UC Irvine

UC Irvine Previously Published Works

Title

Robust Optimal Control of a Natural Gas-Fired Burner for the Control of Oxides of Nitrogen(NO_x)

Permalink

<https://escholarship.org/uc/item/042458z6>

Journal

Combustion Science and Technology, 128(1-6)

ISSN

0010-2202

Authors

JOHN, DAVID ST
SAMUELSEN, SCOTT

Publication Date

1997-10-01

DOI

10.1080/00102209708935702

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Robust Optimal Control of a Natural Gas-Fired Burner for the Control of Oxides of Nitrogen(NO_x)

DAVID ST. JOHN and SCOTT SAMUELSEN

UCI Combustion Laboratory, University of California, Irvine, CA 92717

(Received 13 February 1997; In final form 9 May 1997)

Tightening requirements on industrial boilers and furnaces will require hands-free techniques to (1) assure peak performance with respect to emission, and (2) assure an ability to achieve peak performance throughout a load duty cycle. In the present paper, robust optimal control of a model industrial, swirl-stabilized, natural gas-fired burner is explored as a strategy to attain and maintain low flue-gas nitrogen oxide concentration ($[\text{NO}_x]$) concomitant with high combustion efficiency (η_c). A performance index, J , is defined such that the maximization of J correlates to optimal burner performance, with respect to $[\text{NO}_x]$ and η_c . Two parameters, swirl intensity (S') and excess air (EA), are made amenable to control and incorporated as variable burner inputs. For a given load, the settings of EA and S' are automatically adjusted by a specialized search algorithm in order to maximize the performance index, thereby optimizing η_c and $[\text{NO}_x]$. The robustness of the approach is demonstrated and evaluated by initiating a change in load and observing the reaction of the modified control system. The control scheme is shown to effectively increase and maintain overall burner performance. Implementation of robust optimal control to practical systems is discussed in terms of challenges outstanding and opportunities to integrate with overall system performance.

Keywords: Combust control; combustion in practical systems (furnaces; incinerators); environmental combustion (NO_x)

INTRODUCTION

While combustion of fossil fuels provides most of the world's energy, it also produces most of the world's air pollution. One straightforward technique to reduce these emissions from stationary combustion systems is to switch the fuel being burned, substituting a cleaner-burning fuel where a more polluting type is being used (i.e., switching from coal to oil or from oil to natural gas).

Although natural gas combustion generates significantly lower emissions of sulfur oxides and soot than coal or oil, reducing the emission of oxides of nitrogen (NO and NO_2 , collectively referred to as NO_x), a major contributor to photochemical oxidant ("smog"), remains a challenge.

Many techniques are employed in the task of controlling NO_x emissions from stationary combustion applications. Some controls seek to prevent NO_x formation during combustion, such as staged combustion, flue gas recirculation, catalytic combustion, etc. Other control processes destroy NO_x in a post-combustion reaction; these include selective catalytic reduction (SCR), selective non-catalytic reduction (SNCR), and non-selective catalytic reduction (NSCR) (U.S.E.P.A., 1992).

In any of these processes, a set of static input parameters (fuel load, equivalence ratio, etc.) will correspond to particular values for each of a set of output parameters (NO_x emission, heat loading, combustion efficiency, etc.). For a given combustion process there will be at least one set of input parameters that produces an *optimum* set of output parameters. Identifying the input parameters that produce this optimum condition is not trivial. Furthermore, these optimum input parameters will change as boundary conditions vary due to changes in load, fuel type, inlet air properties, or even subtle changes due to equipment degradation.

Using a natural gas-fired, 100,000 Btu/hr, model industrial burner, research at the University of California, Irvine, has shown that certain values of swirl intensity and excess air (equivalence ratio) can significantly reduce the NO_x concentration in the exhaust gases without reducing combustion efficiency (St. John and Samuelsen, 1994).

Previous research in the area of combustion control has dealt with the problem of reducing pressure oscillation (McManus *et al.*, 1993). The present work explores the potential of applying a robust, on-line, active optimization scheme to a combustion process for the control of the emission of nitrogen oxides, over a range of conditions. A specialized optimization algorithm, static by operation, is applied to control the fuel-air mixing process via swirl intensity and excess air in order to optimize burner performance, which is defined in terms of combustion efficiency and NO_x concentration, measured in the exhaust gas.

