Pergamon

Int. J. Heat Mass Transfer. Vol. 37, No. 6, pp. 893–908, 1994

0017-9310/94 $6.00 + 0.00

0017-9310(93)E0020-H

Genetic algorithms for thermosciences research: application to the optimized cooling of electronic components

N. QUEIPO, R. DEVARAKONDA and J. A. C. HUMPHREY†

Computer Mechanics Laboratory, Department of Mechanical Engineering, University of California at Berkeley, Berkeley, CA 94720, U.S.A.

(*Received* 7 *June* 1993 *and in final form* 18 *October* 1993)

Abstract—Genetic algorithms are adaptive search procedures loosely based on the Darwinian notion of evolution. They use rules of natural selection to investigate highly complex, multidimensional problems and have been employed successfully in a variety of search, optimization and machine learning applications in science and engineering where other more traditional methods fail. In this study genetic algorithms are presented and discussed within the framework of an adaptive solution methodology for investigating otherwise intractable optimization problems in the thermosciences. The exposition focuses on their application to an electronics cooling problem where it is required to find optimal or nearly optimal arrangements of convectively cooled components placed in-line on the bottom wall of a ventilated two-dimensional channel. The present application is specific only for purposes of illustration. The power of the methodology rests on its generality of application and an indifference to the source of data (experimental, analytical or numerical) used in the optimization process. The study shows that genetic algorithms allow a cost-effective approach for investigating highly complex numerical and/or experimental thermosciences problems where it is desirable to obtain a family of *acceptable* problem solutions, as opposed to a single optimum solution.

1. INTRODUCTION

IN THE thermosciences it is frequently the case that relatively complex flow and heat transfer problems can be accurately formulated and often solved by means of closed form analysis or numerical computation. *Accuracy of formulation* implies a correct physico–mathematical representation of the phenomena investigated in terms of the equations that describe the conservation laws, constitutive relations, equations of state, boundary and initial conditions, etc. *Accuracy of solution*, whether analytical or numerical, implies mathematically correct manipulation and calculation of the equations representing the physical phenomena. When fundamental understanding and, therefore, physico–mathematical representation are incomplete, or when the analytical/numerical tools available are not equal to the mathematical complexities of the solution task, experimentation may yield partial, sometimes complete, solutions of the problems. In this case, in addition to accuracy it is indispensable to maintain the precision required to minimize random uncertainties in the measurements. In the present numerical study it is presumed that precision (in the sense associated with experimentation) is not a source of calculation uncertainty and that it can be maintained at the level required. In addition, special attention is paid to the accuracy of problem formulation, presuming that computational speed and memory requirements can be met.

There is a growing class of thermosciences problems that can be accurately formulated and solved, or precisely experimented, for which it can be stated a priori that, except for trivial or rare circumstances, the optimal solutions cannot be found with a reasonable amount of work and in a reasonable amount of time. This is because the multidimensional nature and ranges of the associated variables and parameter spaces, and the non-linearities embedded in the mathematical representations of the problems, combine to render impossible (in a practical sense) an exhaustive search for optimal solutions. In contrast, with considerable less effort many of these problems yield *nearly optimal* solutions which come close to minimizing a predetermined cost function, or maximizing a corresponding performance measure, while simultaneously satisfying a set of imposed constraints.

We are concerned here with developing a methodology that will facilitate the search for optimal or nearly optimal solutions of complex multidimensional thermosciences problems admitting accurate formulation and calculation, and/or precise experimentation. While this is a fundamental activity that could be couched in highly abstract and complicated mathematical terms, we have kept practical engineering applications in mind and, instead, have developed and tested some interconnected concepts for purposes of illustration. The activity communicated in this paper is part of an ongoing study on Interactive Computational–Experimental Meth-

† Author to whom correspondence should be addressed.

NOMENCLATURE

A_i	exponent in equation (1) for component 'i' [K]
B_i	preexponential factor in equation (1) for component 'i' [fr Mh^{-1}]
C_0	scaling constant in equation (3)
f	cost function in equation (2)
f_{max}	maximum value of f in a generation
F	fitness function or performance measure given by equation (2)
FR	total failure rate, $\Sigma\lambda_i$ [fr Mh^{-1}]
g	gravitational acceleration, 9.8 m s^{-2}
Gr	Grashof number, $g\beta(T_{max}-T_0)L1^3/\nu^2$
h_{ave}	component surface-averaged heat transfer coefficient [W m^{-2} K^{-1}]
$h1$	height of small component [cm]
$h2$	height of large component [cm]
H	channel height [cm]
k	component thermal conductivity [W m^{-1} K^{-1}]
$L1$	component length [cm]
$L2$	separation between components [cm]
Ld	distance between channel exit plane and the downstream surface of the last component [cm]
L_{ij}	wiring length between components 'i' and 'j' [cm]
Lu	distance between channel inlet plane and the upstream surface of the first component [cm]
N	number of components in an in-line arrangement
$P1$	parent one
$P2$	parent two
$P(t)$	current population at time t
q_i	heat flux from component 'i' [W m^{-2}]
Re	Reynolds number, UH/ν
t	time [s]
T_i	maximum surface temperature of component 'i' [K]
T_0	uniform temperature of inlet flow [K]
T_{max}	maximum T_i [K]
u	velocity component in x direction [m s^{-1}]
v	velocity component in y direction [m s^{-1}]
U	uniform velocity of inlet flow [m s^{-1}]
$x1$	first location for crossover operator
$x2$	second location for crossover operator
x	streamwise coordinate
y	transverse coordinate
X, Y	chromosomes (also called strings or candidate solutions) denoting arrangements of N in-line components.

Greek Symbols

β	coefficient of volume expansion, $1/T_0$ [K^{-1}]
λ_i	failure rate of component 'i' given by equation (1) [fr Mh^{-1}]
ν	kinematic viscosity [m^2 s^{-1}]
Σ	summation sign.

odologies (ICEME) discussed in depth by Humphrey *et al.* [1].

Because of its fundamental and practical importance, we have chosen to develop and illustrate the methodology in connection with the thermofluids-optimized packaging of electronic components. According to Weiss *et al.* [2], advances in circuit integration and packaging have increased the power dissipation of typical electronic components on circuit boards from 5 to 50 kW m^{-3}, approximately, over the past few years. As a result, the potential for premature thermally-induced failures of electronic components has also increased with an attendant impact on cost. For example, according to the Department of Defense it is estimated that a 5°C reduction in the temperature of the air for cooling aircraft electronics could save 10 million dollars annually in service and maintenance costs (Weiss *et al.* [2]). Thus, the achievement of cost-effective packaging designs is critical to the electronics industry.

Objective of this investigation

In searching for efficient ways to manage thermal dissipation, both researchers and practitioners have found that the heat transfer from electronic components is very sensitive to variations in geometry when convection is the dominant heat transfer mechanism (Azar *et al.* [3]). As a result, the thermofluids-optimized placement of heated components in ventilated enclosures, while simultaneously satisfying geometical, electrical, manufacturing and other problem constraints, has also become an activity of considerable interest (and urgency) among thermosciences engineers.

The specific objective of the present study can be abstractly stated as follows: for a given convective flow configuration, find the optimal or nearly optimal values for a set of variables and/or their related parameters that mimimize a cost function, or maximize a corresponding performance measure, while satisfying all imposed problem constraints. For example, the variables could be: the velocity, presure and temperature fields at specific locations; the geometrical shape, relative locations and orientations of interconnected components; the shear stresses and heat fluxes at the component surfaces. Related parameters could be: the Reynolds, Grashof, Prandtl and Nusselt numbers; the form drag and skin friction coefficients;

other dimensionless groups involving geometry length scales. The performance measure to be maximized is usually a complicated multidimensional function of quantities like the above and, often, it is multimodal, possessing several local maxima or minima. For the non-linear thermosciences problems of interest to this work an experiment could be performed to measure the variables and parameters used to guide the optimization process or, as in the present case, a set of transport equations for mass, momentum, and energy, together with the appropriate boundary and initial conditions, can be solved.

The *practical* intractability of this class of problems is due to the typically large number and ranges of the variables and parameters that characterize the problems, and to the sensitivity of solutions to variations in boundary conditions and parameter values. The sensitivity arises as a result of the non-linearities embedded in the equation set describing the problems. The multimodal nature of the performance measure function renders ineffective most conventional optimization techniques, such as hill-climbing algorithms when they apply, which fail because of their local nature (Jacoby *et al.* [4]). Statistical sampling techniques do not suffer from the local nature of the hill-climbing algorithms but their computation times grow rapidly with the dimensionality of the problem (Hill [5]). The state-of-the-art of solution methods for the kinds of optimization problems of interest here involves forms of random searches which make little (if any) use of the information available prior to, or acquired during, the search process.

