

A. Despierre, P.J. Stuttaford, P.A. Rubini. "Gas turbine combustor design using genetic algorithms". Int. Gas Turbine and Aeroengine Congress & Exhibition, Orlando, Florida, June 2-5, 1997. ASME Paper 97-GT-72

PRELIMINARY GAS TURBINE COMBUSTOR DESIGN USING A GENETIC ALGORITHM

Arnaud Despierre, Peter J. Stuttaford, Philip A. Rubini #

School of Mechanical Engineering, Cranfield University, Cranfield, Bedfordshire, MK43 0AL, UK
Tel: + 44 1234 750111, FAX: + 44 1234 750728, Email: p.a.rubini@cranfield.ac.uk

corresponding author

ABSTRACT

A genetic algorithm, coupled with a versatile preliminary design tool, is employed to demonstrate the concept of an autonomous design procedure for gas turbine combustors with user specified performance criteria.

The chosen preliminary design program utilises a network based approach which provides considerable geometric flexibility allowing for a wide variety of combustor types to be represented. The physical combustor is represented by a number of independent, though interconnected, semi-empirical sub-flows or elements. A full conjugate heat transfer model allows for convection, conduction and radiative heat transfer to be modelled and a constrained equilibrium calculation simulates the combustion process. The genetic algorithm, whose main advantage lies in its robustness, uses the network solver in order to progress towards the optimum design parameters defined by the user. The capabilities of the genetic program are demonstrated for some simple design requirements, for example zone fuel/air ratio, pressure drop and wall temperatures.

NOMENCLATURE

C_p	specific heat at constant pressure
d	external mass flow into node
δ	schema defining length
ΔE	energy source term
f	fitness function
F	population fitness
H	symbolic schema
h	convective heat transfer coefficient
J	total number of network nodes
m	schema instance
n	generation number
\dot{m}	mass flow rate
p_m	mutation probability

p_c	crossover probability
P	pressure
Q	volumetric flow rate
ΔQ_{rad}	radiative energy source
R	conductive heat transfer coefficient
T	temperature
x_k	GA real value for parameter k
x_{Tk}	GA real target value for parameter k
ρ	density

Subscripts

i	= centre node
i,j	= branch element j to node i
$n_{i,j}$	= branch node

INTRODUCTION

Gas turbine combustor preliminary design programs are typically analysis programs driven by the user, with little automatic capability. Design parameters are often highly correlated and performance targets strongly conflicting. This leads to complex multi-parameter optimisation problems, solutions for which are difficult to obtain using common optimisation techniques.

Preliminary design programs typically require considerable empirically determined information. However, these models have the advantage of being computationally efficient, a necessary criteria for any form of optimisation procedure. A common disadvantage is the limitation to simple geometries. Difficult to set-up, and suffering from critical convergence problems when applied to complex geometries their use can appear restrictive. Network approaches provide the ability to effectively model a wide variety of geometries with fewer numerical difficulties, keeping computational time well below other common methods, and substantially less than 3-D computational fluid dynamics (CFD) codes.

Gas turbine combustor design techniques are initially based upon overall performance requirements such as overall pressure loss, combustion efficiency, exit temperature traverse, and emissions. The application of a characteristic time model in achieving these goals was demonstrated in Mellor and Fritsky (1990), where the model provided full details of primary and dilution air requirements. However, before undertaking rig testing a more complete three-dimensional analysis of the flametube should be achieved.

Empirical design methods are an essential component of the combustor design process. Unfortunately, when new technological features are involved they show their limitations. Techniques such as CFD allow for the complete modelling of gas turbine combustors flows but they suffer from large computational time penalties when very accurate results are required (Mongia, 1994). These three-dimensional tools are often only used to visualise combustors for which a preliminary design has already been achieved.

Sizing and designing a combustor with one-dimensional models has been successfully demonstrated by Burrus et al (1987). They demonstrated computing flow areas for target air flow distributions, or alternatively, given the areas the various flow splits and pressure-drops were calculated. The model includes several empirically derived correlations and experimental data which results in an improvement to the simplified overall governing equations (Mellor, 1990).

