

Extremizing Feedback Control of a High-Speed and High Reynolds Number Jet

Aniruddha Sinha^{*}, Kihwan Kim[†], Jin-Hwa Kim[‡], Andrea Serrani[§], and Mo Samimy^{**}
Gas Dynamics and Turbulence Laboratory, The Ohio State University, Columbus, OH, 43235

We present results of the development and application of extremizing feedback control to high-speed and high Reynolds number axisymmetric jets. In particular, we demonstrate control authority on the near-field pressure of a Mach 0.9 jet with a Reynolds number based on jet diameter of 7.8×10^5 . Open-loop forcing experiments are presented wherein localized arc filament plasma actuators are shown to have two distinct effects on the near-field pressure, similar to their effect on the far-field acoustics reported earlier. At low forcing Strouhal numbers (St_{DF} 's) near the jet column mode instability, a large amplification in the pressure fluctuations is observed. At higher St_{DF} 's (close to the initial shear layer instability) a broad attenuation is observed in the near-field pressure fluctuations, especially in the axisymmetric mode. Previous experiments have shown that forcing the jet with these low and high frequencies result in jet mixing enhancement and far-field noise reduction, respectively. Two different gradient-free extremizing feedback control algorithms have been implemented, each of which can perform online minimum-seeking as well as maximum-seeking. Both methods demonstrate fast convergence to the optimum followed by steady operation.

Nomenclature

a_n, a_1, a_{min}	= generic, initial, and minimum value of step size in a stochastic optimization process
<i>AgeLimit</i>	= user-specified age limit of successively retained best vertex in simplex-based searches
B, C, C_1, C_2, E	} = various designated simplex vertices pertinent to simplex-based searches
R, S_1, S_2, W	
c_n, c_1, c_{min}	= generic, initial, and minimum value of perturbation in a stochastic optimization process
D	= nozzle exit diameter
f_F	= forcing frequency
h	= generic cost-function for optimization
I/O	= input/output, typically referred to a plant in a control system
KWA, mKWA	= original and modified Kiefer-Wolfowitz algorithm for stochastic optimization
LAFPA	= localized arc filament plasma actuator
<i>MaxS, MinS</i>	= user-specified simplex size constraints for simplex-based searches
m_F	= azimuthal mode number of forcing
NMA, mNMA	= original and modified Nelder-Mead algorithm for direct search
N_F	= number of actuators arranged in a symmetric azimuthal array at the jet lip
p	= near-field pressure signal
$p^{[m_p]}$	= near-field pressure signal filtered at azimuthal mode m_p
Re_D	= Reynolds number based on jet nozzle exit diameter and exit velocity

^{*} Graduate Student, Gas Dynamics and Turbulence Laboratory, AIAA Student Member.

[†] Postdoctoral Researcher, Gas Dynamics and Turbulence Laboratory; Currently, with POSCO Technical Research Laboratories, Korea .

[‡] Research Associate, Gas Dynamics and Turbulence Laboratory, AIAA Member.

[§] Associate Professor, Dept. of Electrical and Computer Engineering, AIAA Member.

^{**} The Howard D. Winbigler Professor of Engineering, Director of Gas Dynamics and Turbulence Laboratory, Dept. of Mechanical Engineering, AIAA Fellow, Corresponding author, samimy.1@osu.edu

S	= set of vertices in simplex-based searches
SISO	= single-input single-output control system
S_{s_0}	= initial simplex size for simplex-based searches
St_D	= Strouhal number based on spectral frequency, f , and nozzle exit diameter = fD/U_j
St_{DF}	= Strouhal number based on forcing frequency and nozzle exit diameter = $f_f D/U_j$
t	= time
U_j	= jet exit velocity
x	= streamwise distance downstream from nozzle exit
y_k	= independent variable at the k^{th} step of a stochastic optimization process
α, β, γ	= various simplex transformation coefficients pertinent to a simplex-based search
λ	= exponential decrement for perturbation levels in a stochastic optimization process

I. Introduction

WE present experimental results of feedback control of a high-speed high Reynolds number jet. In particular, we demonstrate control over the near-field pressure with two different goals: increasing the pressure fluctuations for mixing enhancement, or decreasing the fluctuations for potential far-field noise attenuation. In this section we will motivate the discussion and establish the background.

A. Motivation

The feedback control strategy implemented is shown in Fig. 1. An axisymmetric jet was operated at Mach 0.9 with $Re_D = 7.6 \times 10^5$. The jet was forced using a localized arc filament plasma actuator (LAFPA) system consisting of eight electrode pairs placed azimuthally near the jet nozzle exit. Each actuator generates square-pulse-like localized heating with fixed amplitude; the forcing frequency, the azimuthal mode, as well as the duty cycle can be independently controlled. The forcing frequency was chosen as the control input for implementing the feedback controller. The near-field pressure was measured using an azimuthal ring array of eight transducers. The azimuthal-mode-filtered pressure fluctuation level was estimated in real-time and formed the plant output to be fed back to the controller. The controller, in turn, provided the appropriate forcing frequency to the actuators based on an optimization algorithm. The control scheme used can be setup either to seek the maximum fluctuation level (corresponding to mixing enhancement) or the minimum level (corresponding to potential noise attenuation). In the rest of this section, each of the choices made above will be rationalized based on previous research.

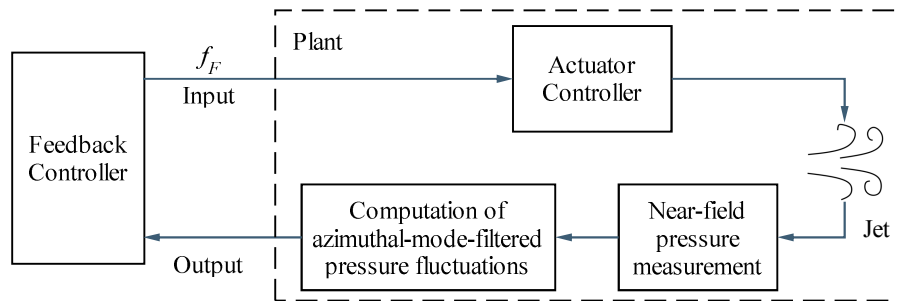


Figure 1. Feedback control system for a jet using plasma actuators.