We claim that in order to effectively obtain optimal or nearly optimal solutions of complex multidimensional thermosciences problems it is necessary to implement an *adaptive problem-solving methodology*. By this we mean a strategy that dynamically accumulates information and uses it to improve problem-solving performance. In this study we present an adaptive solution methodology based on *genetic algorithms*. These are search procedures loosely based on the Darwinian notion of evolution and rules of natural selection. The methodology is illustrated by reference to the problem of finding optimal or nearly optimal arrangements of convectively cooled heat sources, representing electronic components, on the bottom wall of a ventilated two-dimensional channel. This is similar to the problem investigated analytically by Dancer and Pecht [6] using a lumped formulation or 'tank-and-tube' approach. In contrast, in this study we pay close attention to resolving space- and time-variations of the velocity and temperature fields in order to satisfy the condition mentioned earlier concerning accuracy of problem formulation. The present problem solution approach is numerical in nature and is based on the finite difference transient flow calculation procedure described by Schuler *et al.* [7].

It is important to note that, because the genetic algorithm approach is independent of the source of information upon which the optimization is based, what has been accomplished here using numerical solutions of a physico–mathematical model can be as readily achieved by reference to a properly designed and carefully executed experiment. Thus, the methodology developed is equally applicable to numerical and experimental sources of data. In fact, Humphrey *et al.* [1] point out that in some difficult cases there may be considerable advantages to combining complementary experimental and numerical data bases for the purpose of accelerating problem solution.

2. THE ADAPTIVE SOLUTION METHODOLOGY

The present adaptive solution methodology operates in much the same way as a standard feed-back loop in classical control theory. In this representation a complex process is connected to an adaptive strategy via a feedback loop. In this work the process is the thermofluids configuration of interest, represented by a physico–mathematical model which allows detailed calculations of its characteristics. This includes the conservation laws, constitutive relations, equations of state, boundary conditions and any other relevant information. In an application involving measurements the process would correspond to the real physical system, that is, the experiment. Subject to the input data and control variables provided, the process generates output data and performance measures. In the present numerical application the input data are an initial set of known calculation conditions. In an experiment the input data might also include variables or parameters over which there is little or no control. The output data consist of values for field variables such as velocity, pressure and temperature, and associated quantities such as form drag and skin friction coefficients and Nusselt numbers which are used to calculate the performance measures of the solutions obtained.

Part of the output from the complex process provides the input to the adaptive strategy used to optimize problem solving. The adaptive strategy is the procedure responsible for the dynamical accumulation of decision-making information through the feedback portion of the loop. It also generates the control actions that modify a current set of variables or parameters with the expectation that the changes will improve the process performance. Following De John [8], in the present methodology a given class of genetic algorithms is used as the adaptive strategy.

3. OUTLINE OF GENETIC ALGORITHMS

The genetic algorithm (GA) requires an initial set of possible solutions to commence the search for optimal solutions. Through a process based on concepts taken from evolution, using rules of natural selection the GA improves upon these solutions. In any calculation cycle, the set of candidate solutions at time t, $P(t)$, operated upon by a GA is called the *population* and

each member of this set or *generation*, when encoded as a string of symbols, is called a *chromosome*. Originally pioneered by Holland [9], a GA may be abstractly represented by the following sequence of operations:

$t = 0$;
Initialize $P(t)$;
Evaluate $P(t)$;
while (termination condition not satisfied) do
begin
$t = t+1$;
Select $P(t)$;
Recombine $P(t)$;
Evaluate $P(t)$;
end.

In this representation each iteration in the 'while' loop produces a new generation of candidate solutions, also encoded as chromosomes. The expectation is that if a set of candidate solutions is properly encoded, and the 'Select' procedure and the GA operators of the 'Recombine' procedure are appropriately chosen, each generation of *parents* will produce a generation of *children* (the new set of candidate solutions) which, in general, will have an average performance better than the parent generation. It is the role of the GA operators to construct and propagate the features or *schema* of those chromosomes responsible for the improved performance of some candidate solutions relative to others. A schema is like a similarity template, serving to reveal the subset of chromosomes possessing similarities at certain chromosome positions (Goldberg [10]). The schemata derived from good chromosome solutions within a generation provide the *building blocks* from which to synthesize improved solutions in the offspring generation.

In a typical GA, the initial set of candidate solutions is usually selected randomly. There are no definite rules regarding the desired initial population size for a given problem but guidelines are provided by Grefenstette [11]. The candidate solutions are encoded as fixed-length chromosomes for which different encoding schemes, such as binary and integer coding, have been used. For example, a chromosome indicating the sequential positions of four identical heated components distributed among ten possible locations in an enclosure could be represented as the string 0011001010 where a '0' represents an empty location and a '1' an occupied location. Generally, it is to be expected that the heat transfer characteristics of the physical arrangement represented by this chromosome will differ from other arrangements given by, for example, 1010100010, 1001100100, etc. Further details concerning the usefulness and effectiveness of chromosome encoding in binary notation are available in Goldberg [10].

The purpose of the 'Evaluate' procedure in the sequence is to calculate the *fitness* of each chromosome. The fitness is a measure of performance associated with each candidate solution. It is a very important quantity since the probability that a chromosome in the parent population will contribute its schema to the offspring generation is proportional to the chromosome's relative fitness. Thus, a chromosome representing a candidate solution with a high measure of fitness according to some pre-established criterion should have a high probability of parenting children in the next generation of candidate solutions. The function of the 'Select' procedure is to specify the actual number of offspring that each parent chromosome contributes to the following generation based on the relative performance of that chromosome. For a discussion of different selection mechanisms see Baker [12].

The 'Recombine' procedure contains the GA operators that are expected to construct and propagate the schema responsible for good peformance. The most prominent GA operators are *crossover* and *mutation*. The crossover operator acts on two chromosomes at a time, on average generating fitter offspring by combining the schema in each parent. The mutation operator usually involves randomly (but infrequently) altering the value of one or more bits in a chromosome. The crossover and mutation operators for binary chromosomes are discussed at length by, for example, Goldberg [10].

While binary notation is optimal for the chromosome encoding of candidate solutions for many types of problems it may be inappropriate for others (Oliver *et al.* [13]). There are certain cases where the symbols in the encoded solution may not be repeated and the operations of the GA operators must be suitably adapted. An example of this of special interest here is the encoding of different sequences of N distinct objects arranged in-line. In this case integer encoding is used and special GA operators are necessary. In integer encoding each of the N objects is assigned a unique integer in the range 1 to N, and any string of these integers denotes a particular sequence or arrangement. Since the objects are distinct, repeated integers are forbidden in this encoding scheme. The result is that the crossover operator must be of a permutation or reordering type.

Goldberg [10] discusses in detail three basic types of permutation or reordering operators. These are: (i) partially matched crossover; (ii) order crossover; and, (iii) cycle crossover. The operators are similar in terms of their actions and only the partially matched crossover (PMX) operator used here is described. With reference to Table 1, two chromosomes (X, Y) representing candidate solutions or sequences of $N = 8$ objects are aligned one over the other. A matching section consisting of a sequence of overlapping positions common to both strings is identified by randomly selecting two crossover sites. In Table 1 the matching section runs from position III to position V. As a result, elements 6, 5 and 8 in string X are paired with 1, 4 and 3 in string Y. The rule is that elements 6, 5 and 8 *within each string* should exchange places respectively with elements 1, 4 and 3 *within*

Table 1. Genetic operators for integer string representation of chromosomes: (a) partially matched crossover (PMX) operator, (b) mutation operator. See discussion in text

Candidate Solution	Position I	II	III	IV	V	VI	VII	VIII	
X	1	7	6	5	8	3	4	2	
Y	8	2	1	4	3	6	5	7	
			↑		↑				PMX (a)
X'	6	7	1	4	3	8	5	2	
Y'	3	2	6	5	8	1	4	7	
			↓			↓			
X'	6	7	1	4	3	8	5	2	
X''	6	7	8	4	3	1	5	2	mutation (b)

the same string. The resulting (X', Y') chromosomes represent new sequences of the same objects and it is exactly as if they had been generated as the result of the exchange of chromosome material. The expectation is that over a few generations strings generated this way will gradually evolve schema representing desirable features. Table 1 illustrates the corresponding mutation operator on the X' chromosome. This involves randomly selecting two positions, III and VI, in this string and exchanging their corresponding elements. The result is the altered X'' chromosome. Because crossover and mutation can also disrupt desirable schema, it is important to specify appropriate settings for these two operators in a GA (Goldberg [10] and Davis [14]).

