromise regions across all dimensions. Such agents can then inform the designer of any possible need for caution when developing assumptions from the visual images.

The aim of this work, however, is to support a better understanding of objective interaction and conflict through graphical representation rather than providing a succinct and accurate representation of compromise regions or the Pareto frontier. In this sense, the technique again supports the generation of information pertaining to the problem at hand, in which variables and objectives can vary as problem knowledge expands. The approach therefore takes into consideration the uncertainties and poor definition inherent in the utilization of preliminary design models and in the degree of initial understanding of the problem domain. The basic notion of “garbage in, garbage out” must be taken into consideration. Much time could be spent on conducting a more definitive analysis to identify Pareto optimal points that prove erroneous upon the introduction of more definitive problem models. The strategy therefore indicates a probable best way forward rather than the global solution to a design problem that, at this stage, is poorly defined.

The flexibility of the GUI allows objectives to be included or disregarded while also allowing variable ranges to be altered in order to support the investigation of specific regions and objective or variable interaction. Further machine-based support for such activity will be necessary, however, to avoid cognitive overload. An investigation of the agent-based systems briefly described is currently addressing this area.

5.3. The preference component

The techniques of the previous section support a better understanding of objective interactions and of the degree of difficulty likely to be encountered in the satisfaction of initial objective preferences. A highly flexible interface for the introduction of such preferences to a more definitive multiobjective search process can now be introduced. Objective preference relates to a ranking of importance of included objectives.

A methodology that supports the on-line variation of design preferences has been developed. As efficient explora-

tion across the many possible different design variants is considered of more interest than the identification of single optimal solutions, the system should be able to support such exploration while also suggesting the best design direction.

It is generally accepted that it is easier for the decision maker to give qualitative ratings to objectives, that is, “Objective A is much more important than objective B,” than to set the weight w_A of objective A to, say, 0.1 or 0.09. The method of fuzzy preferences (Fodor & Roubens, 1994) and induced preference order can be utilized to introduce such rule-based preference representation for transformation into appropriate objective weightings. The following predicates (Cvetkovic & Parmee, 2000) can be introduced (Table 1):

These, together with the complementary relations $>$ and \gg , can help build the relationship matrix R that is necessary for a “words to numbers” transformation. For this transformation, concepts of “leaving score” (Fodor & Roubens, 1994), among other techniques, can be employed.

It is first necessary for the engineer to rank the objectives in terms of relative importance, but numeric weightings are not required. If transitivity is assumed, then the number of required questions to establish overall preference ratings is reduced. The preference algorithm has been described fully in a number of publications (e.g., Parmee et al, 2000), and I do not intend to reproduce it here. Examples from the GUI follow that illustrate the preference procedure. The problem domain again relates to BAE Systems preliminary airframe design.

The user first selects those objectives that require investigation. In this case, take-off distance, landing speed, FR, and mass at take-off have been selected from a possible 13 outputs from the MINICAPS model. Having selected the objectives, it is necessary to establish equivalence classes that classify them in terms of relative performance. In the example in Figure 6, two of the objectives are considered to be equally important and two further objectives are considered to have different levels of importance. Three different levels of importance are therefore processed.

The interface facilitates the on-line change of objective preferences, which allows further exploratory runs, thus providing information relating to possible problem redefinition. A more detailed description of preference integration can be found in Cvetkovic and Parmee (2001). Having selected objectives, established objective preferences, and

Table 1. Preference relationships

Relation	Intended Meaning
\approx	Is equally important
$<$	Is less important
\ll	Is much less important
\neg	Is not important
!	Is important
#	Do not know or do not care