B. Axisymmetric Jet Mixing Layer and Far-field Noise

The turbulent axisymmetric jet is known to be susceptible to three different instability mechanisms: (i) initial shear layer instability, (ii) jet column mode instability, and (iii) azimuthal mode instability. The initial shear layer instability amplifies disturbances in a narrow range of frequencies that scale with the boundary layer momentum thickness at the nozzle exit; the associated Strouhal number is ~ 0.01 to 0.02 . The initial waves created thereby at the nozzle lip roll up into large-scale coherent structures. These structures entrain fluid into the shear layer from both the high speed potential core and the ambient irrotational near-field and play a major role in the bulk mixing of the fluids. The jet column mode instability determines the preferred passage frequency of large-scale structures at the end of the potential core; this frequency scales with the nozzle exit diameter and the associated Strouhal number is $St_D \sim 0.2$ to 0.6 . Recent experimental results in our laboratory for high Reynolds number and Mach 0.9 and 1.3 jets show that $St_D \sim 0.3^{1,2}$. The azimuthal instability causes certain helical structures to dominate in the mixing layer. The reader is referred to Samimy *et al.*¹ for a detailed review of jet mixing layer instabilities.

Since the discovery of large-scale coherent structures in the late 1960s and early 1970s, researchers have recognized their importance as potential contributors to the radiated far-field noise in addition to the random turbulence. However, in spite of sustained effort ever since, the exact mechanism of influence has not been pinned down as yet. The following source mechanisms have emerged as popular candidates though: vortex-pairing, wavy-wall type mechanisms, and vortex eigen-oscillations. All these effects are believed to be active in the initial mixing-layer region of the flow. An additional postulated mechanism is the violent intermittent events associated with the collapse of the potential core. A recent review article by Jordan & Gervais³ provides a detailed account of these developments. An important character of the far-field noise is its directivity. In particular, the noise radiated in the downstream direction is louder than the noise radiated at right angles to the jet exit. Further, large-scale coherent structures are implicated in the radiation of noise preferentially in the downstream direction (see Tam⁴).

From the above discussion, it is clear that the large-scale coherent structures in the jet shear layer play an important role in bulk mixing as well as noise radiation. Feedback control attempts to manipulate these structures to achieve either one of two desired goals: mixing enhancement or potential noise attenuation.

C. The Near-Field Pressure and its Correlation with Bulk Mixing and Far-field Noise

In general, jet actuation techniques strive to alter the characteristics of the initial shear layer since this latter is most receptive to perturbations; for feedback control, a real-time measurement of this effect of forcing is needed. Present technology does not allow real-time 3D measurement of the velocity field in the jet mixing layer. It is also not practical to require direct sensing of far-field noise levels for control applications. In this subsection, we invoke past research efforts to show that the pressure in the incompressible irrotational near-field of the jet offers an estimate of both the mixing-layer velocity field as well as the far-field noise. The measurement of the near-field pressure presents some unique opportunities: (i) it is a relatively non-intrusive technique; (ii) pressure, being a scalar variable, is considerably easier to measure than the velocity field inside the flow; and (iii) unlike some of the velocity measurement tools like PIV and LDV, pressure sensors are robust and can be feasibly transferred from the laboratory setting to real-world applications.

The mean-square pressure in the irrotational near-field of the jet is related to the Reynolds stresses in the mixing layer by a weighting function that is inversely proportional to the 4th power of the wavenumber.⁵ The nature of this rapid roll-off indicates that the pressure spectrum will be dominated by the larger turbulent scales; this is termed wavenumber filtering. It is well known that the near-field pressure consists of two distinct components – hydrodynamic and acoustic (see the review by Jordan & Gervais³ and the references therein). The hydrodynamic pressure fluctuations carry the convective footprint of the underlying turbulence in the mixing layer, whereas the acoustic pressure fluctuations are characterized by spherical waves propagating with sonic speed. The far-field is dominated by the acoustic component due to the rapid decay of the hydrodynamic component with distance. However, in the immediate periphery of the jet, the hydrodynamic component supersedes the acoustic component. Due to the wavenumber filtering effect, the hydrodynamic pressure itself is dominated by the large-scale coherent structures in the jet mixing layer. Thus, in order to detect the behavior of the actuation-modified large-scale structures, the pressure sensors must be placed close to the shear layer without actually being in the rotational field.

An efficient way of looking at the near-field pressure is via its Fourier azimuthal modal decomposition. Hall *et al.*⁶ measured the near-field pressure of a Mach 0.85 and $Re_D = 9.8 \times 10^5$ jet using an azimuthal array of 15 transducers; they found the azimuthal pressure spectrum to be low-dimensional. In particular, they showed that the instantaneous pressure signal $p(t)$ can be almost completely reconstructed from the sum of the axisymmetric mode $p^{[0]}(t)$ and the first helical mode $p^{[1]}(t)$. The energy of the lower azimuthal pressure modes was found to increase with downstream distance up to $x/D \sim 3$ and decrease subsequently up to $x/D \sim 6$, before increasing slightly again. The previous discussion has already showed that the stronger the large-scale structures in the mixing layer, the higher the pressure fluctuations in the near-field. Since stronger coherent structures are associated with enhanced entrainment and thereby bulk-mixing, this justifies the choice of RMS ($p^{[0]} + p^{[1]}$) in the near-field for real-time estimation of the level of mixing in the jet shear layer.

Hall *et al.*⁶ also found that the near-field $p^{[0]}(t)$ signal correlates better with the far-field noise than just the pressure signal by itself. In particular, they reported a maximum normalized correlation coefficient of 0.34 between the signal from a far-field microphone placed at a polar angle of 30° and the $p^{[0]}(t)$ signal at $x/D = 7.5$. (For later reference, the corresponding correlation coefficient at $x/D = 3$ was 0.11.) In the present work, we do not probe this relationship further; instead we use the foregoing discussion to justify the choice of RMS ($p^{[0]}$) in the near-field as a surrogate for the far-field noise.

D. Control of Axisymmetric Jets

Flow control is usually divided into two general categories: passive and active. Passive control does not add energy to the flow and is normally accomplished by geometric modifications. In active control, energy is added to the flow to excite inherent instabilities or generate new structures (e.g., streamwise vortices). Active control is further divided into open-loop and closed-loop. In open-loop control, the actuation takes place based on an operator's command or a predetermined input. In closed-loop (or feedback) control, information from a sensor or sensors in the flow, possibly along with a flow model, guides the actuation process⁷⁻¹⁰.