Network analysis based on one-dimensional models, although lacking the resolution of three-dimensional models, can provide sufficiently accurate results in a short computational time. The method allows for simple modelling of complex geometries whilst allowing for the continued development of the sub-models embodied within these procedures, e.g. equilibrium models, film-cooling models, radiation models. Stuttaford and Rubini (1996,1997) demonstrated and further extended the successful application of a network solver to the design process of a gas turbine combustor.

In the context of the present application the combustor analysis program is to be used within an overall constrained optimisation process, whereby a number of predefined parameters are varied in order to satisfy a set of design targets. For example, pressure drop, zone fuel/air ratios and wall temperatures may be specified as design targets against which dilution ports and cooling flow areas are to be varied.

There has been considerable development of numerical algorithms for constrained and unconstrained optimisation, see Dennis and Schnabel (1983), Fletcher (1987), Rao (1996). Such classical methods are particularly successful for optimisation problems where the response function is smooth. However, they are less robust for problems where the response functions are highly non-linear. Furthermore, many of these approaches are gradient-based so they are less than satisfactory for typical engineering design problems where the shape of the response function is often unknown.

There has been a growing interest in recent years concerning the use of artificial intelligence techniques in order to develop new optimisation concepts. Among these are genetic algorithms, (GA), parameter search algorithms based upon the natural concept of genetic selection. In the field of optimisation, GAs are known as a "weak" method, which means they are not particularly dedicated to one type of problem but the method can be applied to a wide variety of problems. This feature makes them particularly attractive for applications such as engineering design where little is known about the response function, or its derivatives. An introduction and overview of the capabilities of GAs can be found in KrishnaKumar (1992) who

also discussed the ability of GAs to tackle problems relevant to aerospace system optimisation. Gage and Kroo (1993) typify the successful application of this technique in a design role, for the preliminary design of a non-planar wing.

Many of the classical optimisation methods move from point to point in the decision space according to some deterministic transition rules. Such a point-to-point method is prone to trap the algorithm in a false peak when the search landscape is highly multi-modal, i.e. when several local optimum are present. Genetic algorithms overcome this problem by two means. Firstly, the principle of the algorithm is to evaluate many solutions at the same time. This parallelism in the evaluation is a fair guarantee, but still risks the hazard of mistaking a local optimum for a global one. The use of the mutation operator, described in more details later in this paper, is the ultimate guarantee for the GA not to be trapped. The operator is basically a periodical introduction of a random variation in the solution group. It ensures a regular random sampling on the response landscape, introducing diversity to check whether or not the best solution found so far is the global one.

Genetic algorithms do have some disadvantages. Firstly, to perform their calculations they need a relatively high number of response evaluations, which can lead to a considerable computational expense. Secondly, even though they can locate quite quickly the region of the global optimum, they may have difficulty in reaching the exactly optimal point. If rapidly finding a sufficiently good solution is acceptable then a GA has the ability to surpass other methods on a broad range of large-scale optimisation problems.

In the context of combustor design, the use of a rapid analysis program is essential, numerically intensive methods such as CFD would be impractical. The network based solver (Stuttaford and Rubini, 1996) satisfies the requirements of versatility, accuracy and computational economy.

COMBUSTOR NETWORK ALGORITHM

Each individual physical feature, for example, dilution ports or cooling rings, in the geometry can be represented by elements, linked together by nodes to form a network.

The network model divides the combustor into a series of elements described by one-dimensional sub-flows, containing independent semi-empirical governing equations, appropriately modelled according to the feature. The sub-flows are linked together by the overall governing equations, to obtain a complete solution of the entire flow field. In this manner mass flow splits and pressure drops are obtained throughout the combustor. A pressure correction methodology is utilised to solve the coupled continuity equations and element pressure drop/flow rate relationships (Greyvenstein and Laurie, 1994). For more details see Stuttaford and Rubini (1996,1997).