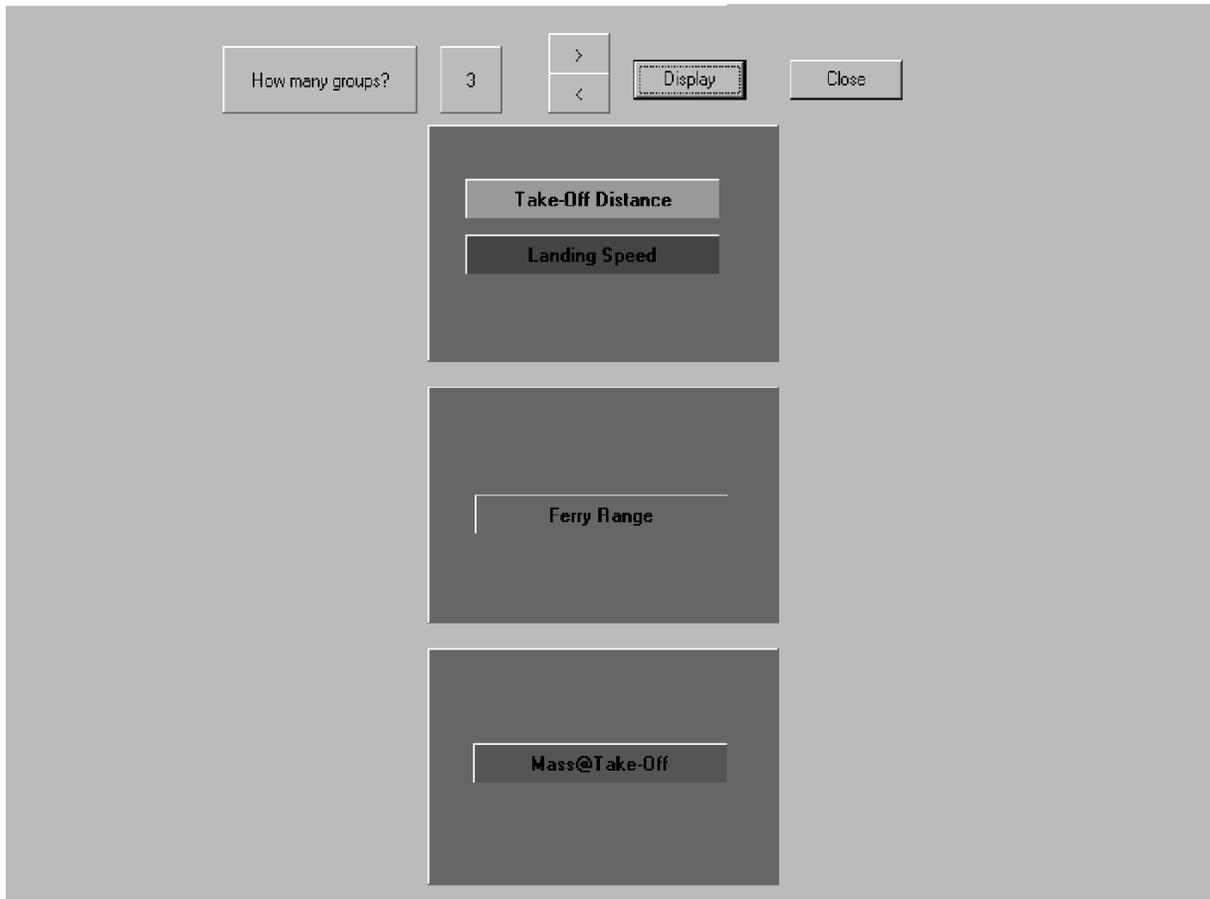


Fig. 6. Establishing equivalence classes.

performed the machine-based words to numbers transformation, the numeric weightings can be passed to the coevolutionary process module described in the following section.

5.4. The coevolutionary multiobjective approach (mode 3)

The preference component has been linked with core coevolutionary processes. The goal is to explore HP solutions relating to several objectives while providing maximum information concerning the following:

- interesting regions of complex, multidimensional, Pareto surfaces;
- single objective optimal solutions; and
- a number of solutions that best satisfy a range of ideal scenarios.

This approach is an alternative to the generation of high-dimensional trade-off surfaces comprising very large numbers of nondominated solutions that can be identified using standard EC-based Pareto approaches. The intention is to generate information that supports a better understanding of the multiple criteria aspects of the problem. The iterative

nature of the IEDS and the associated, user-instigated changes to variable, objective, and constraint space would necessitate the constant regeneration of such high-dimensional Pareto surfaces. The alternative approach adopted here identifies good compromise regions of the overall search space relating to a number of differing objectives.

The distributed method involves individual GAs for the optimization of each objective. Search relating to an individual objective takes place on each evolutionary process. Subsequently, through the application of penalties, the coevolving processes are drawn into that region of the overall space that offers the best compromise relating to all objectives and their preset preferences (Parmee et al., 2000, 2001).

During the coevolutionary run the fitness for each objective is normalized relative to the maximum and minimum values found for each GA with constant adjustment as new upper and lower limits are identified. In each generation, the variables of solutions relating to each objective are compared with those of the best individual from the other coevolving GA populations. If a variable is outside a range defined by a range constraint map, it is adjusted by a penalty function. The range constraint map reduces maximum

allowable distances between variables at each generation. Initially, the map must allow each GA to produce a good solution based on its own specified objective. As the run progresses inflicted penalties increasingly reduce variable diversity to draw all concurrent GA searches from their separate objectives toward a single compromise design region where all objectives are best satisfied. This process is

illustrated in Figure 7, which shows the individual evolution of each objective. In this case all objectives are of equal importance.

The machine-generated numeric weightings resulting from the preference ranking introduced by the user can now modify the penalties that are inflicted. A heavy penalty inflicted upon a much more important objective is therefore moder-