The reader is referred to Samimy *et al.*² for a survey of the previous research in open-loop control of axisymmetric jets. We note here that the majority of the investigations using open-loop control have been carried out in relatively low-speed and low Reynolds number jets. As the speed and the Reynolds number of the jet increase, so do the background noise, the instability frequencies, and the flow momentum. Therefore, actuators must provide excitation signals of much higher amplitude and frequencies. We have recently developed a class of plasma actuators called LAFPAs that can provide excitation signals of high amplitude and high bandwidth for high-speed, high Reynolds number flow control^{1,2,11,12}. These actuators are deployed in an azimuthal array just upstream of the jet nozzle exit. The firing of each actuator is governed by a separate square-pulse-train whose frequency, phase, and duty cycle can be controlled independently. The intense localized heating perturbations generated by the LAFPAs can be used to selectively excite the three different instability mechanisms mentioned in Section I-B.

Open-loop forcing experiments were conducted on a Mach 1.3 and $Re_D = 1.1 \times 10^6$ axisymmetric jet to assess the effect of LAFPAs on the mixing layer¹. Laser-based planar flow visualizations, pressure measurements inside the mixing layer, and two-component PIV measurements were used to evaluate the effects of forcing. The jet responded to the forcing over the entire range of frequencies, but the response was optimal (in terms of generating large-scale coherent structures and hence mixing enhancement) around the jet column mode Strouhal number of $St_{DF} = 0.33$. The jet also responded to the various forcing azimuthal modes that could be explored with eight actuators, viz. $m_F = 0, 1, 2, 3, \pm 1, \pm 2$, and ± 4 . Forcing the jet with $m_F = \pm 1$ at the jet column mode Strouhal number provided the maximum mixing enhancement, with a marked reduction in the jet potential core length and a significant increase in the jet center-line velocity decay rate beyond the end of potential core. The flow visualization, growth and decay of perturbations observed in pressure measurements, and PIV data, all together showed that the plasma actuators have control authority over such a high Reynolds number and high-speed flow.

Experiments were also conducted on a Mach 0.9 and $Re_D = 7.6 \times 10^5$ jet to study the effect of LAFPAs on far-field noise²; this is also the jet on which we implemented closed-loop control. The far-field jet noise was measured using two microphones located at polar angles of 30° and 90° relative to the jet axis. The first noticeable effect of forcing was the appearance of the forcing tones and its harmonics in the acoustic spectra, more prominently for $m_F = 0$ than for $m_F = 3$. More pertinently, all forcing azimuthal modes explored resulted in the attenuation of low-frequency noise attended with an amplification of high-frequency noise, at both measurement stations. The directive nature of the far-field noise was also in evidence with the SPL at 30° being at least 10 dB higher than that at 90° .

A standard metric for comparing broadband noise levels is the overall sound pressure level (OASPL). Samimy *et al.*² compared the OASPL for the forced jet to the baseline case (the difference is denoted $\Delta OASPL$) at the two afore-mentioned stations, for a range of forcing frequencies and azimuthal modes. Several noteworthy features were observed. (i) A well-defined attenuation of noise was evident at high forcing frequencies for all azimuthal modes at both stations. (ii) The forcing Strouhal number corresponding to the minimum OASPL was independent of the forcing azimuthal mode, but varied with the measurement station. At the 30° microphone location, a relatively sharp minimum in $\Delta OASPL$ of ~ -1.2 dB was noted at $St_{DF} \sim 2.0$; at the 90° location, a broader minimum of ~ -0.6 dB occurred at $St_{DF} \sim 3.5$. (iii) The highest azimuthal forcing mode employed ($m_F = 3$) was found to be more effective for attenuating noise, compared to the other m_F 's. (iv) An amplification of noise was also noted at low forcing frequencies for all azimuthal modes; however, in this regard, the different azimuthal forcing modes had markedly distinct effects.

The above review of open-loop forcing experiments shows that the in-house developed plasma actuators have significant ability to either enhance mixing or attenuate far-field noise, as desired. We next discuss the development and implementation of extremizing feedback control for these tasks.

E. Schemes for Online Optimization

The previous results reviewed above suggest that the noise attenuation problem can be cast as an optimization problem wherein we seek the forcing frequency of the plasma actuators that minimizes the RMS ($p^{(0)}$) in the near-field. Likewise, the mixing enhancement problem can be posed as a problem of seeking the maximum of RMS ($p^{(0)} + p^{(1)}$) in the near-field. In either case, the optimum forcing frequency will be a function of the operating conditions; for instance, the Mach number and temperature ratio. Open-loop control is unable to cope with this uncertainty in

operating conditions since it operates at a preset forcing frequency. On the other hand, feedback control can add robustness to the performance in the presence of variations in the operating conditions.

We have implemented two different feedback controllers for online optimization of the plant output. Here, “plant” means the portion of the closed-loop system inside the dashed box Fig. 1. Optimization can mean either maximization or minimization and both would be pursued in this work. We will discuss the controllers from the perspective of maximization; for minimization one simply performs maximization on the negative of the plant output. In the past we have implemented a gradient-based extremum-seeking control strategy¹³; although its steady state operation compared well to the optimal open-loop results, the convergence (transient behavior) was found to be very slow. This prompted the exploration of the gradient-free algorithms that are the subject of this article. Most feedback controllers require a dynamic model of the actual system; however both our controllers are model-free and assume a static relationship between the plant input and output.

Teel & Popović¹⁴ showed that the tools of nonlinear programming can be employed for online optimization of a large class of dynamic systems with appropriate assumptions. Using this idea as the starting point, we investigated the field of offline optimization in search of an appropriate algorithm for our application. In particular, we were looking for methods that are gradient-free, robust to noise, and implementable for real-time operation. Among the plethora of available algorithms in this area, Nelder & Mead’s simplex-based direct search method¹⁵ has a proven record of performance in a deterministic setting; surprisingly, the algorithm lacks a convergence proof. Barton & Ivey¹⁶ recommended modifications to the original algorithm for robust behavior in the presence of noise. Torczon¹⁷ designed another direct search technique, the Multi-Directional Search Method, for provable convergence in a deterministic setting. She also provided empirical evidence of reliable optimization of noisy cost functions. Our first controller implementation is a direct search algorithm based on all the above results with several additional modifications for use in online optimization; the details are discussed in Section V-D.