An example of such a network is shown in Fig. 1. for a double annular combustor, simulating the effects of both the combustion zones. The example is a relatively simple one as it does not include the components required for a detailed wall heat transfer analysis. In this figure circles represent nodes and squares elements.

The continuity equation may be specified as:

$$\sum_{j=1}^J \rho_{i,j} Q_{i,j} S_{i,j} = -d_i \quad i = 1, 2, \dots, J \quad (1)$$




Figure 1 Double annular combustor network

and the pressure-drop/flow-rate relationship as:

$$\Delta P_{i,j} = P_{n_{i,j}} - P_i = s_{i,j} z_{i,j} u_{i,j} w_{i,j} \quad (2)$$

where,

$$u_{i,j} = u_{i,j}(\rho_{i,j})$$

$$w_{i,j} = w_{i,j}(|Q_{i,j}|)$$

$$z_{i,j} = \frac{Q_{i,j}}{|Q_{i,j}|}$$

$$s_{i,j} = 1 \quad \text{if the flow is in the positive direction}$$

$$s_{i,j} = -1 \quad \text{if the flow is in the negative direction}$$

These relationships represented by semi-empirical expressions are specific for each individual combustor feature or internal flow region. Further information regarding these formulations, such as the Darcy-Weisbach equation and flow in duct sections, can be found in Jeppson (1977) and Saad (1985).

An enthalpy balance is performed at each node in the network for the energy equation to be satisfied. The form of the energy balance varies depending on the location in the domain, i.e. within the fluid, within the solid, or on a solid/fluid boundary. This may be specified at nodes with branch elements containing mass transfer as:

$$T_i = \frac{\sum_{j=1}^J (\Delta E_{i,j} + \dot{m}_{i,j} C_p T_{n_{i,j}})_{\text{inflow bits}}}{\sum_{j=1}^J (\dot{m}_{i,j} C_p)_{\text{inflow bits}}} \quad (3)$$

To compute node temperatures on boundaries or within walls, i.e. at nodes where the branch elements contain no mass flow, a semi-implicit formulation is used. Where the three modes of heat transfer are present, i.e. on the flow boundaries, the node temperature can be calculated by:

$$T_i = \frac{\left(\sum_{j=1}^J k_{i,j} T_{n_{i,j}} \right)_{\text{conduction}} + \left(\sum_{j=1}^J h_{i,j} T_{n_{i,j}} \right)_{\text{convection}} + \Delta Q_{\text{rad}}}{\left(\sum_{j=1}^J k_{i,j} \right)_{\text{conduction}} + \left(\sum_{j=1}^J h_{i,j} \right)_{\text{convection}}} \quad (4)$$

and within the solid where only conduction is present:

$$T_i = \frac{\left(\sum_{j=1}^J k_{i,j} T_{n,i,j} \right)_{\text{conduction}}}{\left(\sum_{j=1}^J k_{i,j} \right)_{\text{conduction}}} \quad (5)$$

The convective and conductive heat transfer coefficient terms in Eq. 4. are evaluated using semi-empirical correlations and experimental data for various cooling features found in gas turbine combustors. The axial location of an element within a feature is accounted for when determining heat transfer coefficients, thus allowing for film decay to be modelled. Values within an element are assumed to be constant, therefore resolution of such effects is controlled by the number of elements selected by the user. A wide range of cooling features can be modelled, including Z-rings, lipped-rings, slots, effusion patches, transply patches and film cooling. Examples of such correlations can be found in Lefebvre (1983). Further description of the heat transfer model may be found in Stuttaford and Rubini (1997).

GENETIC OPTIMISATION ALGORITHM

Since the late sixties biological evolution has become an attractive source of inspiration to address problems where computational parallelism is required. Holland et al (1975) were among the first to focus their attention on the possibility of developing a search procedure based on the concepts of natural genetics, the so-called genetic algorithm. In this Darwinian survival-of-the-fittest approach a population of chromosomes (combustors in our case) is moved from generation to generation according to the rules of natural selection, including the use of genetic-based operators. Each combustor is represented by a string of zeros and ones simulating a DNA structure.