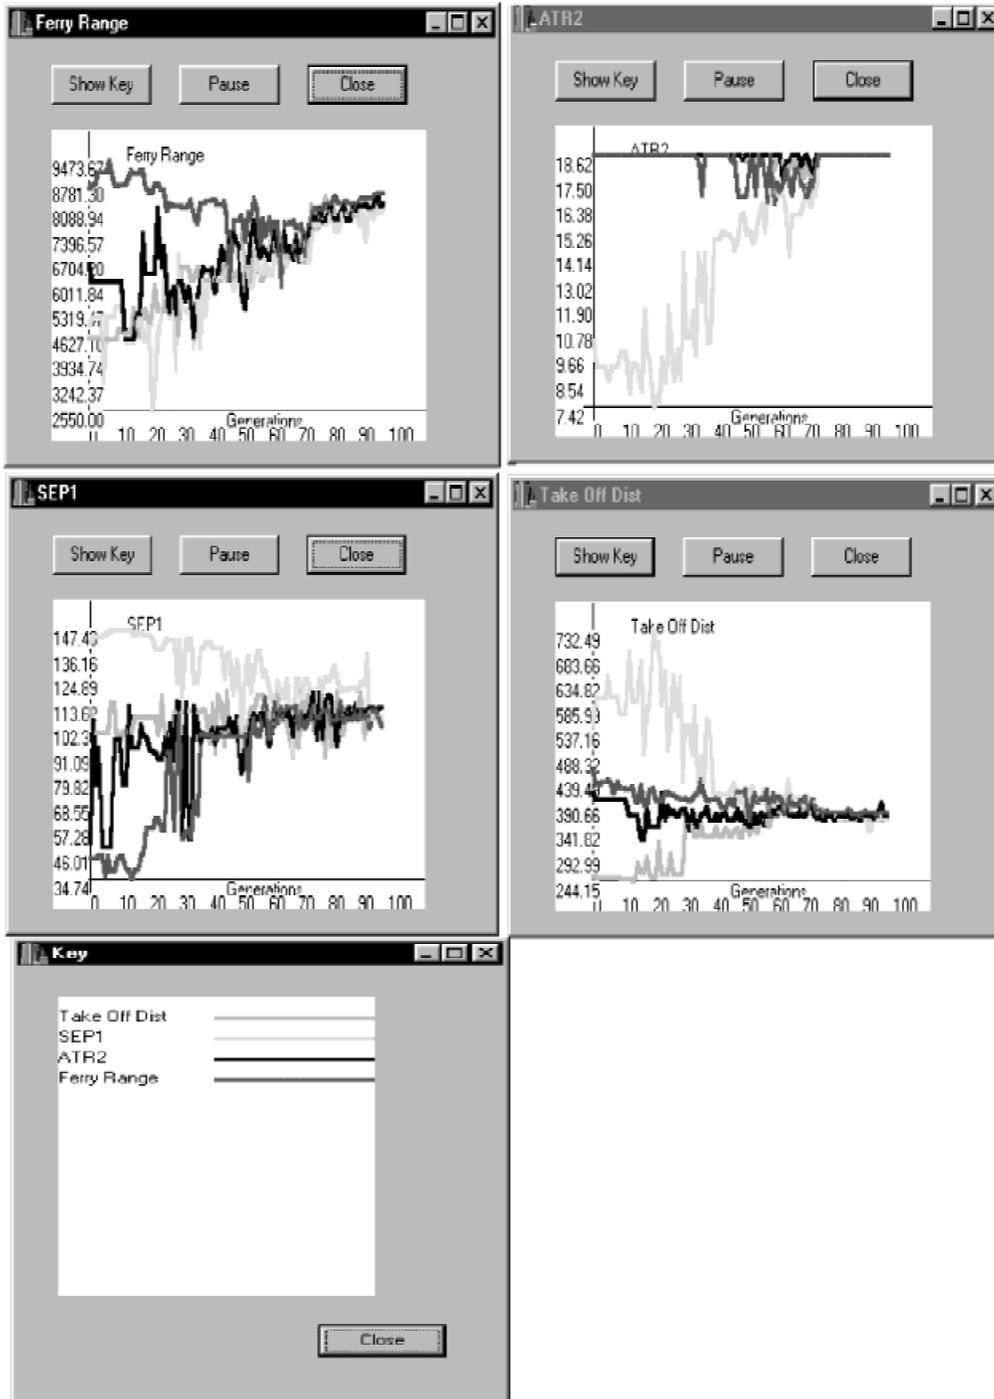


Fig. 7. The graphs show the convergence of the four evolving objectives on a best compromise region of the design space. Each graph provides data relating to a particular objective. Objectives are coded as shown in the key.

ated in order to allow some influence upon the coevolutionary search, whereas such penalties on objectives of lesser importance may not change or may be modified to take the objective ranking into account. The effect of varying the relative importance of the FR objective via the preference component is shown in Figure 8(a–c).

In most real decision-making situations, variables will have differing degrees of influence on any given objective. An on-line sensitivity analysis that ranks variables according to their influence on each objective has been introduced. This design sensitivity ranking is then used to adjust the fitness of each solution to ensure that the values of the most influential variables are within the range defined by the constraint map. Solutions are assigned the highest fitness penalty in cases where their most influential variables lie outside the current constraint map range. This ensures that populations contain high levels of compromise solutions in terms of the most influential variables and relatively redundant variables have little or no effect on overall solution fitness. The Taguchi method was selected to determine the sensitivity of each variable (Peace, 1992).

Again, the concentration is upon information gathering and visual representation rather than the identification of

Pareto optimal points. The coevolutionary multiobjective approach provides the following information within one run of the process:

- HP solutions relating to the individual objectives,
- evolutionary “tracks” that trace the Pareto surface to some extent,
- the bounds of a compromise region where all objectives will likely be best satisfied, and
- the identification of influential and redundant variables relating to each objective.

A more in-depth description of the process with further results can be found in Parmee and Watson (1999).