An *active control* scheme has been postulated in order to continuously monitor NO_x concentration ($[\text{NO}_x]$) and combustion efficiency (η_c), and adjust the fuel-air mixing process to maintain optimum performance of the burner as boundary conditions vary. Given a constant burner geometry, and a set of variable input parameters (in this case swirl intensity, S' , and excess air, EA), the active control system should be able to find some combination of

those parameters such that a desired performance, the optimum condition, is *attained and maintained*.

APPROACH

The approach adopted implements the active control hypothesis in four steps: (1) development of the experiment; (2) definition of performance in quantitative terms; (3) achievement of a "proof-of-concept" phase demonstrating the viability of the active control scheme; and (4) exploration of a more advanced control algorithm and of the control scheme's reaction to a large scale change in boundary conditions (in this case, fuel load).

Experiment

The burner facility and associated control hardware are shown schematically in Figure 1. Given a fixed fuel flow (load), the two inlet parameters (*EA* and *S'*)

FIGURE 1 Burner facility.

are varied by adjusting the amount of air flowing through the axial and swirl air streams. First, the sum of the two air streams ($m_\phi + m_\theta$) determines the overall excess air (EA) provided for combustion. Second, the percentage of flow through the swirl air stream with respect to the total amount of air flow uniquely defines the swirl intensity (S') for this burner.

To facilitate computer control of the air and fuel flow, sensor/valve packages were installed in the fuel line, the axial air line, and the swirl air line. Each sensor and valve is referred to collectively as a mass flow controller (MFC).

Figure 2 is a block diagram representing the general control scheme employed in this study. The burner is the plant, or object under control. The sensor is composed of a bank of gas analyzers, similar to continuous emissions monitoring systems installed in many industrial/commercial burner operations today. A continuous sample is extracted and pulled through five analyzers which measure $[CO]$, $[CO_2]$, $[HC]$, $[O_2]$ and $[NO_x]$.

The controller consists of a 486-based computer which reads emissions signals from the analyzers and calculates NO_x concentration (corrected to 3% O_2), and combustion efficiency. Using a specialized optimization algorithm, the controller determines new values of S' and EA , and sets the air flow accordingly.

Performance Definition

For a given burner, optimum performance can be qualitatively defined as the value of swirl intensity, S' , and excess air, EA , where $[NO_x]$ is relatively low and combustion efficiency remains relatively high. In order to make an

FIGURE 2 General control scheme.

objective evaluation, performance is quantified by a weighted sum called a performance index, which is denoted as J . This performance index is a function of $[\text{NO}_x]$ and η_c . The performance index is defined such that an increase in combustion efficiency and a decrease in NO_x concentration both lead to an increase in J . That is, performance of the burner is considered optimized when J is maximized. The performance index, J , can be defined in terms of *any* measurable parameters of interest. For the purposes of active control demonstration, the performance index is defined as

$$J = g(\eta_c) + f([\text{NO}_x]) \quad (1)$$

$$f([\text{NO}_x]) = \begin{cases} 1 - 0.75 \cdot \left(\frac{[\text{NO}_x]_{3\%}}{[\text{NO}_x]_{\text{limit}}} \right)^4 & ; [\text{NO}_x]_{3\%} \leq [\text{NO}_x]_{\text{limit}} \\ (1 - 0.75) \cdot \frac{[\text{NO}_x]_{\text{max}} - [\text{NO}_x]_{3\%}}{[\text{NO}_x]_{\text{max}} - [\text{NO}_x]_{\text{limit}}} & ; [\text{NO}_x]_{3\%} > [\text{NO}_x]_{\text{limit}} \end{cases} \quad (2)$$

$$g(\eta_c) = \left(\frac{\eta_c - \eta_{c,\min}}{100\% - \eta_{c,\min}} \right)^{1/2} \quad (3)$$

