

Evolutionary Optimization of Feedback Controllers for Thermoacoustic Instabilities

Nikolaus Hansen, André S.P. Niederberger, Lino Guzzella*, and Petros Koumoutsakos**

ETH Zurich, 8092 Zurich, Switzerland

Abstract. We present the system identification and the online optimization of feedback controllers applied to combustion systems using evolutionary algorithms. The algorithm is applied to gas turbine combustors that are susceptible to thermoacoustic instabilities resulting in imperfect combustion and decreased lifetime. In order to mitigate these pressure oscillations, feedback controllers sense the pressure and command secondary fuel injectors. The controllers are optimized online with an extension of the CMA evolution strategy capable of handling noise associated with the uncertainties in the pressure measurements. The presented method is independent of the specific noise distribution and prevents premature convergence of the evolution strategy. The proposed algorithm needs only two additional function evaluations per generation and is therefore particularly suitable for online optimization. The algorithm is experimentally verified on a gas turbine combustor test rig. The results show that the algorithm can improve the performance of controllers online and is able to cope with a variety of time dependent operating conditions.

1 Introduction

Modern gas turbines have to comply with continually more stringent emission regulations (NO_x , CO etc.). This fact led to the development of lean premixed combustion systems. They operate with excess air to lower the combustion temperature which in turn decreases the NO_x levels. The lean regime however makes the combustor prone to thermoacoustic instabilities which arise due to a feedback loop involving fluctuations in acoustic pressure, velocity and heat release. Thermoacoustic instabilities may cause mechanical damage, higher heat transfer to walls, noise and pollutant emissions. This phenomenon is observed also in rocket motors, ramjets, afterburners, and domestic burners. One way to substantially reduce such thermoacoustic instabilities is active control [1, 2]. A feedback controller receives input from pressure sensors and commands a secondary fuel injection. Adjusting the controller parameters into a feasible working regime is an optimization problem with two important properties. First, the stochastic nature of the combustion process leads to a considerable amount of uncertainty in the measurements. Second, changing operating conditions ask for online tuning of the controller parameters.

* A.S.P. Niederberger and L. Guzzella are with the Measurement and Control Laboratory

** N. Hansen and P. Koumoutsakos are with the Institute of Computational Science

Evolutionary algorithms are population-based optimization methods, which are considered to be intrinsically robust to uncertainties present in the evaluation of the objective function. The main reason for this robustness is the use of a population [3, 4]. To improve their robustness against noise either the population size is increased [5, 6] or multiple objective function evaluations of solutions are conducted and an appropriate statistics is taken, usually the mean value. Both methods increase the number of function evaluations per generation typically by a factor between 3 and 100, which is prohibitive for online applications. Consequently, we suggest an alternative approach to optimize the parameters of a Gain-Delay and an \mathcal{H}_∞ controller online with an evolutionary algorithm. A noise handling method is introduced, that distinguishes between noise measurement and noise treatment. The noise measurement is suited for any ranking-based search algorithm, needs only a few additional function evaluations per generation, and does not rely on an underlying noise distribution. The noise measurement is combined with two noise treatments that aim to ensure that the signal-to-noise ratio remains large enough to keep the evolutionary algorithm in a rational working regime.

The next section introduces the noise-tolerant CMA evolution strategy. Section 3 reports experiments on a test rig with the different controller structures for two operating conditions and Section 4 gives a summary.

2 A Noise-Resistant Evolutionary Algorithm

The evolutionary algorithm serves to minimize a time dependent stochastic objective function L (also loss or cost function)

$$L : \mathcal{S} \times \mathbb{R}_+ \rightarrow \mathbb{R}, \quad (\mathbf{x}, t) \mapsto f(\mathbf{x}, t) + N_f(\mathbf{x}, t) . \quad (1)$$

The algorithm is based on a $(\mu/\mu, \lambda)$ Covariance Matrix Adaptation (CMA) Evolution Strategy (ES) [7–9] with the default parameters from [9]. The $(\mu/\mu, \lambda)$ -CMA-ES is predestined because, first, it is a non-elitist continuous domain evolutionary algorithm. Non-elitism avoids systematic fitness overvaluation [3] and possible subsequent failure. Second, the selection is solely based on the ranking of solutions providing robustness in an uncertain environment. Third, the covariance matrix adaptation conducts an effective and efficient adaptation of the search distribution to the landscape of the objective function. Fourth, the CMA-ES can be reliably used with small population sizes allowing for a fast adaptation as it is highly desirable in an online application. Here, we introduce a noise-handling (NH) method that can be applied to any ranking based search algorithm and is combined with the CMA-ES into the NH-CMA-ES. The noise handling preserves all invariance properties of the CMA-ES, but biases the population variance when too large a noise level is detected. The noise measurement and the noise treatment are described in turn.