In a GA application, the Select, Recombine and Evaluate processes just described are repeated from generation to generation until some convergence or termination criterion is met.

4. PROBLEM DEFINITION, PROCEDURES AND TESTING

4.1. *Case study problem*

The flow configuration investigated is shown in Fig. 1. The task is to find optimal or nearly optimal arrangements that minimize the total thermal failure rate of N in-line convectively cooled electronic components in an air-ventilated channel. This is a combinatorial minimization problem not suited to gradient based optimization techniques. A simpler (analytically tractable) version of this problem, based on a lumped formulation approximation, was previously investigated by Dancer and Pecht [6]. Because we do not invoke this approximation, the present problem does not admit a closed form solution and, instead, must be calculated numerically. The flow configuration serves as a model for the placement of chips on a board or for the location of larger electronic components inside a ventilated enclosure. Thus, with reference to Fig. 1, we ask: given N heated components of rectangular cross-section to be distributed among N equally-spaced locations on one wall of a ventilated two-dimensional channel, what are some of the arrangements of the components, encoded as integer strings, that maximize a given performance function while satisfying all imposed boundary conditions and known system constraints?

To answer this question we note that the failure rate of an individual electronic component usually depends strongly on temperature and is assumed to vary according to the Arrhenius equation for the specific rate constant of a chemical reaction. Thus, calling λ_i the failure rate per megahours (fr Mh^{-1}) of an electronic component i, we take

$$\lambda_i = B_i \exp^{-A_i/T_i}. \tag{1}$$

In equation (1) T_i is the maximum temperature of component 'i'. The quantities B_i and A_i are the analogs to the frequency (or preexponential) factor and activation energy of reaction, respectively. In the present application these two coefficients are assumed to be independent of temperature but may vary among component types and must be determined experimentally or taken from data available in the literature.

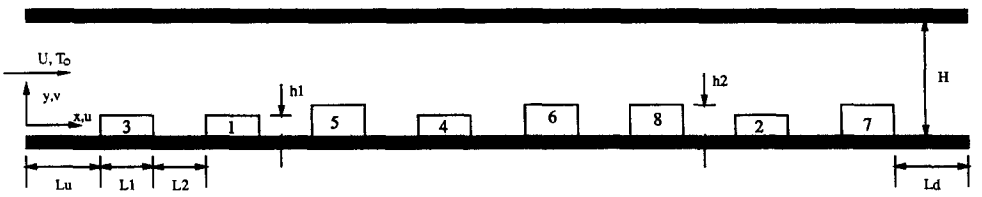

FIG. 1. Schematic of the thermofluids problem of interest showing eight heated components randomly placed at equally spaced locations on the bottom wall of an air-ventilated two-dimensional channel. The conditions are: $U = 0.5$ m s^{-1}, $T_0 = 300$ K, $H = 2$ cm, $h1/H = 0.1$, $h2/H = 0.2$, $L1/H = 1$, $L2/H = 1$, $Lu/H = 2$, $Ld/H = 4.5$. The figure is not drawn to scale.

A quantity f must be defined which represents the cost function to be minimized. However, because the GA works to maximize a non-negative performance measure or fitness function, a mapping of the cost function, f, to a fitness function, F, is required. The transformation used in this study is

$$F = f_{\max} - f \tag{2}$$

where $f_{\max}$ is the largest value of f associated with one of the chromosomes in the current population. This function has the desirable property that it apportions higher values of fitness to chromosomes with lower cost functions. In a typical thermal application f is a function of the λ_i for which concrete expressions are provided in the examples discussed below.

For the special case of geometrically identical components, the optimal solution for the present Case Study Problem is trivial in two cases *if the analysis is based on a lumped formulation approach* such as performed by Dancer and Pecht [6]. Such an analysis predicts that: (i) if the thermal sensitivities, B_i, of the components are all the same but the components differ in their heat fluxes, then they should be arranged in an increasing order of their heat fluxes; and, (ii) if, however, the heat fluxes of the components are all the same but they differ in their thermal sensitivities, then they should be arranged in a decreasing order of their thermal sensitivities. (While these two 'rules of thumb' have proven useful in the electronics industry, we will show below that they can lead to non-optimal arrangements for conditions where the lumped formulation approach is inapplicable.)

Of interest here is the more general and much more difficult case when the heated components differ in size, thermal sensitivity and heat transfer characteristics. For example, finding the optimal solution, by enumeration, for a configuration involving $N = 8$ different components would alone require $8! = 40\,320$ distinct numerical simulations (or the equivalent amount of experimentation) to determine the characteristics of all the possible component arrangements. Clearly, even within the context of a numerical approach such as the present one, this represents a prohibitively large calculation effort. As we will show, the effort can be reduced substantially using a GA to guide the search for optimal solutions.

4.2. *The genetic algorithm code*

The GA code used in this study is a substantially modified version of the SGA (Simple Genetic Algorithm) code presented in Goldberg [10]. The main differences between the SGA and the present code, MSGA (Modified-SGA), are the following. First, the Case Study Problem described above requires a permutation representation for the chromosomes instead of the binary representation provided in SGA. As a result, the GA operators in the 'Recombine' procedure had to be redefined. Specifically, the crossover operator in the MSGA is the so called partially matched crossover of Goldberg and Lingle [15] instead of the one-point crossover operator used in the SGA. Additionally, the mutation operator was redefined so that, when mutation is specified, two positions in the permutation representation of the chromosome are selected and their values exchanged. Second, the MSGA employs an *elitist strategy* in its 'Select' procedure. In this strategy the fittest chromosome in a current generation is automatically passed on *unchanged* to the new generation where it replaces the least fit chromosome in that generation.

The 'Select' procedures in both the MSGA and the SGA use a weighted random sampling technique discussed by Goldberg [10] and Davis [14] as the reproduction mechanism. Essentially, this technique is equivalent to spinning a roulette wheel that is divided into a number of sectors equal to the number of chromosomes in a generation, with the size of each sector proportional to the fitness of the chromosome it represents. In this way, the probability of selecting a particular chromosome for reproduction is proportional to its fitness.

The MSGA requires the following input parameters: population size, chromosome length, number of generations, crossover and mutation rates. A report is generated when the number of generations specified by the user has been reached. The report contains two types of information: (i) a record of the input values introduced by the user and some statistics related to the initial population; (ii) a specification of the chromosomes, their fitness values, crossover-related information and relevant statistics for each generation.

Part of a sample report generated by the MSGA is provided in Table 3, discussed in more detail further below. The report shows, for instance, that in Generation 2 the second string 53178462 has a fitness value $F = 0.00045$ corresponding to a failure rate of $FR = 0.01179$ fr Mh^{-1}. This chromosome had strings #6 and #7 from Generation 1 as its parents. The parents were subjected to a partially matched crossover operator between positions 1 and 2 as indicated in columns '$x1$' and '$x2$'. The resulting offspring chromosome was subsequently subjected to a mutation between two randomly selected locations (not indicated in the table). In these reports, if a chromosome (such as string #1, 84657321, in Generation 2) is preceded by values of 0 in the parents and crossover columns, this indicates that it was passed on unchanged from the parent generation to the offspring generation as a result of the elitist strategy.

4.3. *The field calculation method*

In order to evaluate the fitness values of the integer strings representing solutions of the Case Study Problem we must perform measurements or calculations. In either approach, the impact of recirculating flow regions on fluid motion and heat transfer must be resolved. In this study fluid motion and heat transfer were calculated numerically using the fully elliptic, time-marching CUTEFLOWS algorithm of Schuler

et al. [7]. The numerical procedure is based on spatially second-order accurate finite-difference approximations of the conservation equations derived using a control volume integration approach. A set of ordinary differential equations for momentum and energy is solved in time using a second-order accurate Runge–Kutta predictor–corrector method (RK2).