6. AGENT-BASED SUPPORT

The amount of information extracted from the evolutionary process is potentially very large and considerable attention must be paid to the manner in which it is filtered, processed, and presented to the user. The utilization of agent-based approaches was briefly introduced in previous sections; in addition, a range of single agents (Brown & Dunskus,

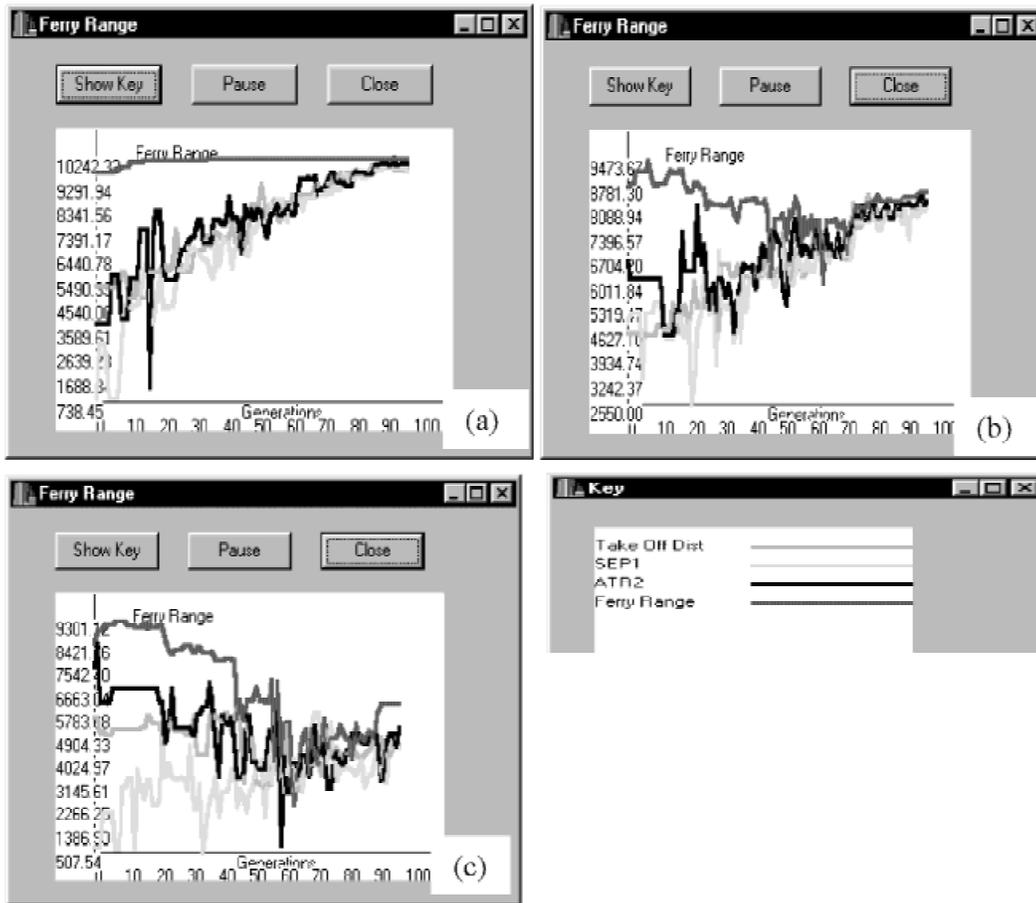


Fig. 8. The (a) ferry range is much more important. (b) All objectives are of equal importance. (c) The ferry range is much less important.

1995) were developed and integrated with the system. Stenmark (1999) classifies agents in a general sense as interface, system, advisory, filtering, retrieval, navigation, monitoring, recommender, and profiling agents. In terms of conceptual design and the IEDS, the following classes of agents appear to offer utility: *interface agents*, which help the designer deal with a system and which (if the designer wishes it) hide some low-level noninteresting details from him or her; *search agents*, which cover the process of optimization, cooperation, population monitoring, jumping out of regions, constraint questioning, and so forth; *information agents*, which deal with the information obtained, look for interesting solutions and filter uninteresting ones, make decisions with regard to what and where to explore, resolve conflicts, and so on. These agents are described in more detail in Cvetkovic (2000) and Parmee (2001). Examples of two particular aspects of agent utilization follow.

6.1. Interesting solutions

Stochastic population-based search generates a mass of information, much of which is discarded. Experimental software agents have been introduced that monitor solutions generated from single or coevolutionary processes and identify those that may be considered interesting by the decision maker. The notion of interesting may relate to, for instance:

- a good solution with a large Hamming or Euclidean distance from the majority of the population,
- a good solution that may satisfy the majority of constraints or objectives but is not satisfactory in a few, or
- a not particularly HP solution where the constituent variable values lie within user-preferred ranges.

The concept of interesting may be largely subjective, which suggests that a degree of machine learning may be appropriate, whereby the responsible agents learn from user reaction to possible interesting solutions presented to them. This is an area requiring extensive further research; however, the investigation of this machine-learning aspect is now underway with the experimental introduction of evolutionary learning classifier systems to the information gathering component. The perceived function of these systems initially relates to the generation of rules relating to variable interaction and objective interaction. Such rules, when presented to the user, should clarify the situation concerning various conflicts, providing a better understanding of what is and what is not achievable. This work is at a very early stage, but preliminary results can be found in Bull et al. (2002).

6.2. Negotiating agents

Experimental negotiating agent systems utilizing the rule-based preferences were established by Cvetkovic (2000) for the identification of solutions that satisfy a range of

design scenarios relating to multiple objectives and ideal variable values. For instance, the designer is likely to have several ideal scenarios such as: "I would like objective A to be greater than 0.6 and objective C to be less than 83.5; objectives B, D, and E should be maximized; variable 2 should have a value of between 128.0 and 164.5; a value greater than 0.32 is preferred for variable 7," and so forth.