The noise measurement is based on measured rank changes induced by reevaluations of solutions. The algorithm outputs a noise measurement value s and reads

1. Set $L_i^{\text{new}} = L_i^{\text{old}} = L(\mathbf{x}_i)$, for $i = 1, \dots, \lambda$, and let $\mathcal{L} = \{L_k^{\text{old}}, L_k^{\text{new}} | k = 1, \dots, \lambda\}$, where λ is the number of offspring in the CMA-ES.

2. Compute λ_{reev} , the number of solutions to be reevaluated; $\lambda_{\text{reev}} = f_{\text{pr}}(r_\lambda \times \lambda)$ where the function $f_{\text{pr}} : \mathbb{R} \rightarrow \mathbb{Z}$, $x \mapsto \begin{cases} \lfloor x \rfloor + 1 & \text{with probability } x - \lfloor x \rfloor \\ \lfloor x \rfloor & \text{otherwise} \end{cases}$. If $r_\lambda \times \lambda < 1$ and $\lambda_{\text{reev}} = 0$ for more than $2/(r_\lambda \times \lambda)$ generations, set $\lambda_{\text{reev}} = 1$ to avoid too long sequences without reevaluation.
3. Reevaluate solutions. For each solution $i = 1, \dots, \lambda_{\text{reev}}$ (assuming the solutions of the population are i.i.d., we can, w.l.o.g., choose the first λ_{reev} solutions for reevaluation)
 - (a) Apply a small perturbation: $\mathbf{x}_i^{\text{new}} = \text{mutate}(\mathbf{x}_i, \varepsilon)$ where $\mathbf{x}_i^{\text{new}} \neq \mathbf{x}_i \iff \varepsilon \neq 0$. For the CMA-ES we might apply $\text{mutate}(\mathbf{x}_i, \varepsilon) = \mathbf{x}_i + \varepsilon \sigma \mathcal{N}(\mathbf{0}, \mathbf{C})$, where $\mathcal{N}(\cdot)$ denotes a multi-variate normal distribution and σ and \mathbf{C} are the step-size and the covariance matrix from the CMA-ES.
 - (b) Reevaluate the solution: $L_i^{\text{new}} = L(\mathbf{x}_i^{\text{new}})$
4. Compute the rank change Δ_i . For each chosen solution $i = 1, \dots, \lambda_{\text{reev}}$ the rank change value, $\Delta_i \in \{0, 1, \dots, 2\lambda - 2\}$, counts the number of values from the set $\mathcal{L} \setminus \{L_i^{\text{old}}, L_i^{\text{new}}\}$ that lie between L_i^{old} and L_i^{new} . Formally we have

$$\Delta_i = \text{rank}(L_i^{\text{new}}) - \text{rank}(L_i^{\text{old}}) - \text{sign}(\text{rank}(L_i^{\text{new}}) - \text{rank}(L_i^{\text{old}}))$$

where $\text{rank}(L_i)$ is the rank of the respective function value in the set $\mathcal{L} = \{L_k^{\text{old}}, L_k^{\text{new}} | k = 1, \dots, \lambda\}$.

5. Compute the noise measurement, s . Therefore the rank change value, Δ_i , is compared with a limit $\Delta_\theta^{\text{lim}}$. The limit is based on the distribution of the rank changes on a random function L and the parameter θ (see text). Formally we have

$$s = \frac{1}{\lambda_{\text{reev}}} \sum_{i=1}^{\lambda_{\text{reev}}} \left(2|\Delta_i| - \Delta_\theta^{\text{lim}} \left(\text{rank}(L_i^{\text{new}}) - \mathbb{1}_{L_i^{\text{new}} > L_i^{\text{old}}} \right) - \Delta_\theta^{\text{lim}} \left(\text{rank}(L_i^{\text{old}}) - \mathbb{1}_{L_i^{\text{old}} > L_i^{\text{new}}} \right) \right), \quad (2)$$

where $\Delta_\theta^{\text{lim}}(R)$ equals the $\theta \times 50\%$ ile of the set $\{|1 - R|, |2 - R|, \dots, |2\lambda - 1 - R|\}$, that is, for a given rank R , the set of absolute values of all equally probable rank changes on a random function L (where f and N_f are independent of \mathbf{x}).

6. Re-rank the solutions according to their rank sum, *i.e.* $\text{rank}(L_i^{\text{old}}) + \text{rank}(L_i^{\text{new}})$. Ties are resolved first using the absolute rank change $|\Delta_i|$, where the mean $\Delta_i = \frac{1}{\lambda_{\text{reev}}} \sum_{j=1}^{\lambda_{\text{reev}}} |\Delta_j|$ is used for solutions $i > \lambda_{\text{reev}}$ not being reevaluated, and second using the (mean) function value.