To satisfy continuity and simultaneously obtain the pressure field, at each new time step the expected divergence-free velocity field is decomposed into a pseudo-velocity field that is independent of pressure and a separate contribution that is due to pressure. Using the RK2 algorithm, the pseudo-velocity is computed directly. The pressure contribution is defined by the discrete Poisson equation that results from the imposition of the divergence-free condition on the velocity field. This contribution is calculated at the end of each half-time step using the conjugate gradient method. With the velocity field known it is a simple matter to calculate temperature in a forced convection flow. When coupling arises between momentum and energy as a result of thermally-induced buoyant motions, the procedure described by Iglesias *et al.* [16] is employed. In this way three (in the present case two) velocity components, pressure and temperature are obtained on a staggered grid. The calculation procedure has undergone rigorous testing in several two and three-dimensional unsteady flows described in Schuler *et al.* [7], Iglesias *et al.* [16], Schuler [17], Treidler [18], Humphrey *et al.* [19] and Tatsutani *et al.* [20, 21].

For the present calculations, constant physical properties of air at 300 K were used. As shown in Fig. 1, the two-dimensional calculation domain was bounded at the top and bottom by fixed channel walls. No slip, impermeable wall boundary conditions for the u and v velocity components were imposed along all solid surfaces. At the channel inlet $u = U$ and $v = 0$ were uniformly prescribed, corresponding to a plug flow velocity profile. At the outlet we set $v = 0$ and used the wave equation boundary condition, employed by Arnal *et al.* [22], to calculate u. The convection–diffusion form of the energy equation was solved for temperature. For this, a uniform value of $T = 300$ K was fixed at the inlet plane and the wave boundary condition was specified at the outlet. The top and bottom walls of the channel were prescribed as adiabatic. A constant heat flux was specified along each of the three exposed surfaces of each heated component. The values of the component heat fluxes are listed in Table 2 together with other relevant parameters. All the calculations were peformed with $U = 0.5$ m s^{-1}, yielding $Re = 630$ for the value of the Reynolds number.

The present study is primarily a proof of concept investigation. Buoyancy contributions to the balance of momentum were small and for expediency they were neglected in the bulk of the calculations. In addition, a relatively coarse non-uniform computational grid consisting of ($x = 90$, $y = 22$) nodes was used. The calculation time step was 0.001 s. The grid non-uniformity allowed a more accurate resolution of the flow and temperature fields around the components. While insufficient to claim grid-independent results, this refinement was more than adequate for demonstrating the GA-based adaptive solution methodology.

4.4. *Validation run*

Before considering the general case involving heated components of arbitrary size, thermal sensitivities and heat fluxes, the *performance correctness* of the MSGA was established by reference to a case for which the arrangement of components minimizing their total failure rate was known a priori. In the present validation run, $N = 8$ heated components with the characteristics listed in Table 2 are to be arranged in the channel shown in Fig. 1 so that the cost function f, defined as $f = \Sigma\lambda_i$, is minimized or, equivalently, the fitness function F given by equation (2) is maximized. In the expression for f the summation is over all components $i = 1, N$. As shown in the table, the components have identical heat transfer and geometrical characteristics but differ in their thermal sensitivities, B_i. We know from a lumped formulation analysis for identical components with identical heat fluxes that to minimize their total failure rate they should be arranged in *decreasing order* of thermal sensitivity and this is what we wish to show.

A few preliminary runs of the MSGA were conducted in order to select appropriate values for the crossover and mutation rate GA parameters. The fitness value for each candidate solution string was calculated using equation (2). Except for their thermal sensitivities, the components were identical in every respect for this validation run. Thus, it was necessary to execute the numerical algorithm only once to obtain the velocity and temperature fields. The maximum temperature for each component as a function of its position in the arrangement was provided as part of the input to the MSGA. (We note that, due to the constant heat flux condition imposed on the three surfaces of any component, the maximum temperature of a component must arise on one of its three surfaces. For the conditions of this work, this maximum always occurred on the vertical downstream surface of a component, in the inter-component recirculating flow region near the bottom wall of the channel.)

Table 3 presents a summary of the results obtained using the MSGA with a population size equal to 7, the number of generations equal to 7, and crossover and mutation rates set to 0.6 and 0.1, respectively. The results show that the MSGA uncovers *one of several possible* optimal arrangements of the components, chromosome 47362815, by Generation 7. Thus, within seven generations the MSGA has captured one of several possible schema reflecting a decreasing order of thermal sensitivity which, as pre-

Table 2. Geometrical and thermal characteristics of the optimally arranged components in the validation run and example cases. All components were assigned the value $A_i = 4635$ K

	Validation run			Examples		
Component	Height (cm)	q_i (W m^{-2})	B_i (fr Mh^{-1})	Height (cm)	q_i (W m^{-2})	B_i (fr Mh^{-1})
1	0.2	200	200	0.2	200	400
2	0.2	200	800	0.2	300	400
3	0.2	200	2000	0.2	200	1600
4	0.2	200	2000	0.2	300	1600
5	0.2	200	200	0.4	200	400
6	0.2	200	800	0.4	300	400
7	0.2	200	2000	0.4	200	1600
8	0.2	200	800	0.4	300	1600

Table 3. MSGA report corresponding to the validation run: (a) input parameters for the MSGA; (b) initial population statistics; (c) detailed results for Generations 1 and 2; (d) detailed results for Generations 6 and 7; (e) statistics for Generations 2 and 7

MSGA input parameters	Value
Population size	7
Chromosome length	8
Maximum generations	7
Crossover probability	0.6
Mutation probability	0.1

(a)

Initial population statistics	Value
FR_{max}	0.01228
FR_{min}	0.01141
FR_{avg}	0.01179

(b)

Generation 1

String	Fitness	FR	$P1$	$P2$	$x1$	$x2$
52368174	0	0.01228	0	0	0	0
25136478	0.00001	0.01227	6	7	1	2
84657321	0.00087	0.01141	4	7	1	4
26348517	0.00069	0.01159	4	7	1	4
57864321	0.00065	0.01163	4	7	4	6
53428617	0.00062	0.01165	4	7	4	6
81375462	0.00055	0.01173	3	4	2	8

Generation 2

String	Fitness	FR
84657321	0.00087	0.01141
53178462	0.00045	0.01179
81372564	0.00037	0.01188
26345718	0.00080	0.01144
26374815	0.00107	0.01118
41358762	0.00090	0.01134
15348672	0.00015	0.01210

(c)

Generation 6

String	Fitness	FR	$P1$	$P2$	$x1$	$x2$
27314865	0.00011	0.01107	2	3	3	4
46372815	0.00055	0.01063	2	3	3	4
27364815	0.00099	0.01092	3	3	0	0
26374815	0	0.01118	3	3	0	0
26374815	0	0.01118	2	3	2	7
26473815	0	0.01118	2	3	2	7
26374815	0	0.01118	3	1	1	6

Generation 7

String	Fitness	FR
47362815	0.00081	0.01037
26374815	0	0.01118
27364815	0.00026	0.01092
27364815	0.00026	0.01092
26374815	0	0.01118
47362815	0.00081	0.01037
27314865	0.00011	0.01107

(d)

Statistics	Generation 2	Generation 7
FR_{max}	0.01210	0.01118
FR_{min}	0.01118	0.01037
FR_{avg}	0.01159	0.01086

(e)

dicted by the lumped formulation analysis, is the feature responsible for good performance. By Generation 7 other, less optimal, candidate solution chromosomes also show evidence of this desirable feature in their respective schema.

The computational effort in terms of the number of fitness function evaluations required for obtaining an optimal sequence using the MSGA for the validation run was 49. This number is smaller, by a factor of about 11, than the number of expected function evaluations required by a purely random search strategy. This is because the probability of finding any one of the optimal sequences randomly is $3! \times 3! \times 2!/8!$ (or 1/560) compared to 1/49 for the present MSGA application. Thus, relative to a random search strategy the MSGA is more cost effective in its ability to find an optimal component sequence. Examples of other equivalently optimal sequences are given by the following arrangements: 34728615, 47368251 and 43762851. In combination with the crossover operator, it is the function of the GA mutation operator to ensure that these sequences are also uncovered in the course of further generation calculations. Whether or not the additional solutions are necessary depends on the level of satisfaction of the nearly optimal solution(s) already obtained.

5. RESULTS AND DISCUSSIONS

Three variations of the Case Study Problem were investigated. In the first example minimization of the total thermal failure rate was the sole criterion for determining optimal candidate solutions. In the second example the additional requirement was imposed that the total interconnectivity length among wired components should also be minimal. These two examples assume forced convection flow. In a third example we examine the influence of buoyancy on the findings of the first two.