An incremental agent operates as follows:

1. It uses the designer's original preferences for both objectives and scenarios and runs the optimization process.
2. If some scenarios are not fulfilled, the agent suggests an increase in the importance of these scenarios.
3. If some scenarios are still not fulfilled, even when classed as most important, the agent suggests a change to ideal variable ranges in the scenario.
4. If some scenarios are still not fulfilled, the agent reports to the designer and asks for assistance.

Systems have also been considered that involve several agents, each trying to optimize a single objective. Each agent is aware of the quality of his or her own solution. If agent 1's solution is inferior and contradicting to others, agent 1 should compromise and accept a worse solution to benefit the group as a whole. If agents cannot agree, the user is consulted. If the user resolves the conflict, the agents remember the decision for next time.

The incremental agent strategies have been integrated with the preferences and coevolutionary multiobjective components on an experimental basis. Both the scenario and incremental agent components sit between the designer and the preference module drawing information from both (Fig. 9). For a more detailed description of the processes and initial results based upon a MINICAPS example, the reader is directed to Cvetkovic (2000).

7. CONCLUSIONS

The objective of the paper was to introduce the potential of user interaction with evolutionary design search and exploration processes. This was achieved through the presentation of initial results from the practical implementation of a prototype IEDS and related discussion. The development of the system and its individual components was described in detail in a number of publications, all of which are referenced in the text. This paper presents an overview of the system, providing an illustration of the manner in which it may be utilized to support decision making during the early stages of product design.

COGAs are introduced as information-gathering processes that identify HP regions of a complex design space and support the analysis of such regions by ensuring good solution cover across them. COGAs are now a well-developed method, and recent developments (Bonham et al.,

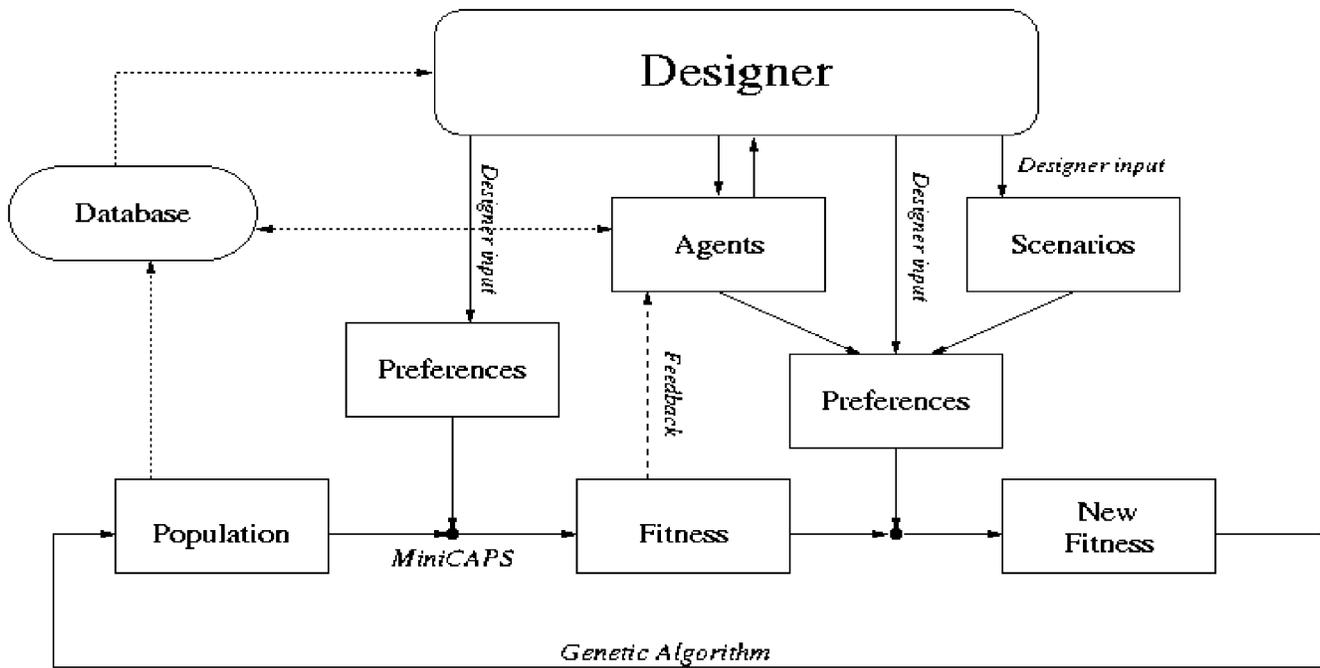


Fig. 9. The integration of negotiating agents for ideal scenario satisfaction.

1999a, 1999b) have significantly improved their performance. The included results illustrate the capabilities of the COGA approach, and the manner in which the results are presented provides a very graphic illustration of the manner in which the following can be accomplished:

- the variable space can be reduced by the setting of more appropriate variable bounds, and
- the decomposition of the space can be achieved through the selection of particular HP or most appropriate variable combinations and transforming those variables into fixed parameters.

The identification of groupings of HP solutions can also increase confidence in the validity of results from preliminary design models. Simple analysis of the solutions within each region can provide an indication of the nature of the multidimensional surface and the sensitivity of the solutions to mild perturbations of their constituent variables.