The second controller that we implemented was a modified stochastic optimization technique. In 1951, Robbins and Monro¹⁸ pioneered the field of stochastic approximation, the study of search techniques that successfully reach a pre-assigned goal in spite of noise. Their routine was designed to find the root of a noisy function. In 1952, Kiefer and Wolfowitz¹⁹ adapted the idea to the problem of finding the extremum of a unimodal function obscured by noise. Although the original algorithm was gradient-based, a later variant was developed for gradient-free stochastic optimization²⁰. The additional modifications that had to be incorporated to adapt this offline technique for online optimization will be detailed in Section V-E.

We also briefly tried some other algorithms for our application, only to reject them for some perceived shortcoming. The self-tuning methodology²¹ seemed to be particularly appropriate for our application; however the noise in our application proved to be overwhelming. We also considered the triangular-search method of Zhang²² and found it to be quite similar to our implementation of the direct-search algorithm. However, Zhang’s method lacks the variety of parameters that allowed us to fine-tune the latter for our specific application.

This brings us back to the motivation for this paper. The axisymmetric mode fluctuations of the near-field pressure, being well correlated to the far-field noise, will be treated as the plant output for potential feedback minimization of far-field noise. On the other hand, the fluctuations of the sum of the axisymmetric and first helical modes of near-field pressure would be maximized online for potential mixing enhancement. In either application, we will use one of the two controllers discussed above to determine the optimal forcing frequency in real-time. The LAFPA control system will use this information to force the jet. The effect of the forcing will be measured with pressure sensors in the irrotational near-field of the jet, which will be processed and fed back to the controller.

II. Experimental Setup

A. Flow Facility

All experiments were conducted in the Gas Dynamics and Turbulence Laboratory (GDTL) at The Ohio State University. The ambient air was compressed, dried, and stored in two cylindrical 36 m³ tanks at up to 16 MPa. The compressed air was supplied to the stagnation chamber and conditioned before entering into a nozzle. We used an axisymmetric converging nozzle with an exit diameter of $D = 25.4$ mm (1.0 inch) and operated at a Mach number of 0.9. The air was discharged horizontally through the nozzle into an anechoic chamber.

The Reynolds number of the jet based on the jet diameter was $Re_D = 7.6 \times 10^5$ for the Mach 0.9 jet. The boundary-layer at the exit of the nozzle was very thin, making it challenging to obtain a boundary-layer profile to determine its momentum thickness and its state. Kastner *et al.*²³ used a similar converging nozzle and measured a few points within the boundary layer. They estimated the boundary layer at the exit to be turbulent, with a thickness of about 1 mm and a momentum thickness of about 0.1 mm. The characteristics of the boundary layer in the current experiments are expected to have been quite similar.

B. Plasma Actuators and Plasma Generator System

The plasma generator system is detailed elsewhere^{1,2,11,12}; here we briefly describe some of its highlights. Each plasma actuator (LAFPA) consists of a pair of 1 mm diameter tungsten pin electrodes, with a tip separation of 3 mm. The electrodes are symmetrically distributed around the perimeter of a boron nitride nozzle extension, approximately 1 mm upstream from the extension's exit plane. A 0.5 mm deep and 1 mm wide ring groove is used to house the electrodes and to shield and stabilize the plasma.

The plasma generator enables simultaneous powering of up to 8 actuators with independent frequency, duty cycle, and phase control. Each actuator is connected in series with a fast-response high-repetition-rate high-voltage MOSFET switch, two 15 k Ω ballast resistors, and a 10 kV - 1 A DC power supply. Two of these power supplies are used to energize the 8 actuators. A Labview program controls the switches using square pulses through an eight-channel National Instrument (NI) PCI-6713 DAC card running at an update rate of 780 kHz. The switches are capable of producing high-voltage pulses at repetition rates from a few hertz up to 200 kHz, with a very short pulse rise/fall time ($\sim 0.1 \mu\text{s}$) and a variable duty cycle (from 0 to 100%). Previous experiments have shown that optimal performance is achieved by selecting the duty cycle as a function of the forcing frequency

$$\text{duty cycle (\%)} = \begin{cases} 0.0006 f_F + 2, & \text{if } f_F \leq 30,000 \text{ Hz} \\ 0.0002857 f_F + 11.4286, & \text{if } f_F > 30,000 \text{ kHz} \end{cases} \quad (1)$$

This formula is implemented in the Labview program. The jet can be forced in helical azimuthal mode $m_F = 0, 1, \dots, N_F / 2 - 1$ by phase-shifting the pulse signals to successive actuators by $2\pi m_F / N_F$ radians. The system can also simulate mixed mode forcing; however these were not employed in the current experiments.

In closed-loop experiments, a dSpace 1103 DSP board operating at 50 kHz sampling rate was used to implement the control algorithm. In such cases, the Labview program was modified to acquire the dictated value of the forcing frequency from the DAC output of the dSpace board using one channel of an NI PCI-6036E ADC card.

C. Near-field Pressure Measurements

Figure 2 shows a symmetric circular array of 8 Kulite pressure transducers (model XCQ-062-25A). The highest azimuthal pressure mode that can be distinguished with this configuration is $p^{[3]}$; this was deemed acceptable in view of the low-dimensionality of the near-field pressure discussed previously. The inner diameter of the ring that holds the transducers is 254 mm (10 inches), ten times larger than the nozzle diameter. A maximum of six additional transducers can also be arranged in a linear array aligned with the jet axis with consecutive sensors separated by 12.7 mm (0.5 inch). This linear array was not employed in the current experiments.

During baseline (unforced) experiments and open-loop forcing cases, the pressure signals were amplified, low-pass filtered at 100 kHz, and acquired using an eight channel NI PCI-6143 ADC card at a sampling rate of 200 kHz. In closed-loop experiments, the dSpace board was required to acquire the pressure signals for feedback. In these cases, each of the 8 channels of output from the 100 kHz low-pass filter was split in two. One signal from each splitter was connected to the NI board for offline computation of the SPL spectrum in steady-state operation. The other set of signals were again low-pass filtered at 25 kHz before being routed to the input ADC channels of the dSpace board.

For all computations of the pressure spectrum, a Labview data capture routine collected 786,432 samples at 200 kHz on the eight channels simultaneously, for a total duration of ~ 3.9 s. The spectra were calculated using a window size of 8192 samples with a 50% overlap, resulting in a frequency resolution of 24.4 Hz. The reference pressure for conversion to dB was the standard value of 20 μPa .




Figure 2. Schematic of near-field pressure sensing system using a circular array of eight and a linear array of six pressure transducers.