For the Reynolds number explored, $Re = 630$, eddy shedding was absent and steady state was achieved in all the component arrangements calculated under the guidance of the GA. The CPU time required to compute the maximum steady state surface temperatures of the heated components was typically about 30 min on an IBM RS-6000/530, for the grid spacing and time step used. The results reported were obtained using the MSGA with a population size equal to 7, the number of generations equal to 7, and crossover and mutation rates set to 0.6 and 0.1, respectively. As in the validation run, the values of the λ_i required to calculate total thermal failure rates were obtained using the maximum surface temperatures predicted by the numerical procedure.

5.1. *First example: optimization according to thermal sensitivity*

In this example $N = 8$ heated components of the size and thermal characteristics listed in Table 2 are to be arranged in the channel so that the cost function $f = \Sigma\lambda_i$ is minimized. Table 4 for this example shows that the average failure rate of the initial population is 0.02094 fr Mh^{-1} while the average failure rate at the end of Generation 7 is 0.01953 fr Mh^{-1}, a decrease of 7%. The average failure rate could be further reduced if the GA parameters such as number of generations, population size, crossover and mutation rates were simultaneously optimized for the present problem (Grefenstette [11]). As expected, the calculations show that the heat transfer from N in-line heated components of different characteristics is quite sensitive to their relative positioning. The variation among failure rates for the various arrangements uncovered by the GA was as high as 25%. The best arrangement achieved by Generation 7 is given by chromosome 41387562, with a failure rate of 0.01856 fr Mh^{-1}.

The isotherms and maximum component temperatures corresponding to this arrangement are provided in Fig. 2. Although not shown, flow streamlines revealed recirculating regions behind all the large components. It is particularly noteworthy that the average heat transfer coefficients of geometrically identical components with the same heat flux differ as a result of the dependence of h_{ave} on the relative location of the components. The dependence is due to the developing nature of the flow. Clearly, the assumption invoked in many electronics cooling applications, that h_{ave} is constant for a component irrespective of its relative location, is generally incorrect and bears upon the issue of problem formulation accuracy discussed in the Introduction.

The characteristics of good candidate solutions are associated with geometry, heat flux and thermal sensitivity. Figure 3 provides a graphical illustration of the component arrangements corresponding to the seven chromosomes in Generation 7 of Table 4. In the figure the arrangements are ranked from top to bottom in decreasisng order of fitness. In each arrangement, components with the higher of the two heat flux values are drawn with a bold line while those with the lower heat flux value are drawn with a light line. Components with the higher of the two B_i values are shaded while those with the lower B_i value are left blank. Visual inspection of these results immediately reveals two underlying patterns. First, that components with the higher value of B_i tend to be clustered near the channel inlet. Second, that components with the higher value of heat flux tend to be clustered at the outlet. However, a close inspection of the arrangement with the highest value of fitness in this generation suggests that a clustering of components according to *both* thermal sensitivity and size, with the small components ahead of the large, appears to outweigh clustering according to heat flux. Interestingly, the worst arrangement in this generation also shows small components clustered ahead of large components but, because clustering according to thermal sensitivity is absent in this arrangement, it has a low value of fitness.

Such findings would be difficult to uncover,

Table 4. MSGA report corresponding to the first example: (a) input parameters for the MSGA; (b) initial population statistics; (c) detailed results for Generations 1 and 2; (d) detailed results for Generations 6 and 7; (e) statistics for Generations 2 and 7

MSGA input parameters	Value
Population size	7
Chromosome length	8
Maximum generations	7
Crossover probability	0.6
Mutation probability	0.1

(a)

Initial population statistics	Value
FR_{max}	0.02321
FR_{min}	0.01888
FR_{avg}	0.02094

(b)

Generation 1

String	Fitness	*FR*	*P*1	*P*2	*x*1	*x*2
52368174	0.00042	0.02279	7	7	1	2
25136478	0	0.02321	7	7	1	2
84657321	0.00365	0.01957	5	7	1	4
26348517	0.00271	0.02050	5	7	1	4
57864321	0.00202	0.02119	4	7	4	6
53428617	0.00278	0.02044	4	7	4	6
81375462	0.00433	0.01888	3	4	2	8

Generation 2

String	Fitness	*FR*
81675432	0.00115	0.02042
81375462	0.00269	0.01888
31574826	0.00222	0.01935
56813472	0	0.02157
26374815	0.00132	0.02025
41358762	0.00195	0.01962
15348672	0.00019	0.02137

(c)

Generation 6

String	Fitness	*FR*	*P*1	*P*2	*x*1	*x*2
41357862	0.00072	0.01948	1	3	5	7
85371462	0.00109	0.01911	1	3	5	7
81375462	0.00132	0.01888	1	2	0	0
87351462	0.00105	0.01916	1	2	0	0
54318762	0	0.02020	1	6	1	2
75384162	0.00089	0.01932	1	6	1	2
85371462	0.00109	0.01911	6	1	3	4

Generation 7

String	Fitness	*FR*
81375462	0.00215	0.01888
41357268	0	0.02103
41357862	0.00154	0.01948
85371462	0.00192	0.01911
75314862	0.00070	0.02033
41387562	0.00247	0.01856
78354162	0.00172	0.01930

(d)

Statistics	Generation 2	Generation 7
FR_{max}	0.02157	0.02103
FR_{min}	0.01888	0.01856
FR_{avg}	0.02021	0.01953

(e)

especially in more complex flows, without an adaptive solution methodology. If a 'rule' were to be extracted from the findings specific to this example it would be that: for conditions similar to those calculated, involving components with comparable values of their heat fluxes, arrangements with high values of fitness (low thermal failure rates) will have small components preceding large ones and components with large values of thermal sensitivity preceding components with small values of this quantity. The observed dependence of solution optimality on component size is not surprising since arrangements consisting of small heated components immersed in the wakes of larger ones are especially unfavorable for heat transfer from the smaller components.

Having detected and ranked the features responsible for good performance in this example, the arrangement in Fig. 3 with highest fitness can be further improved *without* using the MSGA. For example, moving component 1 from its present location in that chromosome to a position between components 7 and 5 yields the chromosome 43871562. In accordance with the rule just extracted, this chromosome has all four components with high B_i ahead of those with low B_i, and these first four components are arranged so that the small ones precede the large. A numerical calculation of the flow and temperature field for the new arrangement yielded a total failure rate $FR = 0.01810$ fr Mh^{-1} which is 2.5% smaller than the failure rate associated with the best chromosome in Generation 7.

The discussion of this example shows that: (a) sufficiently detailed calculations (or measurements) of the convective heat transfer characteristics of complex

h_{avg}	17.4	10.5	9.2	9.9	6.6	6.3	7.5	6.5
T_i	351	339	341	367	356	357	378	372
COMP.	4	1	3	8	7	5	6	2

h_{avg}	17.4	10.5	9.2	8.5	8.0	8.1	7.1	5.2
T_i	351	339	341	348	364	375	381	364
COMP.	4	3	1	7	2	8	6	5

FIG. 2. Isotherms and maximum component temperatures for the best arrangements obtained by the genetic algorithm for the first example (top; 41387562) and the second example (bottom; 43172865). Average heat transfer coefficients, maximum surface temperatures and component identification are shown.

geometrical configurations are necessary to provide the flow field information (velocity, temperature, and associated quantities) from which performance measures are to be determined and the satisfaction of constraints verified; (b) by preserving and exploiting the features associated with good performance, in just a few generations a GA is capable of evolving an initial population of chromosomes, here arrangements of components, into families of solutions with improved performance. It is precisely the ability of a GA to uncover and propagate, on average, the schemata of high performance chromosomes that makes this approach so cost-effective, especially during the early stages of a search.

5.2. *Second example: optimization according to thermal sensitivity and wiring length*

In so far as the channel configuration, flow conditions and the characteristics of the heated components are concerned, the second example is identical to the first example. As before, we wish to find thermofluids-optimized arrangements of the same eight heated components that minimize the sum of their failure rates subject to the *additional* requirement that the total interconnectivity length among the components, specified according to a pre-established interconnectivity matrix, *also* be minimized. Thus there are now two, possibly competing, criteria to be weighted in the optimization process. The present interconnectivity requirement is that all small components and all large components should be wired among themselves, respectively, but that no small component should be wired to any large component.