COGA strategies can also provide, in a very graphical manner, information concerning relationships among several differing objectives and the manner in which simple changes relating to their relative importance in terms of a reduction in solution performance can reduce conflicts between them. There are obvious dangers of reading too much into the various graphical representations, but background analysis of the generated solution vectors can eliminate such dangers or provide cautionary warnings to the user.

Further COGA research is required relating to front-end HCI issues concerning the processing and graphical representation of the extracted data. It is envisaged that this will

involve the development and utilization of processing and advising agents and a more in-depth study of front-end HCI issues. This area has been investigated to some extent and basic single agents have been developed, but concentration has been on the overall development of the IEDS infrastructure and more background interaction issues.

The linguistic representation approach of the preference module fits in well with the utilization of coarse, preliminary design models. The introduction of strict numeric weightings would seem inappropriate when working in uncertainty and with poorly defined domains. It would be possible to refine the preference rule set by introducing further levels of relative importance, but the degree to which this is necessary must be commensurate with the required degree of solution definition. Such a refinement should perhaps be gradually introduced as confidence in the problem domain increases. As it stands, the preference component has the potential to support further exploration of objective space, leading to radical changes in objective ranking and significant changes in design direction.

The results from the coevolutionary multiobjective processes show that the methods outlined here can provide the user with useful information from one run, although this may comprise several evolutionary processes communicating via parallel virtual machine (Breshears, 1995). Initially, each evolutionary process identifies HP solutions that provide the user with an idea of the maximum achievable results for each objective when optimized alone. The gradual introduction of penalties then draws the individual search processes into a best compromise region.

Runs that optimize two objectives can be shown to approximately traverse the Pareto front of the feasible design space. Results from experimentation involving three objectives show the ability of the method to converge on an optimal solution by approximating a Pareto surface from three different starting points (Parmee & Watson, 1999). The on-line identification of sensitive variables via Taguchi analysis aids the search process by ensuring that the most important variables have the greatest influence on the direction of the searches as they move through the design space. Experimentation suggests that the on-line sensitivity analysis has a greater role to play as the number of objectives increases.

Combining the preference component with the coevolutionary multiobjective processes illustrates the manner in which the coevolutionary search focuses on significantly differing parts of the design space depending on the relative importance of the objectives. This combined facility therefore provides another level of exploration.

The focus of the next stage of research will relate to the further integration of agent technologies that will play a major and, to an increasing extent, autonomous role to ensure appropriate communication and information processing capabilities. The overall objective of such agent-based support is to lessen the load on the designer relating to the processing of information concerning multidimensional aspects of the problem at hand. This lessening of the load is already apparent from current work relating to single negotiating agents that have a potential to provide the necessary capabilities to act in an advisory capacity while carrying out extensive analysis that remains hidden to the user. Such a capability is essential to the successful further development of the concept. The prospect of the introduction of machine learning techniques that will contribute to increasing autonomous agent activity is exciting but fraught with possible difficulties in terms of the introduction of appropriate protocols and degrees of autonomy.

The overall IEDS concept moves away from the identification of solutions through the short-term application of evolutionary search techniques. The goal is a continuous, dynamic explorative process that is primarily made possible by the search and exploration capabilities of iterative designer or evolutionary systems. Although ambitious, it is suggested that such a concept could best utilize the processing capabilities of present and future computing technology during complex human- or machine-based design and decision-making activities. The proposed architecture and IEDS concept provides an indication of what is possible in terms of interaction, information gathering, and problem reformulation relating to variables and objectives. Further research and development is required to achieve a relatively seamless development of the problem space, where the decision maker's knowledge becomes embedded within an iterative human or evolutionary computational process.

The requirement for such decision-making support is universal. Any domain where heuristics play a major role could

benefit from an appropriate interactive approach. Population-based stochastic search can provide the level of underlying nonlinear search and exploration required by an interactive system, especially when considering ill-defined, uncertain problem domains that most likely involve multiple objectives (both quantitative and qualitative), varying degrees of constraint, high dimensionality, and high modality. It is suggested that interaction can promote and achieve productive search across a changing fitness landscape that eventually results in the identification of a competitive problem solution. In this sense, the technologies can provide an excellent supporting role that enhances the decision maker's knowledge and capabilities.

ACKNOWLEDGMENTS

The author thanks the UK Engineering and Physical Science Research Council, Rolls Royce, and BAE Systems for funding this research. Several researchers have been involved over a number of years in various aspects of the work. Far greater detail of their particular contributions can be found in the texts referenced in the following section. Particular thanks go to Dragan Cvetkovic, Andy Watson, and Chris Bonham, who contributed greatly in recent years to the IEDS work.