III. Baseline Experimental Results

One of the first issues to be addressed in our feedback control strategy was the positioning of the azimuthal ring array of pressure sensors. In practical applications, one would like to place the sensors as close as possible to the nozzle exit for ease of implementation. However, as mentioned earlier, Hall *et al.*⁶ reported that the correlation of near-field pressure to the far-field noise is low at the jet exit and reaches a maximum at $x/D = 7.5$. To balance these opposing constraints, we chose to locate our ring array at $x/D = 3.0$. We have already pointed out that at this location both the axisymmetric and first helical pressure modes have the maximum energy compared to all other streamwise locations⁶. Recall that at this location the correlation coefficient between the near-field pressure and far-field noise is around 0.11. The choice for the tip-circle radius of the pressure sensors is also a compromise between conflicting constraints. On the one hand, the pressure sensors should not be within the jet to reduce measurement error. On the other hand, as already mentioned, the sensors should be close enough to the mixing layer to capture a strong signature of the hydrodynamic fluctuations in the jet. Kim *et al.*²⁴ performed PIV studies of the same jet in various open-loop forcing experiments; their results show the extent of the jet mixing layer in various forcing conditions. In particular, the maximum radius of the jet at $x/D = 3.0$ was observed when the jet was forced at its column mode frequency and the first flapping azimuthal mode $m_F = \pm 1$; this radius was about $0.9 D$. Hence, the tip-circle radius was chosen as $1 D$. All near-field pressure data presented in this paper were measured at this location.

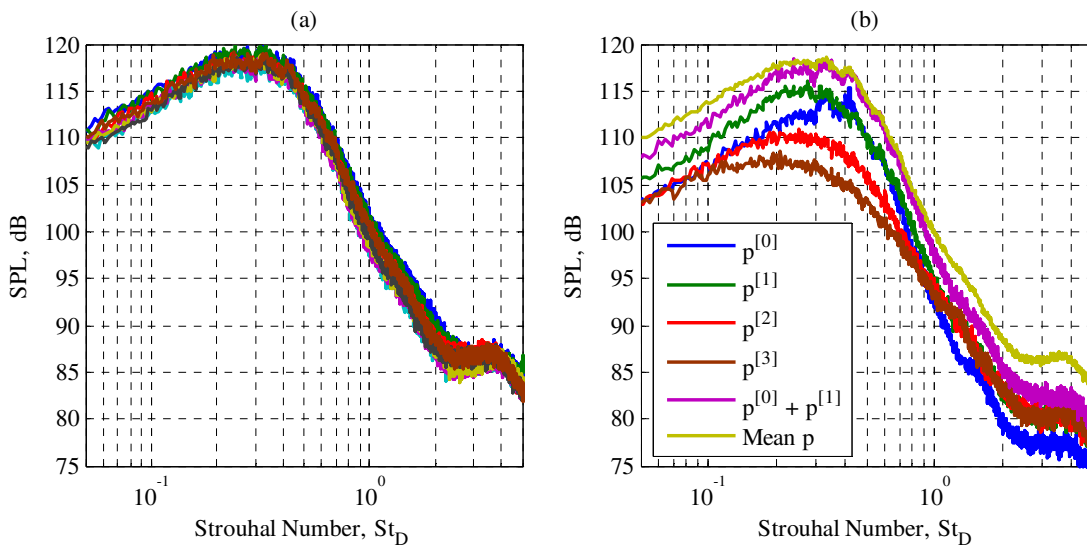


Figure 3. Baseline near-field pressure spectra for (a) the eight individual channels, and (b) the various azimuthal modes.

Figure 3 presents the SPL spectra for the baseline (unforced) case. In Fig. 3(a), we show that the 8 channels record almost identical SPL spectra, verifying that the jet is axisymmetric and the azimuthal array of pressure sensors is coaxially aligned to it. Henceforth, we use “SPL of the pressure” to denote the mean of the 8 channels. Figure 3(b) shows the SPL spectra of the first four azimuthal-mode-filtered pressure signals; this is the most that we can resolve with 8 sensors. At this streamwise location, $p^{[1]}$ is the strongest, followed by $p^{[0]}$. As mentioned earlier, Hall *et al.*⁶ have reported that the instantaneous pressure signal can be almost completely reconstructed from the $(p^{[0]} + p^{[1]})$ signal. In Fig. 3(b), we plot the SPL spectrum of this signal and overlay the mean pressure spectrum; they are seen to have very similar peak amplitudes. Note that the peak pressure fluctuation is around $St_{DF} = 0.3$ which corresponds to the jet column instability frequency.

IV. Open-loop forcing results

The purpose of the open-loop experiments was two-fold. We wished to investigate the effect of forcing on near-field pressure using plasma actuators. We also wanted to trace the static map between the forcing frequency (our control input) and the RMS of the azimuthal-mode-filtered pressure (our control output) to guide the design of the closed-loop controller. Therefore, the data for the open-loop experiments will be presented in terms of the RMS values of the pressure signals and their azimuthal-mode-filtered components.

With eight actuators in the azimuthal array, four different helical forcing azimuthal modes can be generated, viz. $m_F = 0, 1, 2,$ and 3 ; all of these were explored in our experiments. The forcing Strouhal number St_{DF} was varied from

0.09 to 5.44 ($f_F = 1$ to 60 kHz) in varying step sizes. The azimuthal-mode-filtered pressure signals and their associated SPL were computed in post-processing. Figure 4 shows the SPL spectra for two representative open-loop forcing experiments ($m_F = 0$, $St_{DF} = 0.44$ and $m_F = 3$, $St_{DF} = 2.72$) compared to the baseline case. It will be shown that the former corresponds to a maximum in the near-field pressure fluctuations, whereas the latter corresponds to a minimum. The first noticeable aspect in Fig. 4 is the presence of the forcing tone and its harmonics. The more important aspect is the amplification at the lower St_{DF} and the attenuation at the higher one. Similar effects and trends have been observed in the far-field noise also². A peculiar characteristic is the appearance of strong side bands to the forcing tone in the $St_{DF} = 2.72$ case. Although these bands are absent from the primary forcing tone in the $St_{DF} = 0.44$ case, they do appear, albeit less prominently, around the higher harmonics. At this stage, we do not have an explanation for this behavior.