The three basic operators, reproduction, crossover and mutation form the core of the Simple Genetic Algorithm (SGA), the simplest algorithm based upon genetic transitional rules (Goldberg, 1989).

The reproduction operator (also called selection operator) allows combustors from a population to be selected to survive through the next generation depending on their performance relative to the entire population. For this purpose a fitness function is defined which illustrates the performance of the chromosome (combustor) compared with the targets fixed for the optimisation. Then, the higher the fitness of a chromosome, the more instances of this chromosome being present in the next generation.

During reproduction (see Fig. 2), strings are randomly paired with further probability that a crossover occurs at a random location along the strings. This results in two new strings obtained by swapping all the bits between the crossover point and the end of the strings.

Through the crossover operator children may be produced who are completely different from their parents therefore introducing diversity into the genetic process. Crossover and reproduction are the base of every GA.

In parallel with crossover, a mutation is allowed which is simply a random alteration of a string position. In a binary alphabet this means changing a 1 for a 0 or vice versa. The mutation operator is believed to be very important in order to avoid being trapped in a local

minimum as it always introduces random diversity. Used sparingly it can significantly enhance the performance of a GA.




Figure .2. The crossover principle

Holland's concept of a population based algorithm was a major and original breakthrough among the set of optimisation techniques currently available and such techniques have now been successfully implemented for several types of application in fundamental and applied sciences. Furthermore, he was the first to set up a theoretical background for the method by developing the notion of "schemas" (or schematas) which were the basis for all subsequent work on genetic algorithms.

Genetic algorithm theory (Holland, 1975) is based upon the idea that the algorithm is discovering and reproducing good "building blocks" (sequences of bits) in the chromosomes, i.e. particular series of bits which contribute to bring the chromosome in which they are present closer to the optimal point. This leads to the definition of schemas to explain what is understood by building blocks. A schema is a series of zeros, ones and asterisks, the asterisks being wild cards or "don't care" symbols (they can stand for one or zero). Strings can then be created over the ternary alphabet $\{0,1,*\}$. For instance the schema $H = 0***1$ represents the set of all eight bit strings which begin with a 0 and ends with a 1. A schema is usually noted H as it is also known as one hyperplane belonging to the hypercube of the solution space. The schema has two defined bits and is said to be of order $o(H) = 2$, and its defining length (the distance between its extreme defined bits) is $\delta = 3$.

It should be underlined that the asterisk is only a metasympol (a symbol standing for other symbols) used for theoretical analysis and is not processed by the algorithm. By the same token the notion of schema is unknown to the algorithm and schema average fitness is never explicitly calculated by the GA. However, the increase or decrease of instances of specific schemas from generation to generation can be evaluated by using the Schema Theorem (1), (Holland, 1975) broadly acknowledged as the main basis for GA theory.

This theorem gives the expected number of instances m of a schema H at generation $n+1$, $m(H,n+1)$, according to the number of instances of H at generation n , the mutation probability p_m , the crossover probability p_c , and the fitness of schema H , $f(H)$, relative to the fitness of the population $F(n)$:

$$m(H,n+1) \geq m(H,n) \cdot \frac{f(H)}{F(n)} \left(1 - p_c \frac{\delta(H)}{l-1} \right) (1 - p_m)^{o(H)} \quad (6)$$

In processing populations of chromosomes (combustors), the genetic algorithm implicitly evaluates each schema present in the

population and enhances their increasing or decreasing presence using the schema theorem. These simultaneous evaluations are known as implicit parallelism, and the effect is to progressively bias the reproduction procedure for good schemas, or building blocks, to be more and more represented through all the generations.