where $[\text{NO}_x]_{\text{max}}$ and $\eta_{c,\min}$ are set according to the ranges expected for a particular burner geometry. The definition of the efficiency function term is such that an increasing reward (i.e., an increase in J) is applied as combustion efficiency increases. The purpose of the piece-wise definition of $f([\text{NO}_x])$ is to impose a high penalty on J above $[\text{NO}_x]_{\text{limit}}$ and a rapidly decreasing penalty (increasing reward) for measurements below this limit. In other words, the contribution from $f([\text{NO}_x])$ to J is high as long as the measured NO_x concentration is below the specified limit. The value of $[\text{NO}_x]_{\text{limit}}$ can have practical significance, such as the permitted NO_x emission limit for a particular burner application. Hence, a burner with perfect performance would yield a value of J equal to 1.0 ($\eta_c = 100\%$, $[\text{NO}_x] = 0$ ppm).

The intention of this performance index is to provide a single variable (a function of the two variables of interest) that may be evaluated and searched in real-time by a computer, using specialized optimization techniques. The way in which the performance index is defined is specific to the particular burner configuration, and its emission character, under study. It should not, in the form presented above, be used as a universal parameter for comparing the performance of different burners, but it may be useful in comparing the performance of different configurations of the same burner.

This performance index, applied to $[\text{NO}_x]$ and η_c measurements taken across the stability limits of one particular nozzle configuration at 100% load, is plotted in Figure 3. The region of optimum performance (maximum J) is indicated by the white band in Figure 3.

Proof-of-Concept

As a demonstration of the viability of the active control approach, a relatively simple problem was considered: Optimize the performance of the burner for a given geometry, at a static load (100%), using a relatively simple and well-understood search algorithm. That is, determine if the active control system can find the optimum of the surface shown in Figure 3, without having knowledge of the shape of that surface, other than the location of the stability limits. The burner geometry incorporates a nozzle that injects the fuel in the same sense as the flow of the swirling air (co-swirl). This problem was presented in a previous work (St. John and Samuelsen, 1994) and is summarized here for completeness.

The optimization technique employed in this demonstration is known as a direction-set technique. Any technique which works from an initial point in a given search space and then optimizes along each of a set of directions within

FIGURE 3 Performance Index, J , as a function of EA and S' . White region indicates area of best performance.

that space can be classified as a direction-set technique. What distinguishes individual techniques within this classification is the process by which those directions are chosen.

The method of steepest descent is one popular direction-set technique. This method is considered a first order technique because it involves calculating or measuring the local gradient and then optimizing along the line in the direction of steepest descent (or ascent). Upon optimization in one direction, the gradient is again determined and optimization proceeds in the new direction. Although powerful, the time-consuming process of measuring the gradient at each turning point prevented exploration of this method in the current research.

The direction-set technique used in the proof-of-concept stage is known as Powell's method, a zero order technique, because it does not require calculation of a gradient. The approach and results of the proof-of-concept demonstration are included in the Appendix.

Following success of this initial demonstration, more practical studies were undertaken. First, a more advanced optimization method, known as the genetic algorithm, was incorporated and applied to the same problem, and same burner geometry, addressed in the proof-of-concept. The results from this study appear in a previous work (St. John and Samuelsen, 1994) and are presented here in the Appendix.

Genetic algorithms represent a radical departure from traditional forms of optimization, such as the direction-set class of search methods. Based on natural selection mechanics, the description of this method requires language borrowed from that field of study. The process starts with a *population* of *individuals*. The *fitness* of each individual is evaluated and individuals are *selected* for reproduction according to each one's fitness: individuals with higher fitness have a better chance of reproduction. Each individual selected for reproduction can be represented by a character string and functions as a *chromosome*. Each chromosome may undergo *crossover* with its mate based on a finite probability that crossover will occur. In addition, each *allele* – represented by a single character in an individual's chromosome (string) – has a small probability of *mutating* (Goldberg, 1989).

In the present context, a population of individuals is comprised of twelve discrete points in *EA* and *S'* search space, either randomly or uniformly distributed, initially. An individual's fitness is that individual's scaled performance index value. Individuals are presented as chromosomes, undergoing crossover and mutation, by coding each point as a ten bit unsigned integer string (ten ones or zeros).