During feedback control, the signals needed to be band-pass-filtered to remove both the DC offset and the high frequency noise. Based on the well-defined broadband peaks in the spectra in Fig. 4, we designed a Chebyshev filter with a pass-band from $St_D = 0.04$ to 0.73 (0.5 to 8 kHz). For accurate comparison, this filter was also applied to the open-loop pressure signals before computing their respective RMS values.

Figure 5 presents the RMS of the near-field pressure and its various azimuthal-mode-filtered versions for a wide range of forcing Strouhal numbers. The $m_F = 0$ and 3 cases are more pertinent for our application; so these forcing cases were explored on a finer grid of St_{DF} 's. The peculiar two-peak nature of the axisymmetric forcing case thus might be replicated in other modes too. We make the following observations regarding the results. (i) For all forcing




Figure 4. Representative near-field pressure spectra with open-loop forcing.

azimuthal modes explored, the pressure intensity as well as the intensity of each pressure azimuthal mode have similar characteristics – there is a relatively sharp maximum at a low St_{DF} and a broad minimum at a higher St_{DF} . The maximum clearly corresponds to the jet column mode instability frequency; previous flow visualization and PIV results have shown that the jet exhibits the largest response to forcing around this St_{DF} ¹. The first (slightly lower) peak in RMS (p) for $m_F = 0$ also corresponds to the jet column mode; at this time we do not have an explanation for the second peak. The attenuation at higher St_{DF} 's is thought to be associated with the initial shear layer instability frequency which is predicted to be between $St_{DF} = 3.6$ and 4.5 for our jet. (ii) Figure 5(a) shows that the peak in RMS (p) is largest for $m_F = 0$ and smallest for $m_F = 2$ & 3. (iii) For a particular m_F , the maximum intensity is noted in the corresponding pressure azimuthal mode. For example, Fig. 5(c) shows that the maximum value of RMS ($p^{(1)}$) is obtained by forcing the jet at $m_F = 1$. (iv) The RMS of ($p^{(0)} + p^{(1)}$) shows similar characteristics as the RMS of the pressure. (v) The switchover from amplification to attenuation (the crossing of the curves with the baseline pressure intensity line) follows a set pattern for the pressure as well as all its azimuthal modes. In particular, the curves for $m_F = 0$ attenuate at the lowest St_{DF} and those for $m_F = 3$ attenuate at the highest St_{DF} . (vi) The different pressure azimuthal modes also exhibit a progression of crossover locations. For example, all the curves for RMS ($p^{(0)}$) crossover at $St_{DF} \sim 0.9$ (but with differences therein, as mentioned above). However, this crossover happens at $St_{DF} \sim 1.8$ for RMS ($p^{(3)}$). (vii) The effects of different forcing azimuthal modes become indistinguishable at higher forcing Strouhal numbers. (viii) The location of the minimum shifts to higher St_{DF} 's for increasing pressure azimuthal modes. While the minima for RMS ($p^{(0)}$) and RMS ($p^{(1)}$) occur at $St_{DF} \sim 2.7$, the other modes have their minima at $St_{DF} \sim 3.6$. (ix) The minima become flatter for higher pressure azimuthal modes, especially at the high- St_{DF} end.

Since we are relying on the axisymmetric mode of the near-field pressure signal to estimate the far-field noise, it is instructive to compare this result with the effect of forcing on the far-field noise reported by Samimy *et al.*² and mentioned previously in Section I-D. In particular, they found the minimum far-field OASPL at $St_{DF} \sim 2.0$ at the 30° location and at $St_{DF} \sim 3.5$ at the 90° location. On the other hand, Fig. 5(b) shows that in the near-field, $p^{(0)}$ is minimized at $St_{DF} \sim 2.7$. This aspect will not be addressed in this paper; here we would minimize RMS ($p^{(0)}$) without any reference to the actual far-field noise level.