The function f to be minimized now has the following form:

$$f = (\Sigma\lambda_i)^{1/m}(C_0\Sigma\Sigma M_{ij}L_{ij})^{1/n}. \quad (3)$$

In equation (3) C_0 is a scaling factor, calculated for each generation in order to render the average contribution of the interconnectivity term in the product comparable in magnitude to the average contribution due to the total failure rate. The quantity M_{ij} denotes the (i, j) component of the interconnectivity matrix specified, a number equal to 0 or 1. L_{ij} denotes the wiring length between components 'i' and 'j'. In the double summation, the outside sum is over the values $i = 1, N$ and the inside sum is over the values $j = i, N$. The exponents $1/m$ and $1/n$ determine the relative weights of the two terms in equation (3) subject to the constraint that $1/m + 1/n = 1$. In the present example we have set $1/m = 1/n = 1/2$ for illustration purposes.

The use of equation (3) ensures that configurations which perform relatively well under *both* criteria will yield the highest fitness values calculated by equation (2). As before, the numerical procedure provides the maximum surface temperature for each component, necessary for calculating the individual failure rates. The total interconnectivity length for any arrangement is calculated within the MSGA according to the interconnectivity matrix specified. As indicated by the double summation, this involves the sum of the wiring lengths among all wired components, here connected center to center.

The chromosomes resulting from GA-guided calculations of populations evolving over seven generations (not shown here) revealed that both the failure rate and interconnectivity length are sensitive to the positioning of the heated components. These two quantities varied by up to 12 and 54%, respectively, among the various arrangements uncovered by the MSGA. For this second example, where two criteria are simultaneously imposed to establish good candidate solutions, the average failure rate and inter-

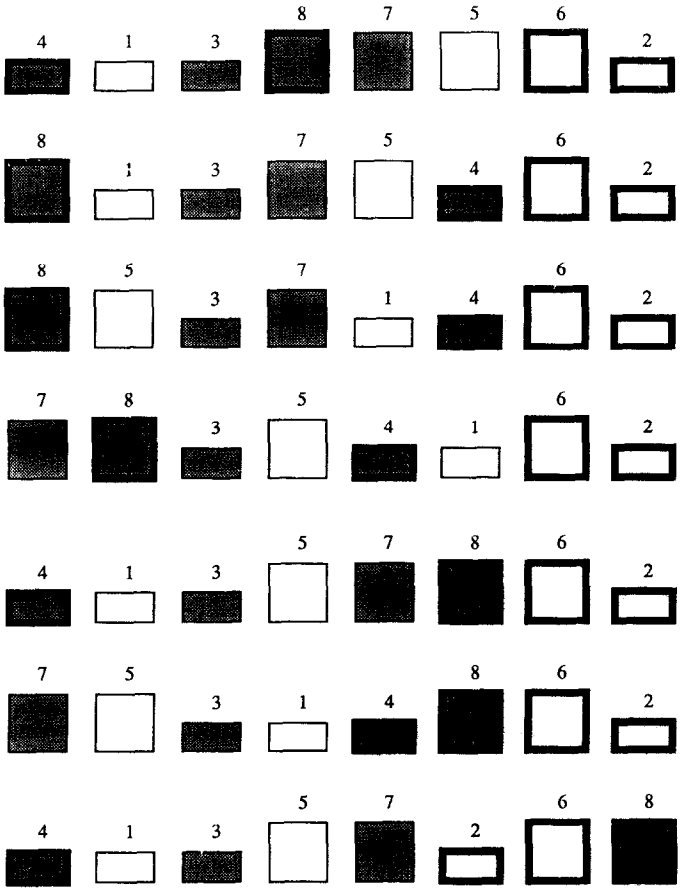

FIG. 3. Different component arrangements obtained at the end of Generation 7 by the application of the genetic algorithm in the first example. The arrangements are listed in order of fitness with the best at the top. In the figure: bold lines enclosing boxes correspond to $q = 300$ W m^{-2}, light lines enclosing boxes correspond to $q = 200$ W m^{-2}, tall boxes correspond to components of height 0.4 cm, short boxes correspond to components of height 0.2 cm, shaded boxes correspond to components with $B_i = 1600$ fr Mh^{-1}, blank boxes correspond to components with $B_i = 400$ fr Mh^{-1}.

connectivity length of the initial population were 0.02020 fr Mh^{-1} and 1.44 m, respectively. By Generation 7 the average failure rate and interconnectivity length (calculated using the different arrangements only) had dropped by 4 and 24%, respectively, compared to the initial population. The best chromosome in Generation 7 was found to be 43172865 and it appeared listed four times in that generation as a result of its dominance. This chromosome represents an arrangement with a total failure rate of 0.01920 fr Mh^{-1} and an interconnectivity length of 1.04 m. The corresponding isotherms and maximum component temperatures are shown in Fig. 2.

Figure 4 provides a graphical illustration of the component arrangements corresponding to the seven chromosomes obtained by Generation 7 of this example. (As in Fig. 3, the arrangements are ranked from top to bottom according to fitness and repeated arrangements are drawn only once. The encodings for size, thermal sensitivity and heat flux in this figure are the same as for Fig. 3). The interpretation of the results in Fig. 4 is facilitated by reference to Fig. 3, where the minimization of thermally-induced failure rates was the sole criterion for optimization. We start by noticing that the best arrangement in Fig. 3 has a smaller failure rate than the best arrangement obtained in Fig. 4, but that it also has an interconnectivity length which is 23% higher. It is clear from the best arrangement in Fig. 3 that a minimal wiring condition would immediately result if component 2 were placed between components 3 and 8. However, this would come at the cost of moving components 8 and 7 with high B_i into a hotter portion of the coolant stream where the risk of thermal failure is larger. (Recall the rule extracted from the results of the first example.) Thus, what would be gained in performance by minimizing the wiring length of the best arrangement in Fig. 3 would be partly lost to an increased thermal failure rate.

The compromise solution evolved in the present example by the time the MSGA reaches Generation 7 is reflected in the best arrangement shown in Fig. 4, where component 2 appears between components 7 and 8. In this arrangement the total wiring length has

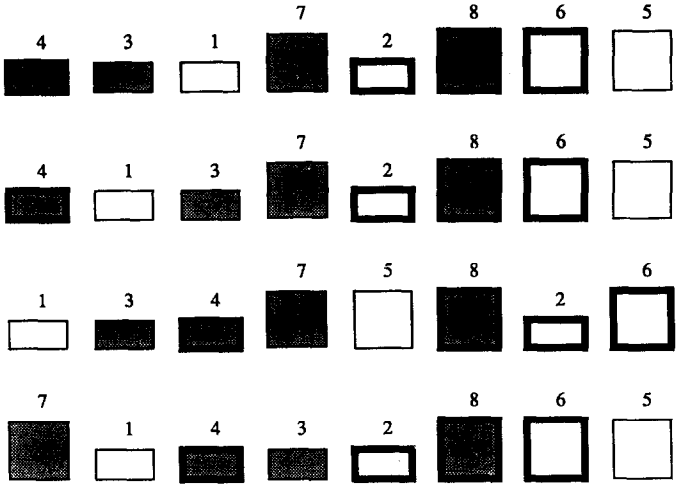

FIG. 4. Different component arrangements obtained at the end of Generation 7 by the application of the genetic algorithm in the second example. The arrangements are listed in order of fitness with the best at the top. The encoding of component size, heat flux and thermal sensitivity is as specified in Fig. 3.

been reduced significantly by respectively clustering small and large components closer together. But by keeping component 7 with high B_i ahead of 2, the (unavoidable) increase in the total thermal failure of this arrangement has been kept small. The best arrangement, obtained by the MSGA, was artificially altered to be 43127865, in order to further reduce the total interconnectivity length from 1.04 to 0.8 m, the minimum value. This changed the cost function given by equation (3) from $f = 0.0084$ for the original chromosome to 0.0075 for the altered one, a decrease of almost 11%. The corresponding effect on the thermal failure rate was to raise it slightly, from $FR = 0.01920$ to 0.01962 fr Mh^{-1}, an increase of 2.2%.

The second example shows clearly the competing effects between component geometrical layout (aiming to minimize interconnectivity) and thermofluids performance (aiming to minimize thermally-induced failure rates). At this level of configuration complexity, numerical calculations (or measurements) of the flow and temperature fields are indispensable for evaluating the maximum component temperatures and interconnectivity lengths required by the fitness function.