REFERENCES

- Bonham, C.R., & Parmee, I.C. (1999a). Improving the performance of cluster-oriented genetic algorithms (COGAs). *Proc. IEEE Congress on Evolutionary Computation*, Washington DC, pp. 554–561.
- Bonham, C.R., & Parmee, I.C. (1999b). An investigation of exploration and exploitation in cluster-oriented genetic algorithms. *Proc. Genetic and Evolutionary Computation Conf.*, Orlando, FL, pp. 1491–1497.
- Breshears, C. (1995). *A Beginner's Guide to PVM—Parallel Virtual Machine*. Knoxville, TN: University of Tennessee, Joint Institute For Computational Science.
- Brown, D.C., & Dunksus, B.V. (1995). SINE: Support for single function agents. *Proc. Applications of AI in Engineering, AIENG '95*, Udine, Italy.
- Bull, L., Wyatt, D., & Parmee, I.C. (2002). Toward the use of XCS in interactive evolutionary design. *Poster Proc. Fifth Int. Conf. Adaptive Computing in Design and Manufacture*, Exeter, UK.
- Corney, D. (2000). Designing food with Bayesian belief networks. *Evolutionary Design and Manufacture: Proc. Fourth Int. Conf. Adaptive Computing in Design and Manufacture*, pp. 83–94. New York: Springer-Verlag.
- Cvetkovic, D. (2000). *Evolutionary Multi-Objective Decision Support Systems for Conceptual Design*. PhD Thesis. University of Plymouth, UK.
- Cvetkovic, D., & Parmee, I.C. (2000). Designer's preferences and multi-objective preliminary design processes. *Evolutionary Design and Manufacture: Proc. Fourth Int. Conf. Adaptive Computing in Design and Manufacture*, pp. 249–260. New York: Springer-Verlag.
- Cvetkovic, D., & Parmee, I.C. (2001). Preferences and their application in evolutionary multiobjective optimisation. *IEEE Transactions on Evolutionary Computation* 6, 42–57.
- Deb, K. (2001). *Multi-Objective Optimisation Using Evolutionary Algorithms*. New York: Wiley.
- Elby, D., Averill, R.C., Punch, W.F., & Goodman E.D. (1998). Evaluation of injection island GA performance on flywheel design optimisation. *Proc. Third Int. Conf. Adaptive Computing in Design and Manufacture*, pp. 121–136. New York: Springer-Verlag.
- Emmerich, M., Grotzner, M., Gross, B., & Schutz, M. (2000). Mixed integer evolution strategy for chemical plant optimisation with simulators. *Evolutionary Design and Manufacture: Proc. Fourth Int. Conf.*