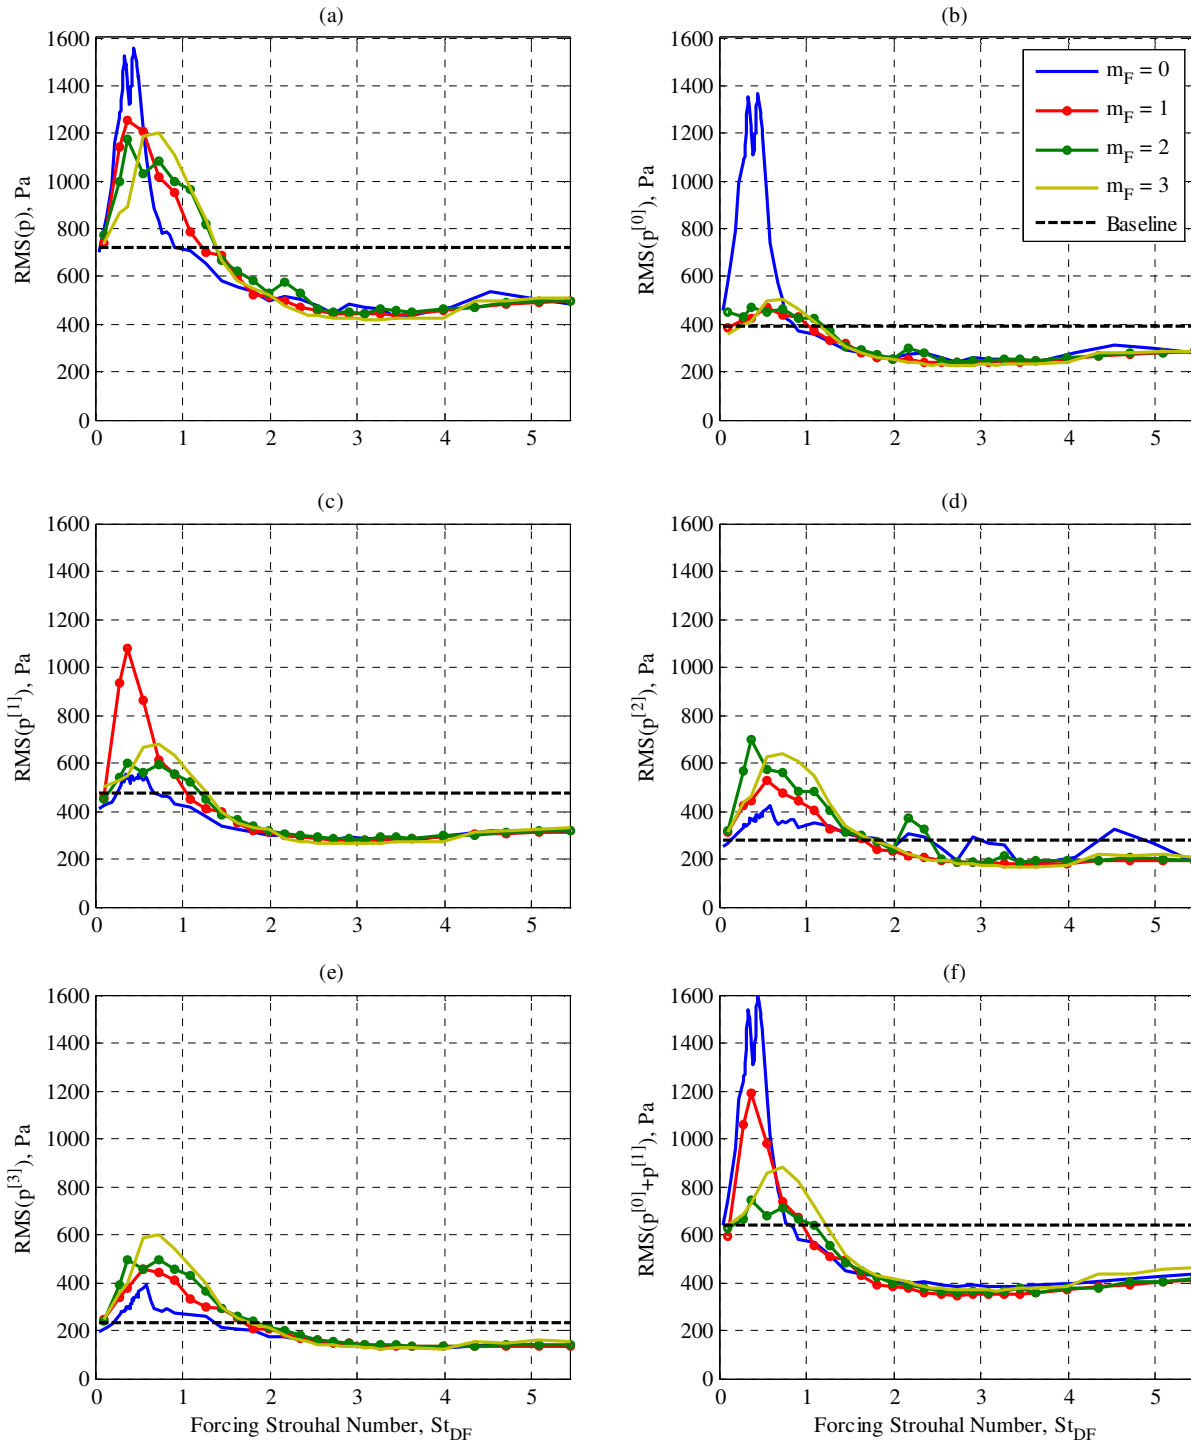


Figure 5. Effect of open-loop forcing on near-field pressure. RMS of (a) pressure, (b) $p^{[0]}$, (c) $p^{[1]}$, (d) $p^{[2]}$, (e) $p^{[3]}$, and (f) $p^{[0]} + p^{[1]}$, vs. St_{DF} . Each graph shows the effect of the four different forcing azimuthal modes explored. Also, the RMS of the respective signals in the baseline case are shown for ease of comparison. Note that the $m_F = 0$ and 3 curves do not have markers since they were collected on a fine resolution.

V. Extremizing Feedback Control Algorithms

In this section, we discuss the two feedback controllers that were implemented for online optimization in our application. We remind the reader that in our SISO problem the cost function depends only on the forcing frequency f_F of the LAFPA's. Also, the cost function is the RMS of a particular azimuthal mode of pressure fluctuation in the near-field of a highly turbulent high Reynolds number jet. Preliminary studies have shown that in spite of the averaging character of RMS, when performed over meaningful window sizes the resulting RMS value can still be noisy. Thus the problem is inherently stochastic. A further issue to be considered is the possible variation of the optimal value of f_F with time; indeed, this is the motivation for implementing *closed-loop* control. Although rapid changes cannot be expected to be handled by our controllers, they should be able to track gradual variations in the location of the optimum due to possible changes in the operating conditions. We must point out that we did not test this portion of our implementation for several practical reasons.

D. Modified Nelder-Mead Algorithm

The first controller is based on the Nelder-Mead algorithm¹⁵ (NMA) with some ideas from the Multi-Directional Search Method of Torczon¹⁷. Both these algorithms are simplex-based direct search techniques for off-line optimization of static deterministic multi-variable cost functions. On the other hand, our application involves the online optimization of a stochastic single-variable cost function; moreover the optimum of the function may vary with time. Several modifications to the original NMA are proposed below to address these differences. The modified Nelder-Mead algorithm (mNMA) will be discussed with reference to function maximization. The flowchart and pseudo-code are presented in the Appendix.

We denote the SISO cost function to be maximized by $h: \mathbb{R} \rightarrow \mathbb{R}$. For this problem, the simplex is a 2-element set of vertices representing the end points of a line; let this set be $S = \{S_1, S_2\}$. The initial choice of these vertices must be specified by the user, or some other algorithm. In the optimization process, the algorithm adapts the size and location of the simplex; however the user specifies the maximum (*MaxS*) and minimum (*MinS*) allowable size of the simplex. The significance of these constraints will be explained subsequently.

The process starts with the evaluation of $h(\bullet)$ at the vertices of the initial simplex before entering the iterator. The loop starts by ranking the vertices based on their functional values. In a maximization problem, the best vertex B would be the one with the highest function value. Since in a SISO problem there are only 2 vertices, the other one would be the worst vertex W . Hence one has

$$B \leftarrow \arg \max_i \{h(S_i)\}, \quad W \leftarrow \arg \min_i \{h(S_i)\}. \quad (2)$$

The use of the mere rank distinguishes the “direct-search” techniques from the gradient-descent methods which would typically use a finite difference approximation of the local gradient of the cost function.

The next step is *Reflection*. It is based on the hypothesis that a more optimal value of the independent variable can be found on the opposite side of B away from W . The reflected vertex R is calculated as

$$R \leftarrow B + \alpha(B - W), \quad (3)$$

where, a typical value of the reflection coefficient α is 1. The expression $h(R)$ is evaluated at this point.