The next step involves application of both search techniques to a single burner geometry, incorporating a step change in load into the experiment, so

that the control system's response to such a change may be evaluated. A second burner geometry was used in this study due to the observation of a more pronounced change in the location of the peak region following a change in load. The geometry of the second burner differs from previous work only in the fuel nozzle; the second nozzle injects fuel against the flow of the swirling air (and is referred to as the counter-swirl nozzle).

A change in load is incorporated into the control trials simply by starting at one load and allowing the system to reach an optimum at that load, then returning the operation of the burner to its initial operating point, initiating a change in load, and allowing the controller to relocate the new region of optimum performance.

RESULTS

Results are submitted in several graphical forms, which merit discussion. First of all is the most basic presentation, absolute performance index value versus iteration. This type of graph is simply a line-plot of performance index evaluations as a function of sequential algorithm iterations. Graphs of this type are generated for both the direction-set and the genetic algorithm techniques. They answer the fundamental question for a particular run: Over time, was performance of the burner improved? These graphs are individually helpful in addressing the ideas of efficacy, or whether or not the control system completed its objective. Two or more of these graphs can be helpful in comparing the efficiency of one technique to another: Did one technique reach a high performance condition sooner than another? This type of graph is referred to as an *absolute performance history curve*.

Another graph type is similar to the first but displays a running average performance rather than an instantaneous performance, versus iteration. This gives an indication of the burner's overall performance for a particular run, and can be used to compare the efficiency of competing optimization techniques. This type of graph is called an *average performance history curve*.

The third type of result presented here is also useful in evaluating the efficacy of a particular direction-set trial. This graph is a contour plot of the performance map, overlaid by the path that a particular direction-set trial run produced. It shows where operation of the burner started, and if the region of measured high performance was indeed attained by the direction-set technique. This plot type is known as a *direction history graph*.

The fourth graph type is similar to the direction history graph but is used for display of the genetic algorithm results. One plot of the performance map is

produced for each generation, with the location of each individual in a particular generation denoted by a data marker. This shows whether or not, over time, the locations of the individuals converged to the region of high performance. It also shows, over several generations, how less-fit individuals are less likely to be selected for reproduction than more-fit creatures. This final graph type is referred to as a *population history* graph.

Results from the proof-of-concept phase, and the initial exploration of the genetic algorithm, are included in the Appendix, and show the ability of the genetic algorithm to find the optimum region of a search space in a practical combustion application. The true test of the active control system, however, is the question of robustness: Can the system respond to a large-scale change in boundary conditions, such as a change in the fuel flow rate?

The absolute and averaged performance history for a representative trial using the direction-set method is presented in Figure 4. The shading change in the plot corresponds to the change in load from 100% to 70% during the trial. This plot shows that, as in the case of the single load, performance of the burner is increased over time. Following the load change, the search begins anew and performance of the burner is again improved. Note that the absolute difference between the value of the performance index at the start and near the end is not as great as it was in the proof-of-concept results (Fig. A2). This difference arises from the fact that the performance of the second geometry (counter-swirl) is better overall than that of the first geometry (co-swirl). There is not as much difference between a low performance condition and a high performance condition for the better performing burner geometry. This difference in scale is also the reason why large

FIGURE 4 Absolute (solid) and average (dashed) performance history, direction-set technique. Shading shift indicates the change in fuel load from 100% to 70%.

excursions are experienced (note, the line is not as flat as it is in Fig. A2) after a peak in performance is reached.

Figures 5a and 5b are the direction history plots for this particular run at 100% and 70% load, respectively. In both cases, search begins in the 153°

FIGURE 5a Typical direction history plot, 100% load.

FIGURE 5b Typical direction history plot, 70% load.

direction, reflects off of a stability limit, and proceeds to another edge of stability. These plots are very telling, for in both cases the global optimum conditions is not reached. The system essentially gets hung up on a local peak, or can not find the direction in which to continue searching (which would be exactly along the stability limit line). This demonstrates two very important weaknesses in the direction-set technique: First, if the controller reaches a local peak, it will remain there; second, if a path of ascent exists from a point, but it is narrow, the algorithm may take a long time finding that path.