5.3. *Third example: the effects of buoyancy*

In the above two examples, the Reynolds number based on the inlet velocity, U, and the channel height, H, was $Re = 630$. The Grashof number based on the length of a component, $L1$, was $Gr = 51\,740$ and, therefore, the ratio Gr/Re^2 was of order 0.1. This led us to assume that the flow field calculations could be performed neglecting buoyant contributions to momentum. We now revise this assumption and establish the attendant impact on the results obtained.

The temperature fields shown in Fig. 2 correspond to the best forced convection arrangements obtained in the first (41387562) and second (43172865) examples, respectively. The fields for these two arrangements were recalculated employing the Boussinesq approximation to model the buoyant contribution to momentum. For this, the CUTEFLOWS numerical procedure was extended as explained in Iglesias *et al.* [16]. Otherwise, as before, constant physical properties for air at 300 K were used. With reference to Fig. 1, two channel orientations were explored for the best arrangement already determined for each example: (i) a horizontal orientation, with the channel aligned in the x coordinate direction and gravity aligned in the y coordinate direction; and, (ii) a vertical orientation, with the channel and gravity both aligned in the x coordinate direction.

The resulting buoyancy-affected temperature fields for each arrangement (not shown here) were very similar to those corresponding to Fig. 2 irrespective of the channel orientation. Notwithstanding, in the presence of buoyancy the maximum surface temperatures predicted for some of the components were sufficiently different to significantly alter the total thermal failure rates of these two arrangements and Table 5 summarizes the findings. The tabulated results show that for the conditions calculated the temperatures of the first four to five components in an arrangement are relatively insensitive to orientation as a consequence of the weakness of the buoyancy-induced flow. However, the last three to four components in an arrangement show a significant dependence on orientation. For these components, the *cumulative* effect of buoyancy has been to enhance the heat transfer from them, but to a larger extent for the vertical orientation than for the horizontal. This is because in the vertical orientation the imposed and buoyancy-driven contributions to momentum are aligned in the same direction; they combine to increase the heat transfer coefficients of the heated components relative to the horizontal orientation where the two contributions to

Table 5. Component maximum surface temperatures and total thermal failure rates for the best arrangements in the first and second examples with (b) and without (nb) buoyancy

First example									
Arrangement/orientation	4	1	3	8	7	5	6	2	*FR*
Horizontal (nb)	351	339	341	367	356	357	378	372	0.01856
Horizontal (b)	350	339	341	366	356	357	377	372	0.01852
Vertical (b)	350	339	341	365	353	353	371	370	0.01741
Second example									
Arrangement/orientation	4	3	1	7	2	8	6	5	*FR*
Horizontal (nb)	351	339	341	348	364	375	381	364	0.01920
Horizontal (b)	350	339	341	347	363	374	380	364	0.01888
Vertical (b)	350	339	341	346	363	370	374	360	0.01756

momentum are mutually perpendicular. In this regard, we note that component arrangements in a vertical channel have failure rates 6–7% lower than the same arrangements in a horizontal channel.

These observations reinforce the point raised in the Introduction, concerning the importance of implementing accurate physico–mathematical formulations and numerical solution methodologies in complex thermofluids optimization problems. Because of the coupling between temperature and velocity, free and mixed convection flow conditions add to the complexity of finding optimal or nearly optimal arrangements of convectively cooled electronic components, further emphasizing the need for adaptive solution methodologies. In flow configurations similar to the one examined here, for values of $Gr/Re^2 > 0.1$ the thermal failure rates of electronic components in channel geometries may depend significantly on buoyancy-induced motions and, hence, on configuration orientation with respect to gravity.

5.4. *Comments on the 'rules of thumb' for arranging electronic components*

The performance correctness of the MSGA was established through the validation run discussed in Section 4.4. It was shown there that the MSGA arranges identical components with identical heat fuxes in decreasing order of thermal sensitivity. That this is correct follows from the fact that all arrangements of components with the same geometry and heat flux will yield the same relative temperature distribution in a ventilated channel; that is, monotonically increasing for both the fluid and the components from the inlet to the outlet planes. Thus, the optimal arrangement of such components is that which associates the smallest values of maximum temperature with the most thermally sensitive components. This results in the ordering according to decreasing values of thermal sensitivities and is the basis for the *second* rule of thumb discussed in Section 4.1.

There is a difficulty with the *first* rule of thumb which states that, if the components are geometrically identical and their thermal sensitivities are the same, they should be arranged in an increasing order of their heat fluxes. As mentioned earlier, this conclusion is correct only for flow configurations allowing a lumped formulation analysis of the flow. To apply, the analysis requires that the flow field should be thoroughly mixed and thermally homogeneous in the space around each component. (This is the 'tank-and-tube' approximation referred to in the Introduction.) The requirement is not met for the conditions of the present study which is typical of the materials and flow conditions in electronic cooling applications. In particular, the presence of recirculating flow regions between components is the cause for the occurrence of non-uniform and relatively high temperatures, particularly on the downstream surfaces of the components in the present ventilated channel configuration.

That the first rule of thumb can lead to *non-optimal* arrangements of components with different heat fluxes was verified numerically. For this we calculated the velocity and temperature fields associated with $N = 8$ components of identical size (the smaller size in Table 2) and identical thermal sensitivity ($A_i = 4635$ K and $B_i = 1600$ fr Mh^{-1}, $i = 1, 8$). The component heat fluxes were $q_1 = 100$ W m^{-2}, $q_2 = 200$ W m^{-2}, $q_3 = 300$ W m^{-2}, ..., $q_8 = 800$ W m^{-2}. Temperature distributions like those shown in Fig. 2 were obtained for two arrangements: one in which the components were arranged in increasing order of their heat fluxes; another in which they were arranged in decreasing order. From the calculated component maximum surface temperatures, T_i, the total failure rates, $f = \Sigma\lambda_i$, were obtained for each of the two arrangements.

Table 6 summarizes the results. This shows that, contrary to the recommendation of the first rule of thumb, it is the ordering according to *decreasing* heat flux which yields the arrangement with smaller total thermal failure rate (0.153 fr Mh^{-1} vs 0.183 fr Mh^{-1}). Inspection of the individual component temperatures and their associated failure rates shows why. When placed in order of increasing heat flux, the first four components acquire smaller values of T_i than

Table 6. Maximum surface temperatures and thermal failure rates for eight components of identical size (small size in Table 2) and thermal sensitivity ($A_i = 4635$ and $B_i = 1600$, $i = 1, 8$) arranged according to increasing (top) and decreasing (bottom) order of their heat fluxes: $q_1 = 100$ W m^{-2}, $q_2 = 200$ W m^{-2}, $q_3 = 300$ W m^{-2}, . . ., $q_8 = 800$ W m^{-2}. The total failure rates are $\Sigma\lambda_i = 0.183$ for case (a) and $\Sigma\lambda_i = 0.153$ for case (b)

	q_1	q_2	q_3	q_4	q_5	q_6	q_7	q_8
T_i	317.0	335.5	355.5	376.4	398.2	420.9	444.4	468.9
λ_i	0.001	0.002	0.004	0.007	0.014	0.026	0.047	0.082

	q_8	q_7	q_6	q_5	q_4	q_3	q_2	q_1
T_i	434.6	432.1	424.0	412.3	398.2	382.4	365.2	347.1
λ_i	0.037	0.035	0.029	0.021	0.014	0.009	0.005	0.003

when placed in order of decreasing heat flux. However, the opposite is true for the last three components which acquire larger values of T_i. In particular, the last two components of the arrangement ordered according to increasing heat flux have temperatures *significantly larger* than the temperatures of the hottest components arranged according to decreasing heat flux. Since, according to equation (1), component failure rate is exponentially dependent on temperature, this results in a higher total thermal failure rate for the arrangement ordered according to increasing heat flux.

6. CONCLUSIONS

This study demonstrates how an adaptive methodology can be used to efficiently search the solution spaces characterizing complex thermofluids configurations in order to establish optimal performance conditions. The demonstration was based on recursively solving the conservation equations and boundary conditions describing an electronics cooling problem under the guidance of a GA. It could also have been based on performing the equivalent experiment. In the process of uncovering possible solutions, a GA dynamically gathers information, detects features responsible for good performance and exploits these features to generate improved solutions for the problems considered. Because of the high dimensionality and strong non-linearities typical of complex thermofluids problems, it would be impracticable to search through the multiplicity of their solutions without an adaptive technique.