For a computational genetic algorithm to be coupled with a design program, certain information should be provided by the user. Firstly, for every design parameter a targeted value must be fixed, or a targeted range should be allowed for. In order to reach this optimal user-defined point or area, varying parameters have to be chosen. For every varying parameter a minimum and maximum allowable value should be defined as boundaries to keep the design realistic. All these real-valued variables x_k are then non-dimensionalised and coded as binary strings, for instance using 16-bit integers, the real value x_k is represented by the integer:

$$\text{integer value } x_k \equiv \frac{(x_k - x_{min})}{(x_{max} - x_{min})} \times (2^{16} - 1) \quad (7)$$

Thus in this case the resolution is roughly 0.0015 percent over the range of x for each variable, as a result a slight loss of information occurs due to the binary coding technique used. Real number coding might be more suitable but it is a relatively new technique and less theoretical background has been devoted to it, thus a classical binary coding has been adopted.

In order to evaluate each combustor relative to the targets defined by the user, the fitness function, f , is defined so that the closer the combustor to the objective design, the higher the fitness for the reproduction operator to be performed. This fitness function is:

$$f = \frac{1}{\sum_k \frac{x_k - x_{T_k}}{x_{T_k}}} \quad (8)$$

SAMPLE GENETIC OPTIMISATION

To illustrate the capabilities of a genetic algorithm an example is presented requiring the minimisation of a simple function. The function is specified as:

$$t(x, y) = 1 - 2x - 2y - 4xy + 10x^2 + 2y^2 \quad (9)$$

and the minimum $t(x, y) = 0$ is given for $(x = 0.25; y = 0.75)$.

The run is made for 100 chromosomes per generation during 50 generations, with a crossover probability of 0.1 and a mutation probability of 0.01. Variables x and y are allowed to vary from -50 to 50. The best element was found at generation number 28, with: $x=0.2495$ and $y=0.74998$ leading to a function result of $t(x, y) = 0.000003$.

Notice that the precision reached for x and y is on the order of the maximum precision given by the encoding method, so no real benefit could be expected from a longer calculation.

Figure 3 presents a plot of fitness against generation, typical of a genetic optimisation. It must be remembered that the higher the fitness the closer the element is from the optimal solution. The curve featuring average fitness for each generation proves there is an overall

optimisation process occurring and the plot of the best element of each generation shows that individual elements use the good behaviour of the population to get closer to the optimum.




Figure .3. Simple function optimisation fitness evolution

Figure 3 highlights the optimisation capacity of genetic algorithms, but because a user can never be sure of reaching the absolute optimum solution a termination criterion must be chosen. Usually, it is a maximum number of generations (as used in this function optimisation case). Another termination strategy could involve a population convergence criteria such as the variance in the population fitness, or a lack of improvement in the best solution over a specified number of generations. Strategies can also be used in conjunction with each other.

Chromosomes treated by a GA can be function variables but in our design process they are defined as combustors.

COMBUSTOR DESIGN STUDY

An annular combustor of relatively simple geometry has been chosen to carry out the optimisation. The network set up for this combustor is shown in Fig. 4. This illustration is not to scale, the network has been expanded for clarity. The solid network of the elements and nodes refers to the flow computational cells, in which the flow and energy equations are solved. The dashed components of the network represent the overall heat transfer from the flametube through the liner to the annuli, accounting for conduction, convection (including various film cooling) and radiation effects. Heat transfer through the combustor liner is accounted for by groups of three elements representing the combined influence of convection, conduction and radiative fluxes.

The combustor is divided into three zones separated by primary and dilution cooling holes. Effusion cooling occurs all along the liner. The mass flow splits and pressure drop results for this design were compared and validated against results obtained from a proven industrial one-dimensional combustor algorithm (Stuttaford and Rubini, 1996). The study showed close agreement for mass flow splits and the difference in overall pressure drop was within 0.05%.