The performance history results (absolute and average) for a typical genetic algorithm search with a change in load are shown in Figure 6. Note the jagged nature, which is typical of the genetic algorithm's character. It's difficult to tell much from this plot, other than performance of the burner has been improved over time, and this improvement is repeated following a change in load.

Figures 7a and 7b are the population histories for this same trial for 100% and 70% load, respectively. Some very interesting observations can be made from these figures. First of all, Figure 7a shows the familiar genetic algorithm behavior: starting from a random distribution of twelve individuals, more-fit individuals propagate and less-fit individuals die off. By the fifth generation, in fact, all but three individuals are clustered about the region of the global optimum, where performance is at its peak for this burner. The outlying individuals are due to the mutation effect.

Looking now at Figure 7b, the load has changed (along with the stability limits) and the sixth generation is re-initialized in a random distribution. The survival of the fittest behavior of the genetic algorithm can be followed

FIGURE 6 Absolute (solid) and average (dashed) performance history, genetic algorithm. Shading shift indicates the change in fuel load from 100% to 70%.

FIGURE 7a Typical population history plot, 100% load. White numerals denote generation, gray numerals indicate clustering.

through the generations again, and by the tenth generation, *all* of the individuals are crowded about the region of peak performance.

For a more in-depth comparison of the two techniques, refer to the results summarized in the Appendix.

SUMMARY

The goal of developing an active control approach that will attain operation of a burner at its optimum performance, and that will maintain peak operation

FIGURE 7b Typical population history plot, 70% load. White numerals denote generation, gray numerals indicate clustering.

following a change in boundary conditions, has been realized. The approach presented herein has been shown to adjust burner inlet parameters in order to optimize NO_x emissions and combustion efficiency. This was accomplished by defining a trade-off between combustion efficiency and NO_x concentration in the form of a performance index. The performance index functions as a single search criterion, to which two search techniques were applied in the current work.

Further, the simple genetic algorithm has been shown to be a superior search technique, in the case of continuous combustion optimization, as

compared to a zero-order direction-set method. The main advantage of the genetic algorithm seems to be its ability to locate a global optimum more reliably than a direction-set technique, without the tendency to settle on a local ridge.

While these results are encouraging, care must be taken in applying them to the general problem of practical burner optimization. Certainly, each burner application will have peculiar characteristics that must be accounted for in the development of a control scheme. This research does not indicate that either the direction-set or the genetic algorithm search technique can be applied to the practical control of a burner. Rather, this study should be seen as a first step, providing guidance to future research into the development of a practical application of active control. The achievement of such a goal will benefit from a refined search mechanism, and an improved emissions and stability sensor.

The system developed in the present work continuously searches for the optimum operating condition of a burner, and successfully achieves optimum performance even following a change in load. While the present system optimizes emissions, it does so without any knowledge of that burner's particular emissions character. The only requirement, in the present case, is knowledge of the stability limits. The successful behavior of the control scheme following a large-scale change in boundary conditions (fuel load) implies that the system would respond to smaller-scale changes in boundary conditions as well (fuel composition, equipment degradation, etc.).

Acknowledgements

This program is supported by the California Institute for Energy Efficiency under a Memorandum of Understanding, #4901510, with program liaison provided by Dr. Jim Cole. The Southern California Gas Company is a supporting sponsor as well with program oversight by Mr. Cherif Youssef. The authors gratefully acknowledge the contributions of Steve Dunn, Dave Denny, and especially Matt Miyasato to this effort.

References

- Goldberg, D. E. (1989) *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley Publishing Company, Inc., Massachusetts.
- Powell, M. J. D. (1964) An Efficient Method for Finding the Minimum of a Function of Several Variables without Calculating Derivatives, *J. Computer*, 7(4), pp. 303–307.
- St. John, D. and Samuelsen, G. S. (1994) *Twenty-Fifth Symposium (International) on Combustion*, The Combustion Institute, pp. 307–316.

Summary of NO_x Control Technologies and Their Availability and Extent of Application. U.S. Environmental Protection Agency, Research Triangle Park, NC, Publication NO. EPA-450/3-92-004, February 1992.