The parameters are set to $r_\lambda = \max(0.1, \frac{2}{\lambda})$, $\varepsilon = 10^{-7}$, and $\theta = 0.2$.

Two noise treatments are used in this paper. First, increase of the evaluation (measuring) time, t_{eval} , for evaluating the controller's performance. Second, increase of the population variance (step-size σ), that can have three beneficial effects. (a) the signal-to-noise ratio is likely to improve, because the population becomes more diverse; (b) the population escapes search-space regions with too low a signal-to-noise ratio, because in these regions the movement of the population is amplified; and (c) premature convergence is prevented. The noise treatment algorithm applied after each generation step uses noise measurement s , and affects step-size σ and evaluation time t_{eval} .

```

 $\bar{s} \leftarrow (1 - c_s) \bar{s} + c_s s$ 
if  $\bar{s} > 0$  % apply noise treatment
  if  $t_{\text{eval}} = t_{\text{max}}$ 
     $\sigma \leftarrow \alpha_\sigma \sigma$ 
   $t_{\text{eval}} \leftarrow \min(\alpha_t t_{\text{eval}}, t_{\text{max}})$ 
else if  $\bar{s} < 0$  % decrease evaluation time
   $t_{\text{eval}} \leftarrow \max(t_{\text{eval}}/\alpha_t, t_{\text{min}})$ 

```

Initialization is $t_{\text{eval}} = t_{\text{min}}$ and $\bar{s} = 0$ and the parameters are chosen to $c_s = 1$, $\alpha_\sigma = 1 + 2/(n + 10)$, $\alpha_t = 1.5$, $t_{\text{min}} = 1\text{s}$, $t_{\text{max}} = 10\text{s}$.

All parameter settings result from the combination of the noise handling with the CMA-ES and simulations on the sphere function. For the combination with different algorithms a different parameter setting might be useful and necessary.

3 Experimental Results

A lab scale test rig is used for the experiments. Preheated air premixed with natural gas flows into a downscaled model for the ALSTOM environmental (EV) swirl burner that stabilizes the flame in recirculation regions near the burner outlet plane. The pressure signal is detected by a water-cooled microphone placed 123 mm downstream of the burner. A MOOG magnetostrictive fuel injector installed close to the flame is used as control actuator. The operating conditions are a mass flow of 36 g/s, a preheat temperature of 700 K, and a ratio of actual to stoichiometric air/fuel ratio of $\lambda = 1.875$ and 2.1. Two controller types are investigated: a simple phase-shift or Gain-Delay controller where gain and delay are optimized by the evolutionary algorithm; and a model-based robust \mathcal{H}_∞ controller where a frequency shift, gain and delay of a previously designed \mathcal{H}_∞ controller [10, 11] are optimized by the evolutionary algorithm.

The cost function to be minimized is the equivalent continuous level of the sound pressure $L_{eq} = 10 \log_{10} \frac{(p_s^2)_{av}}{p_{\text{ref}}^2}$, where $(p_s^2)_{av}$ is the mean squared pressure and $p_{\text{ref}} = 20\mu\text{Pa}$ is the reference pressure. The sound pressure level L_{eq} is acquired from a measurement during t_{eval} seconds with a given controller parameter setting. The total measurement cycle time consists of ramping the controller gain up and down (about 2 s each), pressure data acquisition time $t_{\text{eval}} \in [1, 10]\text{s}$ (determined by the algorithm), data logging (1 s) and NH-CMA-ES computation time (negligible).

Three cost function landscapes for different time intervals are shown in **Fig. 1**, where the combustor is fired up from ambient temperature (cold start) with an air/fuel ratio of $\lambda = 2.1$ switched to $\lambda = 1.875$ after 4800 seconds, and the Gain-Delay controller is switched on. A trend to less negative values for the gain with the heating up becomes apparent (left versus middle figure) and the general background noise level rises (indicated by areas getting darker). Also the parameters evaluated are narrowed down to the small black polygon. The operating condition at $\lambda = 1.875$ (right) exhibits less thermal drift. The algorithm finds a new minimum, where the gain is more negative.