Figure .4. Annular combustor network used in genetic optimisation, with full flow and heat transfer modelling

The goal of the optimisation study was to simulate an initial preliminary design phase by keeping the actual features of the combustor fixed, except for six varying parameters, to reach eight design targets. The design targets were primary, intermediate and dilution zone average fuel/air ratios and average wall temperatures, overall pressure drop, and injector fuel/air ratio. This last target represents a means of controlling the relight capability of the combustor. The six varying parameters are primary and dilution holes area, injector area, and primary, intermediate and dilution effusion cooling holes area.

Table 1 summarises the eight targeted values chosen for the optimal design. It must be underlined that fuel/air ratios and wall temperatures are zone averaged values, therefore these are much more difficult to exactly optimise. The targeted design is similar to the original, but the original pressure drop was 5.2%, for illustration a value of 4% has been targeted. Lower intermediate zone fuel/air ratio and dilution zone wall temperatures have also been targeted to demonstrate the flexibility of the optimisation tool.

A study regarding the best values for genetic parameters in the combustor design process was carried out in Despierre (1996), the data chosen for this run are in accordance with the results found in the study. The genetic optimisation was carried out for 100 generations with 30 combustors per generation. The crossover probability was 0.5 and the mutation probability was taken to be 0.005. The number of fitness evaluations (and of solver runs) was limited to 3000 in order to keep calculation time below 24 hours on a low end workstation.

	value	value
Injector fuel/air ratio	0.0883	0.0880
Primary zone fuel/air ratio	0.0848	0.0850
Intermediate zone fuel/air ratio	0.0657	0.0600
Dilution zone fuel/air ratio	0.0459	0.0450
Primary zone wall temperature	1097 K	1100 K
Intermediate zone wall temperature	1150 K	1100 K
Dilution zone wall temperature	1162 K	1050 K
Overall pressure drop	5.2 %	4.00 %

Table .1. Design and original targets summary

RESULTS AND DISCUSSION

The genetic algorithm results obtained for the fitness evolution are presented in Fig. 5. As for the previous simple function optimisation example, the average fitness of every combustor globally increases throughout the generations, demonstrating that the optimisation process occurs. The best combustor of each generation also follows the upward trend.

Design parameters	Original	Targeted
-------------------	----------	----------

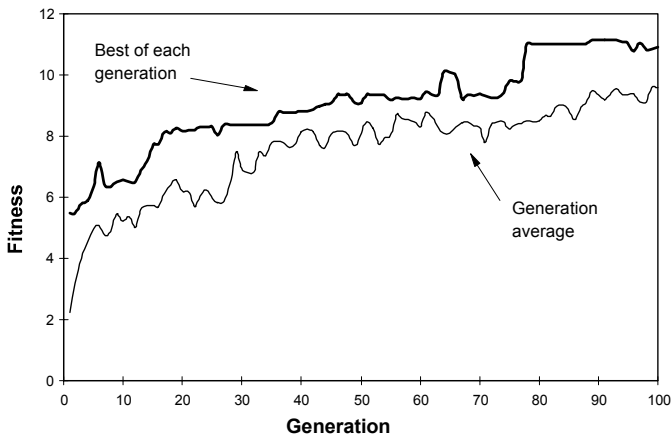


Figure .5. Combustor optimisation fitness evolution

The best combustor was found at generation 93 and it could be argued that it is surely not the best element in the response landscape. This highlights the fact that only an infinite number of generations can certify that the optimal point is reached, at which point the theoretical fitness equals infinity. The designer must decide when a satisfactory design performance has been reached, with no further significant improvement from the algorithm. In this case the calculation was halted after a reasonable period of no improvement. The fitness value on completion is arbitrary and represents the discrepancies between the target and calculated values, see below. Tables 2 and 3 present the results for the best combustor found.