The GA optimization strategy is premised on the notion that features responsible for good performance should emerge and be propagated within a calculable number of generations as a result of selection, crossover and mutation operators applied to the chromosomes encoding this information. In principle, the combined effects of the crossover and mutation operators is to guarantee that the entire solution space is explored irrespective of its topology; that is, irrespective of the number of maxima or minima characterizing the space. In this regard, the GA approach is superior to the more traditional optimization procedures.

Typically, a GA will provide not one but several possible solutions that satisfy problem constraints in different ways. The availability of a family of acceptable solutions provides the flexibility needed to simultaneously satisfy design criteria involving geometrical, electrical, mechanical and related cost constraints. It has been our experience that the first few generations of chromosomes obtained by a GA often suffice to provide qualitative guidance concerning the influence of variables and/or parameters on performance trends and problem solution optimality. This knowledge allows the investigator to shortcut the GA by constructing and testing new chromosomes which are likely to represent acceptable solutions. If, however, they turn out to be unsatisfactory, these improved investigator-generated strings can be seeded as new members of the population in further GA calculations. This approach has the potential for significantly accelerating the attainment of optimal or nearly optimal solutions.

The variations of the case study problem explored in this investigation served to illustrate the way a GA can be used in the context of an adaptive solution methodology. Among the main effects examined was the dependence of solution optimality on the cost function specified, especially in the presence of competing effects. Thus, in the present application to heated electronic components in a ventilated channel, a minimization of thermal failure rates was favored by clustering components according to size and thermal sensitivity. In contrast, a minimization of wiring length was favored by clustering them according to a predetermined interconnectivity matrix. The nearly optimal arrangements uncovered by the GA for $N = 8$ components of different size, heat flux and thermal sensitivity were far from obvious. They point to unconventional ways for arranging components so that flow-surface interactions favor heat transfer. In this regard it has been shown that the neglect of buoyant effects, or the implementation of 'rules of thumb' derived from physico–mathematical models lacking formulation accuracy (or derived from inappropriate experiments), can lead to non-optimal arrangements of components.

In the above examples the times required to obtain optimal solutions were relatively small because of the coarseness of the grids used and the steady two-dimensional nature of the flow fields. More complex unsteady three-dimensional problems using the same methodology will require dense grids, high performance machines, and/or distributed computing environments. An alternative to the physico–mathematical model approach is to conduct a GA optimization search with an experimental apparatus providing the quantities required by the cost or fitness

functions. Of course, hybrids of numerical and experimental approaches under GA control are also possible and this is the subject of continued ICEME research.

In concluding we note the considerable potential in *all areas* of science and engineering for adaptive solution methodologies based on GAs. In particular, the heat transfer community can expect to see a significant increase in pioneering applications of such methodologies to many complicated thermosciences problems admitting optimization in some sense. These exciting applications are being facilitated by the increased availability of high performance computers, distributed computing environments and improved guidelines for the specification of the necessary GA parameters.

Acknowledgements—One of us (J.A.C.H.) gratefully acknowledges a RISC System/6000 Workstation provided by the IBM Corporation through the efforts of Dr H. Greenberg of the Almaden Research Center (San Jose, California), as well as funding also provided by IBM through V. Kamath (Boca Raton, Florida) and R. Linton (Austin, Texas), to pursue this investigation. Very special thanks go to these three persons for their unflagging support of the ICEME concept, furthered here in the context of an adaptive solution methodology. Thanks are also due to the Computer Mechanics Laboratory of the Department of Mechanical Engineering at the University of California at Berkeley and to the Fundacion Gran Mariscal de Ayacucho, and the Universidad de Zulia, Venezuela, for their support of R.D. and N.Q., respectively. Some of the lengthier numerical calculations of this study were performed at the SDSC at the University of California, San Diego, with a grant received from that institution.

REFERENCES

1. J. A. C. Humphrey, R. Devarakonda and N. Queipo, Interactive computational-experimental methodologies (ICEME) for thermofluids research: Application to the optimized packaging of heated electronic components. In *Computers and Computing in Heat Transfer Science and Engineering* (Edited by K. T. Yang and W. Nakayama, pp. 293–317). CRC Press and Begell House, New York (1993).
2. J. Weiss, P. Fortner, B. Pearson, K. Watson and T. Monroe, Modeling air flow in electronic packages. *Mech. Engng* 56–58, Sept. (1989).
3. K. Azar, S. E. Develle and V. P. Manno, Sensitivity of circuit pack thermal performance to convective and geometric variation, *IEEE Trans. Components, Hybrids, Manufact. Technol.* **12**(4), 732–740 (1989).
4. S. L. S. Jacoby, J. S. Kowalik and J. T. Pizzo, *Iterative Methods for Non-linear Optimization Problems*. Prentice-Hall, Engelwood Cliffs, New Jersey (1972).
5. J. D. Hill, A search technique for multimodal surfaces, *IEEE Trans. Systems, Sci. Cybernetics* **SSC-3**(1), 2–8 (1969).
6. D. Dancer and M. Pecht, Component placement optimization for convectively cooled electronics, *IEEE Trans. Reliability* **38,** 199–205 (1989).
7. C. A. Schuler, E. B. Treidler and J. A. C. Humphrey, *CUTEFLOWS (Computing Unsteady Three-dimensional Elliptic Flows): A User's Guide*. Computer Mechanics Laboratory Report # CML-91-001, Mechanical Engineering Department, University of California at Berkeley (1991).
8. K. De John, Adaptive system design: A genetic approach. *IEEE Trans. Systems, Man Cybernetics* **SMC-10**(9), 566–574 (1980).
9. J. H. Holland, *Adaptation in Natural and Artifical Systems*. MIT Press, Cambridge, Massachusetts (1975).
10. D. E. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, New York (1989).
11. J. J. Grefenstette, Optimization of control parameters for genetic algorithms, *IEEE Trans. Systems, Man Cybernetics* **SMC-16**(1), 122–128 (1986).
12. J. E. Baker, Reducing bias and efficiency in the selection algorithm, *Proc. Second Int. Conference on Genetic Algorithms* pp. 14–21. Lawrence Erlbaum, Hillsdale, N.J. (1987).
13. I. M. Oliver, D. J. Smith and J. R. C. Holland, A study of permutation crossover operators on the traveling salesman problem. In *Genetic Algorithms and Their Applications: Proceedings of the Second International Conference on Genetic Algorithms*, pp. 224–230. MIT, Lawrence Erlbaum Associates (1987).
14. L. Davis (Editor), *Handbook of Genetic Algorithms*. Van Nostrand Reinhold, New York (1991).
15. D. E. Goldberg and R. Lingle, Alleles, loci, and the traveling salesman problem, *Proc. International Conference on Genetic Algorithms and their Applications* 154–159 (1985).
16. I. Iglesias, J. A. C. Humphrey and F. Giralt, Numerical calculation of two-dimensional, buoyancy-assisted flow past a backward-facing step in a vertical channel. In *Computational Aspects of Heat Transfer, Benchmark Problems* (Edited by B. F. Blackwell and B. F. Armaly), HTD-Vol. 258, pp. 63–72, ASME Winter Annual Meeting, 28 November–3 December, New Orleans, Louisiana (1993).
17. C. A. Schuler, Investigation of the flow between rotating disks in an enclosure, Ph.D. Thesis, University of California at Berkeley (1990).
18. E. B. Treidler, An experimental and numerical investigation of flow past ribs in a channel, Ph.D. Thesis, University of California at Berkeley (1991).
19. J. A. C. Humphrey, C. A. Schuler and I. Iglesias, Analysis of viscous dissipation in disk storage systems and similar flow configurations, *Phys. Fluids A* **4,** 1415–1427 (1992).
20. K. Tatsutani, W. R. Usry and J. A. C. Humphrey, Numerical calculation of two-dimensional laminar flow and heat transfer for a backward facing step using CUTEFLOWS. In *Benchmark Problems for Heat Transfer Codes* (Edited by B. F. Blackwell and D. W. Peffer), HTD-Vol. 222, pp. 1–6, ASME Winter Annual Meeting, 8–13 November, Anaheim, California (1992).
21. K. Tatsutani, R. Devarakonda and J. A. C. Humphrey, Unsteady flow and heat transfer for cylinder pairs in a channel, *Int. J. Heat Mass Transfer* **36,** 3311–3328 (1993).
22. M. P. Arnal, D. J. Goering and J. A. C. Humphrey, The influence of a solid boundary on vortex shedding from a bluff body, *J. Fluids Engng* **113,** 384–398 (1991).