Spectra achieved with the optimized Gain-Delay and \mathcal{H}_∞ controllers are compared to the uncontrolled plant in **Fig. 2**. They are shown for the plant which has been running for several hours and is thus heated up. For $\lambda = 1.875$ (left) the L_{eq} of the uncontrolled plant is 148.72 dB, the Gain-Delay controller reduces it to 146.67 dB, while the




Fig. 1. Cost function landscapes for different time intervals. Pentagrams show the best parameter set obtained from NH-CMA-ES for each generation, the larger they are, the later they have been acquired. The black polygon is the convex hull of all controller parameter values tried in the given time range. Function evaluations, left: 1-150 (0-1300 s); middle: 325-390 (3800-4800 s); right: 395-900 (4900-9800 s). The landscapes are obtained by Delaunay triangulation of a second-order polynomial fit to L_{eq} values for individual delay slices.




Fig. 2. Comparison of the pressure spectra for the uncontrolled, Gain-Delay controlled and \mathcal{H}_∞ controlled plant. Both controllers are NH-CMA-ES optimized. Left: $\lambda = 1.875$, right: $\lambda = 2.1$.

\mathcal{H}_∞ controller reaches 146.16 dB, which is about 15% less. For $\lambda = 2.1$ the values of L_{eq} are 159.87 dB, 147.48 dB and 147.35 dB, respectively. Here the \mathcal{H}_∞ controller performs only slightly better than the Gain-Delay controller, but the control signal contains about 10% less energy.

4 Summary

This study has investigated feedback controllers for secondary fuel injection used on a test rig designed to study thermoacoustic instabilities. To allow for best controller performance in changing operating conditions a self-tuning controller is applied. The main difficulty in optimizing the controller parameters is the uncertainty inherent in the pressure measurements. For this reason a novel noise-handling algorithm is introduced that can be applied to any ranking-based optimization algorithm. The noise-handling algorithm consists of a noise measurement and a noise treatment and is applied to the CMA evolution strategy (NH-CMA-ES), where it preserves all invariance properties

of the original algorithm. In combination with the CMA-ES two additional function evaluations per generation are sufficient to establish a functional noise measurement.

Parameters of Gain-Delay and \mathcal{H}_∞ controllers have been optimized online with the introduced NH-CMA-ES while the combustor was running. The experiments show that the algorithm can optimize different controller types and can cope with changing operating conditions and high levels of noise. Model-based \mathcal{H}_∞ controllers perform best, and can be improved further through the use of the NH-CMA-ES. The optimized solutions deviate remarkably from the originally designed solutions and can make up for uncertainties in the model-building and design process, as well as for time-varying plant characteristics.

Acknowledgments

Support by Daniel Fritsche during the experimental phases and assistance by Caroline Metzler for the technical illustrations is gratefully acknowledged. The authors thank David Charypar for the valuable suggestions and Stefan Kern for providing supporting data. Fruitful discussions with Bruno Schuermans and financial support from ALSTOM (Switzerland) Ltd. are gratefully acknowledged.

References

1. Lieuwen, T., Yang, V.: Combustion Instabilities in Gas Turbine Engines: Operational Experience, Fundamental Mechanisms, And Modeling. Progress in Astronautics and Aeronautics, Vol. 210. AIAA (2005)
2. Dowling, A.P., Morgans, A.S.: Feedback control of combustion oscillations. *Annual Review Of Fluid Mechanics* **37** (2005) 151–182
3. Arnold, D.V.: Noisy Optimization with Evolution Strategies. Volume 8. Kluwer, Boston (2002)
4. Beyer, H.G., Arnold, D.: Qualms regarding the optimality of cumulative path length control in CSA/CMA-evolution strategies. *Evolutionary Computation* **11**(1) (2003) 19–28
5. Arnold, D.V., Beyer, H.G.: Local performance of the $(\mu/\mu_I, \lambda)$ -ES in a noisy environment. In Martin, W., Spears, W., eds.: *Foundations on Genetic Algorithms FOGA*, Morgan Kaufmann (2000) 127–142
6. Harik, G., Cantu-Paz, E., Goldberg, D.E., Miller, B.L.: The gambler’s ruin problem, genetic algorithms, and the sizing of the populations. *Evolutionary Computation* **7** (1999) 231–253
7. Hansen, N., Ostermeier, A.: Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation* **9**(2) (2001) 159–195
8. Hansen, N., Müller, S.D., Koumoutsakos, P.: Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evolutionary Computation* **11**(1) (2003) 1–18
9. Hansen, N., Kern, S.: Evaluating the CMA evolution strategy on multimodal test functions. In Yao, X., et al., eds.: *Parallel Problem Solving from Nature - PPSN VIII*, LNCS 3242, Springer (2004) 282–291
10. Skogestad, S., Postlethwaite, I.: *Multivariable Feedback Control: Analysis and Design*. John Wiley and Sons Ltd., Chichester, New York (1996)
11. Niederberger, A.S.P., Schuermans, B.B.H., Guzzella, L.: Modeling and active control of thermoacoustic instabilities. In: *16th IFAC World Congress*, Prague (2005)