Design parameters	Genetic design	Targeted value
Injector fuel/air ratio	0.0872	0.0880
Primary zone fuel/air ratio	0.0858	0.0850
Intermediate zone fuel/air ratio	0.0630	0.0600
Dilution zone fuel/air ratio	0.0452	0.0450
Primary zone wall temperature	1106 K	1100 K
Intermediate zone wall temperature	1096 K	1100 K
Dilution zone wall temperature	1062 K	1050 K
Overall pressure drop	4.02 %	4.00 %

Table .2. Genetic performance results

The overall design satisfactorily reaches the desired targets. The only parameter which was slightly left apart in the optimisation is the intermediate zone average fuel/air ratio (5 % of error against less than 1.5 % for every other parameter). This was one of the parameters for which the target was set far away from the original value. One way to handle this weakness could be to define a constrained fitness function biasing the optimisation process towards targets difficult to reach.

It is now interesting to have a look at the varying geometrical parameters set for the best combustor found (Table 3).

Combustor features	Original / genetic ratio
Primary port flow area	3.54
Dilution port flow area	0.98
P. zone effusion cooling flow area	0.40
I. zone effusion cooling flow area	2.70
D. zone effusion cooling flow area	2.21
Injector flow area	1.25

Table .3. Genetic design results

The value for the primary port flow area is much smaller than the original case due to the value targeted for the intermediate zone fuel/air ratio being greater than the original, leading to a much reduced primary port flow area. Whilst the injector fuel/air ratio is a constrained target value the injector flow area has decreased to control the overall pressure drop. Effusion cooling flow areas have increased due to the lower target values for wall temperature.

Concerning the performance of the genetically-designed combustor it is interesting to have a look in more detail at the real behaviour of the flow. It should be kept in mind that *average* fuel/air ratios and wall temperatures have been optimised, possibly leading to erratic distributions, fitting only the targets on average. Figure 6 presents as an illustration a comparison of the fuel/air ratio distribution in the flametube for the original and the new combustor.

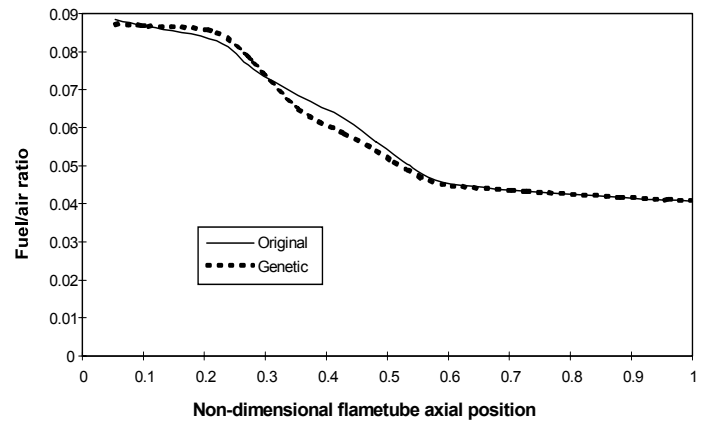


Figure .6. Fuel/air ratio distribution

The new distribution is very similar to the original, though with a lower fuel/air ratio for the intermediate zone as specified as a design target. This demonstrates the capability of a genetic algorithm, coupled with a network solver to quickly achieve a combustor preliminary design task, and illustrates the genetic-based search procedure.

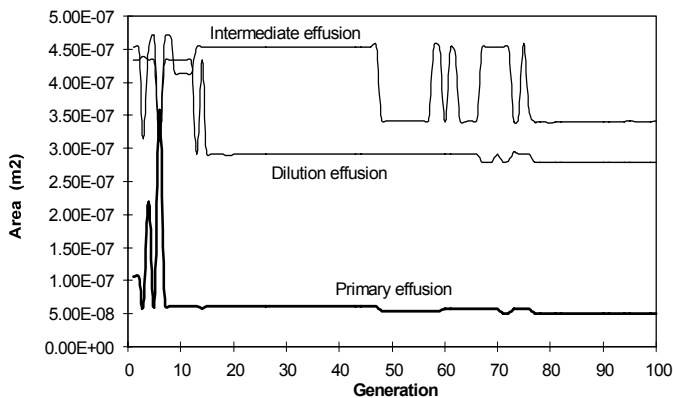


Figure 7. Parameter evolution

Figure 7 presents the evolution of the three effusion cooling flow areas for the best combustor along the generations. Optimum flow areas are found after a few generations and then the value for the intermediate zone oscillates between two regions on the response landscape. The same sort of evolution occurred for the three other varying parameters.