McManus, K. R., Poinso, T. and Candel, S. M. (1993) *Prog. Energy Combust. Sci.*, **19**, 1–29.

APPENDIX: POWELL'S METHOD IN 2 DIMENSIONS

Powell's direction-set technique was first proposed as a method for minimizing a mathematical function (Powell, 1964). To explain Powell's method in two dimensions, first consider an objective function, F , which is a function of two variables, X_1 and X_2 , represented by the vector, $\mathbf{X}$. Now consider a search direction, represented by the normalized vector, $\mathbf{S}$. Thus, the direction along the X_1 coordinate would be represented by $\mathbf{S} = (1, 0)$, and a search in the direction 45° to the X_1 and X_2 directions would result if $\mathbf{S} = (1, 1)$. Given an initial position, $\mathbf{X}_i$, and an initial search direction, $\mathbf{S}_i$, a new position, $\mathbf{X}_{i+1}$ is found by maximizing (or minimizing) the objective function along the line defined by the search direction. Computationally, optimization is accomplished in an iterative fashion by finding the (scalar) value of a_i that produces a value of $\mathbf{X}_{i+1}$, which maximizes the objective function, F , according to

$$\mathbf{X}_{i+1} = \mathbf{X}_i + a_i \mathbf{S}_i$$

This is, in fact, the general optimizing strategy for all direction-set methods. The direction-set method chosen for this research is a zero order technique, because it requires evaluation of the objective function (performance index) only. This method is a modified version of a powerful, popular, and well-understood technique known as Powell's method. Figure A1 is a flowchart describing the modified version of Powell's method employed in this study.

In discussions about search processes of this modified type, some clarification should be made regarding numbering convention. A position in the search space is still denoted by the vector, $\mathbf{X} = (X_1, X_2)$, where $X_1 \equiv S'$ and $X_2 \equiv EA$. In the modified approach, each position in the search space has two numbers associated with it, $\mathbf{X}_{ij}$. The first index, i , corresponds to the search iteration, just as in the discussion of Powell's method. For the initial position, $i = 0$. After maximization along the initial direction, α_0 , the new value of i is 1. And so on. The second index, j , refers to steps within a search direction. Hence, the first position is denoted $\mathbf{X}_{0,0}$, and the first evaluation of performance is $J_{0,0}$. Following the first step along α_0 , the j index is incremented by one, but the i index does not change until the search direction changes. So, if it takes ten

decrease (or a stability limit is encountered) three times. Thus, along a unimodal line this search will generate a value of X that results in a maximum value of J along that line.

This technique has the undesirable characteristic of settling on the first peak encountered. However, the hill climbing technique is employed because examination of the character of the J surface reveals an essentially unimodal surface with respect to swirl intensity and excess air. Hill climbing on a unimodal surface ensures that the first maximum encountered along any line is the global maximum on that line.

PROOF-OF-CONCEPT RESULTS

Figure A2 is an absolute performance history plot from a typical direction-set trial of the burner at 100% load, for the first geometry explored (co-swirl nozzle). Note that the performance index starts off at a relatively low value, increases sharply after about the ninth iteration, and settles out to what is essentially the peak value after about fifteen iterations.

The direction history for the same trial displayed in Figure A2 is shown in Figure A3. Keep in mind that the control program has no knowledge of the shape of the performance map (contours) during its operations; the map is shown for reference only. Operation of the burner starts off in a region of relatively low performance. Initialized along a direction of 333°, the search proceeds until a constraint is encountered (zero percent excess air, in this case), at which point the system changes direction and hunts along a path that is 90° to the first direction. Search along this second direction takes the operation of the burner up to the region of peak performance, where performance

FIGURE A2 Typical absolute performance history, direction-set technique.

FIGURE A3 Typical direction history plot.

is essentially optimized. Search continues, however, in different directions, until a slight increase is encountered. The trial was halted after 83 iterations.

Figure A4 is the average performance history curve for this same trial, indicating average performance of the burner up to a given point.