- Adaptive Computing in Design and Manufacture*. pp. 55–68. New York: Springer-Verlag.
- Fodor, J., & Roubens, M. (1994). Fuzzy preference modelling and multi-criteria decision support. System theory. In *Knowledge Engineering and Problem Solving* (Parmee, I.C., Ed.), Vol. 14. Dordrecht: Kluwer.
- Gero, J. (1998). Adaptive systems in designing: New analogies from genetics and developmental biology. *Proc. Third Int. Conf. Adaptive Computing in Design and Manufacture*, pp. 3–12. New York: Springer-Verlag.
- Gero, J.S., Louis, S.J., & Kundu, S. (1994). Evolutionary learning of novel grammars for design improvement. *Artificial Intelligence for Design, Analysis and Manufacture* 8, 83–94.
- Gero, J.S., & Schnier, T. (1995). Evolving representations of design cases and their use in creative design. In *Preprints Computational Models of Creative Design* (Parmee, I.C., Ed.), pp. 343–368. Sydney, Australia: Key Centre of Design Computing.
- Goel, A.K. (1997). Design, analogy and creativity. *IEEE Expert, Intelligent Systems and Their Applications* 12, 62–70.
- Graf, J., & Banzhaf, W. (1995). Interactive evolutionary algorithms in design. *Proc. Artificial Neural Nets and Genetic Algorithms*, Ales, France, pp. 227–230.
- Hajela, P. (1998). Emergent computational models in structural analysis and design. *Proc. Third Int. Conf. Adaptive Computing in Design and Manufacture*, pp. 13–26. New York: Springer-Verlag.
- Herdy, M. (1997). Evolutionary optimisation based on subjective selection—Evolving blends of coffee. *Proc. Fifth Eur. Congress on Intelligent Techniques and Soft Computing (EUFIT'97)*, pp. 640–644.
- Hillermeier, C., Huster, S., Marker, W., & Sturm, T.F. (2000). Optimisation of power plant design: Stochastic and adaptive solution concepts. *Evolutionary Design and Manufacture: Proc. Fourth Int. Conf. Adaptive Computing in Design and Manufacture*, pp. 3–18. New York: Springer-Verlag.
- Inoue, T., Furuhashi, T., Fujii, M., et al. (1999). Development of nurse scheduling support system using interactive evolutionary algorithms. *Proc. IEEE Int. Conf. Systems, Man and Cybernetics (SMC'99)*, pp. 533–537.
- Koza, J.R., Bennett, F.H., Andre, D., & Keane, M.A. (1998). Evolutionary design of electrical circuits using genetic programming. *Proc. Third Int. Conf. on Adaptive Computing in Design and Manufacture*, pp. 177–192. New York: Springer-Verlag.
- Levine, D., Facello, M., & Hallstrom, P. (1997). Stalk: An interactive system for virtual molecular docking. *IEEE Computer Science Engineering Magazine* 4, 55–65.
- Maher, M.L., Poon, J., & Boulanger, S. (1995). Formalising design exploration as co-evolution—A combined gene approach. In *Preprints of the Second W/IFIPWG5.2 Workshop on Formal Design Methods*.
- Navinchandra, D. (1991). *Exploration and Innovation in Design*. New York: Springer-Verlag.
- Olhoffer, M., Arima, T., Sonoda, T., & Sendhoff, B. (2000). Optimisation of a stator blade used in a transonic compressor. In *Evolutionary Design and Manufacture* (Parmee, I.C., Ed.), pp. 45–54. London: Springer-Verlag.
- O'Reilly, M., & Testa, P. (2000). Representation in architectural design tools. *Evolutionary Design and Manufacture: Proc. Fourth Int. Conf. Adaptive Computing in Design and Manufacture*, pp. 137–148. New York: Springer-Verlag.
- Parmee, I.C. (1966a). The maintenance of search diversity for effective design space decomposition using cluster-orientated genetic algorithms (COGAs) and multi-agent strategies (GAANT). *Proc. Second Int. Conf. on Adaptive Computing in Engineering Design and Control, PEDC* (Parmee, I.C., Ed.), pp. 128–138. University of Plymouth.
- Parmee, I.C. (1996b). The development of a dual-agent strategy for efficient search across whole system engineering design hierarchies. *Proc. Parallel Problem Solving from Nature (PPSN IV), Lecture Notes in Computer Science No. 1141*, pp. 523–532. Berlin: Springer-Verlag.
- Parmee, I.C. (2001). *Evolutionary and Adaptive Computing in Engineering Design*. London: Springer-Verlag.
- Parmee, I.C., & Bonham, C.R. (1999). Towards the support of innovative conceptual design through interactive designer/evolutionary computing strategies. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing* 14, 3–16.
- Parmee, I.C., Cvetkovic, C., Watson, A.H., & Bonham, C.R. (2000). Multi-objective satisfaction within an interactive evolutionary design environment. *Journal of Evolutionary Computation* 8, 197–222.
- Parmee, I.C., Cvetkovic, C., Watson, A.H., Bonham, C.R., & Packham, I. (2001). Introducing prototype interactive evolutionary systems for ill-defined design environments. *Journal of Advances in Engineering Software* 32, 429–441.
- Parmee, I.C., & Vekeria, H.D. (1997). Co-operative evolutionary strategies for single component design. *Proc. Seventh Int. Conf. Genetic Algorithms*, Michigan State University.
- Parmee, I.C., & Watson, A.W. (1999). Preliminary airframe design using co-evolutionary multi-objective genetic algorithms. *Proc. Genetic and Evolutionary Computation Conf.*, Orlando, FL, pp. 1657–1665.
- Peace, G.S. (1992). *Taguchi Methods*. Reading, MA: Addison-Wesley.
- Roy, R., & Parmee, I.C. (1996). Adaptive restricted tournament selection for the identification of multiple sub-optima in a multi-modal function. In *Lecture Notes in Computer Science, Evolutionary Computing*, pp. 236–256. New York: Springer-Verlag.
- Roy, R., Parmee, I.C., & Purchase, G. (1996a). Integrating the genetic algorithm with the preliminary design of gas turbine cooling systems. *Proc. Second Int. Conf. Adaptive Computing in Engineering Design and Control, PEDC* (Parmee, I.C., Ed.), pp. 228–235. University of Plymouth.
- Roy, R., Parmee, I.C., & Purchase, G. (1996b). Qualitative evaluation of engineering designs using fuzzy logic. *Proc. ASME Design Engineering Technical Conf. and Computers in Engineering Conference*, Irvine, CA, 96-DETC/DAC-1449.
- Schnier, T., & Gero, J.S. (1998). From Mondrian to Frank Lloyd Wright: Transforming evolving representations. *Proc. Third Int. Conf. Adaptive Computing in Design and Manufacture*, pp. 207–220. New York: Springer-Verlag.
- Shiraki, H., & Saito, H. (1996). An interactive image retrieval system using genetic algorithms. *Proc. Int. Conf. Virtual Systems and Multimedia (VSMM'96)*, pp. 257–262.
- Sims, K. (1991a). Artificial evolution for computer graphics. In *Computer Graphics, ACM SIGGRAPH Proceedings, Vol. 25*, pp. 319–328.
- Sims, K. (1991b). Interactive evolution of dynamical systems. *First Eur. Conf. Artificial Life*. Cambridge, MA: MIT Press.
- Stenmark, D. (1999). Evaluation of intelligent software agents. Available on-line at: <http://w3.informatik.gu.se/~dixi/agent/agent.htm>.
- Su, N.P. (1990). *The Principles of Design*. New York: Oxford University Press.
- Webb, E. (1997). *MINICAPS—A Simplified Version of CAPS for Use as a Research Tool*. Unclassified Report BAE-WOA-RP-GEN-11313, British Aerospace Plc.
- Zitzler, E., Deb, K., Thiele, L., Coello, C.A., & Corne, D. (2001). Evolutionary multi-criterion optimisation. *Proc. First Int. Conf. Evolutionary Multi-Criterion Optimisation, Lecture Notes in Computer Science, No. 1993*. New York: Springer-Verlag.

Ian Parmee has several years' experience in both the contracting and consultancy sectors of the civil engineering industry. He returned to an academic career in 1991, playing a major role in the development of Plymouth University's EPSRC Engineering Design Centre, which investigates the integration of evolutionary computing technologies with engineering design. He has now joined the University of the West of England, Bristol, where he is currently establishing research in the area of evolutionary design and decision making. His research has resulted in over 100 publications in journals, conference proceedings, and books.