Such evolution is characteristic of the genetic algorithm, multi-parameter optimisation technique, in which implicit parallelism in the response evaluations makes it very robust and difficult to become trapped in a local optimum.

CONCLUSIONS

Preliminary design for gas turbine combustors typically employs computational analysis programs driven by the user, based upon past experience. This paper has presented a new and unique approach to gas turbine combustor design by the application of a genetic algorithm to control an existing analysis code.

The design procedure as reported is not without its drawbacks. It is presently not possible to include a number of critical details for example combustor relight, flame stability and pollutant emissions. In addition the present work has not addressed the issue of optimal choice of targets and design parameters.

The work represents a significant first step towards the autonomous design of combustors with user specified characteristics, as an alternative to the traditional approach of further development and modification of existing designs. By the very nature of genetic algorithms, such autonomous designs are unhindered by past design principles and offer the potential of further improvements in combustor performance.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the support of both the Defence Research Agency and Rolls Royce plc.

REFERENCES

- Burrus, D.L., Shyy, W. and Braaten, M.E., 1987, "Numerical model for analytical predictions of combustor aerothermal performance characteristics", APGARD CP 422, U.S.A.
- Dennis, J.E., Schnabel R.B., "Numerical methods for unconstrained optimization and nonlinear equations", 1983, Prentice-Hall.
- Despierre, A., 1996, "Gas turbine combustor design using genetic algorithms", MSc thesis Cranfield University.
- Fletcher, R., "Practical methods of optimisation", 1987, John Wiley
- Gage, P., Kroo, I. 1993, "A Role for Genetic Algorithms in a Preliminary Design Environment", AIAA Paper 93-3933.
- Goldberg, D., 1989, "Genetic Algorithms in Search, Optimisation and Machine Learning", Addison Wesley.
- Greyvenstein, G.P. and Laurie, D.P., 1994, "A Segregated CFD Approach to Pipe Network Analysis", International Journal for Numerical Methods in Engineering, Vol. 37, pp. 3685-3705.
- Holland, J.H., 1975, "Adaptation in natural and artificial systems", University of Michigan Press.
- Jeppson, R.W., 1977, "Analysis of Flow in Pipe Networks", Ann Arbor Science, Michigan, U.S.A.
- KrishnaKumar, K., 1993, "Genetic Algorithms: an introduction and an overview of their capabilities, AIAA Paper 92-4462-CP.
- Lefebvre, A.H., 1983, "Gas Turbine Combustion", Hemisphere Publishing Corporation, U.S.A.
- Mellor, A.M. (ed.), 1990, "Design of Modern Turbine Combustors", Academic Press, U.S.A.
- Mellor, A.M. and Fritsky, K.J., 1990, "Turbine Combustor Preliminary Design Approach", Journal of Propulsion and Power, Vol. 6, no. 3, pp. 334-343.
- Mongia, H., 1994, "Combustion modelling in design process applications and future direction", AIAA Paper 94-0466.
- Rao, S.S., 1996, "Engineering Optimization: Theory and Practice", John Wiley, U.S.A.
- Saad, M.A., 1985, "Compressible Fluid Flow", Prentice-Hall, U.S.A.
- Stuttaford, P.J., Rubini, P. A., 1996, "Preliminary gas turbine combustor design using a network approach", ASME Paper 96-GT-135, also to appear in *ASME Journal of Engineering for Gas Turbines and Power*.
- Stuttaford, P.J., Rubini, P. A., 1997, "Assessment of a radiative heat transfer model for a new gas turbine combustor preliminary design tool", AIAA Paper 97-0294.

© British Crown copyright 1997/DERA
Published by the American Society of Mechanical Engineers with permission of the Controller of Her Majesty's Stationery Office