Figure A5 shows both the absolute and the average performance histories of a typical application of the genetic algorithm. Note the relatively erratic trace of the absolute performance index as compared to Figure A2. This is due to the genetic algorithm's evaluation of an entire population at once. The genetic algorithm is not concerned with the order in which evaluations are made. It functions, rather, from the whole population at once, taking all performance values in parallel and generating a new population based on the relative performance of the previous population. Hence, the separation of the genetic algorithm's performance history curve into sections is demarcated by vertical lines in the figure. Although probably of concern in a practical sense, the jagged nature of the absolute performance line is not important to the function of the genetic algorithm. What is important is the average performance, after each entire population. This average performance is shown to increase steadily.

Figure A6 is the population history for this same run. This figure is perhaps more illustrative than Figure A5. The first generation is essentially uniform across the search space. In the second generation, it is clear that some less-fit individuals have died off. By the third generation, there is a clear migration toward the region of peak performance and in the fourth generation all but

FIGURE A4 Average performance history, direction-set technique.

FIGURE A5 Absolute (black) and average (white) performance history, genetic algorithm. Vertical lines and white numerals denote generations.

three of the twelve individuals have converged to the same area of high performance.

The three outliers can be attributed to two sources. One of the rogues, at $EA = 17\%$ and $S' = 0.44$, is simply a stubborn individual. This particular point has been carried through all the way from the initial generation, with little variation. The reason this individual persists is because each individual, even a poorly fit one, has a finite probability of being reproduced—it simply got lucky. It's likely that it would die off in subsequent generations, although it is finitely possible that it would never die off. The other two outliers were produced by a separate mechanism: mutation. Recall that after a population's fitness has been evaluated, individuals are selected for reproduction, then they

FIGURE A6 Typical population history plot. White numerals denote generation, gray numerals indicate clustering.

are potentially crossed-over with their mates, and each bit along an individual's coded binary string is subjected to potential mutation. These two individuals are the product of this process.

Thus, the genetic algorithm has been successful in demonstrating the fundamental requirement of an active control system as well: performance of the burner, under a static fuel load and burner geometry, has been steadily increased over time.

Figure A7 is a plot of both average performance of the direction-set and the genetic algorithm cases. Note that even after 48 iterations the average performance of the direction-set method has not surpassed and the overall performance of the genetic algorithm.

The behavior of the two typical trials discussed so far deserve some "what-if" discussion. From other attempts, the results of which are not presented here, it is clear that the genetic algorithm will perform essentially the same as in this typical case: erratic start, converging after a few generations, but almost always with a few outliers. The direction-set technique, however, could perform quite differently. It should be obvious that this method depends heavily on two initial parameters: the starting location and the starting search direction. One can imagine the trivial search that starts in the peak region.

FIGURE A7 Average performance history for both direction-set (black) and genetic algorithm (white).

In this case, the goal is achieved by default, attaining and maintaining peak operations without any search. Or, consider starting a search in a “poor” performing region, but initializing the search direction headed exactly for the region of peak performance. For this scenario, peak performance would be achieved much earlier than in the results presented.

So, the apparent dominance of the genetic algorithm must be tempered with the statement that there are circumstances where the direction-set could perform more efficiently than the genetic algorithm, but that behaviour is not guaranteed. The genetic algorithm, on the other hand, will provide relatively consistent behavior, and has the added feature of locating the global optimum.

FIGURE A1 Powell's method in two dimensions.

steps along the initial search direction to maximize J , then the final position is denoted $X_{0,10}$, and the final evaluation of J is $J_{0,10}$. The final position along a search direction becomes the initial position in the next search direction. For example, $X_{1,0}$ is given the value of $X_{0,10}$, $J_{1,0}$ is evaluated and the search continues in the next direction, α_1 .

In Figure A1, "maximizing J along the direction α_i " can be accomplished through a variety of methods, including polynomial approximation or the golden section method (if the minimum is bracketed). For practical reasons, the method of line maximization used in this study is a simple hill climbing technique. Starting at a given $X_{i,0}$, the settings of EA and S' are stepped along the direction, α_i . Search proceeds in finite steps as long as the value of J continues to increase. If the performance index value decreases, or if a stability limit is encountered, the search direction is reversed. A new direction is determined and a new search initiated only after the value of J is found to