On *Reflection*, if it is found that $h(R) > h(B)$, then it is logical to hypothesize that the optimum is further along the directed line segment \overline{BR} . This would then warrant *Expanding* the simplex, but only if it is smaller than *MaxS*. Also, the *Expansion* step might itself cause the simplex size to exceed its bound, in which case it must be forcibly shortened. Thus, the *Expanded* vertex E is computed as

$$E \leftarrow B + \text{sign}(R - B) \min\{\gamma |R - B|, \text{MaxS}\}, \quad (4)$$

where γ is the expansion coefficient. Typically, $\gamma = 2$. At this point, $h(E)$ is evaluated. If $h(E) > h(R)$, then the *Expansion* step is deemed successful and the new simplex is set to $\{B, E\}$ before looping back to the *Reflection* step. (The return path goes through a *Resampling* step that will be discussed later.) If on the other hand $h(E) \leq h(R)$, then the original NMA would have set the new simplex as $\{B, R\}$ before the ensuing *Reflection*. There is an issue with this scenario that is peculiar to a SISO problem. With the typical values of α and γ as above, the next vertex to be evaluated on *Reflection* of B through R would be E again. Since functional evaluations are typically expensive, one should avoid this successive evaluation of E . A better option is as follows.

Note that this branch of the algorithm is entered only if $h(B) \leq h(R)$ and $h(E) \leq h(R)$; thus the optimum can be hypothesized to lie within B and E . Also, it is logical to *Contract* the simplex to concentrate the search effort in the most promising sub-interval. An additional complication is introduced by the lower bound *MinS* on the simplex size. Let us consider the logical step in case the simplex cannot be *Contracted* any further, i.e. if $|B - R| = \text{MinS}$. (Note that $|B - R|$ cannot be less than *MinS* as $\alpha = 1$ and $|B - W| \geq \text{MinS}$ would have been enforced in the previous iteration.) R being the best vertex should be retained in the simplex. E has been newly evaluated, and hence should

not be re-evaluated. So, the only option is to re-evaluate B . This can be achieved by setting the simplex to $\{R, E\}$ and looping back to a *Reflection*. If on the other hand $|B - R| > MinS$, then we borrow from Nelder and Mead's idea of the *Inside/Outside Contraction*¹⁵. In particular, if $h(E) > h(B)$, then it is logical to *Contract* the $\{R, E\}$ simplex toward R . For this, C_1 is set to R and C_2 to E , before proceeding to the *Contraction* step; the latter is explained below. On the other hand, if $h(E) \leq h(B)$, then the $\{R, B\}$ simplex is *Contracted* toward R by setting C_1 to R and C_2 to B .

Until now, we have discussed the case wherein $h(R) > h(B)$ and the simplex size has not reached its maximum. In case the first condition is satisfied but not the second, then the simplex cannot be *Expanded* further. However, it is hypothesized that the optimum lies on the other side of R away from B (this is what motivated the *Expansion*). In this case, the correct option is to set the new simplex as $\{B, R\}$. Note that in the ensuing *Reflection*, B would be *Reflected* through R to give the E vertex that a normal *Expansion* step would have explored. This completes the discussion for the branch of the algorithm wherein the *Reflection* creates a new best vertex.

In case $h(R) \leq h(B)$, the optimum can be hypothesized to lie within W and R ; this would warrant a *Contraction* per the above discussion. Following the NMA, an *Outside Contraction* is performed if $h(R) > h(W)$ by setting $C_1 = B$ and $C_2 = R$. Otherwise an *Inside Contraction* is effected by setting $C_1 = B$ and $C_2 = W$. There is a (by now familiar) problem with the *Outside Contraction*: if the simplex size is the smallest that it can be, then the *Contraction* would not change the worst vertex R , and it would be evaluated twice in succession. A more preferable option in this case would be to set $\{B, R\}$ as the new simplex before looping back to a *Reflection*. Note that the constraint of the lower bound on the simplex size does not affect the *Inside Contraction*. In that case, if the simplex is at its minimum size, then the *Contraction* would not change the worst vertex W , and it would be evaluated again, but not in twice in succession. This would be the correct option under the circumstances.

We have already described the four routes by which one can arrive at a *Contraction* step. By the end of each such path, we have indicated the setting of C_1 and C_2 . C_1 represents the better vertex that should be retained, while C_2 is the worse vertex which should be replaced by its *Contraction*. The actual *Contraction* step is identical to the one in the NMA except for the enforcement of the minimum size constraint on the simplex, $MinS$. That is, the *Contracted* vertex is calculated as

$$C \leftarrow C_1 + \text{sign}(C_2 - C_1) \max\{\beta |C_2 - C_1|, MinS\}, \quad (5)$$

where a typical value of the contraction coefficient β is 0.5. Now $h(C)$ is evaluated. The *Contracted* simplex $\{C_1, C\}$ is accepted unconditionally and the algorithm proceeds to the next *Reflection* step. The original NMA also includes a *Shrink* step. However, for a SISO problem, it duplicates the *Contraction* step, albeit with a possibly different coefficient; it is therefore not included in the proposed algorithm.

We now move on to a discussion of the *Resampling* and *Age-test* steps that are incorporated to address the noise in our application; the ideas are adopted from Barton & Ivey¹⁶ and Zhang²², respectively. In essence, although the vertex B might not be the true optimum, a single observation of its functional value in a stochastic setting might make it *seem* optimal. So, at the end of an iteration (before the next *Reflection* step) the algorithm notes the age of B , i.e. the number of times it has been retained consecutively as the best. If the user chooses to *Resample* and if the current age of B is greater than a user-specified *AgeLimit*, then B is re-evaluated and its old functional value is replaced with the new one in subsequent rank determinations.

At this point, we explain the incorporation of the constraints on the simplex size; such constraints are absent in the NMA. The mNMA is designed to track gradual variations in the location of the optimum due to possible changes in the operating conditions. A collapsed simplex will not be sensitive to changes in the local gradient of the cost function; the lower bound $MinS$ is applied to keep the process "alive" to such variations. This also corresponds to the notion of persistent excitation used in the adaptive control literature since it ensures a minimum amplitude of oscillation of the plant input. Note that with this modification, we cannot have a stopping criterion unlike the original NMA. In fact, the possibility of variations in the optimum argues against stopping the algorithm except through a manual override.

The upper bound $MaxS$ is implemented as a safety precaution. In a stochastic setting, a few contiguous erroneous observations can result in a very large simplex due to repeated *Expansions*. First of all, this might result in actuator saturation. Also, the process may require a large number of iterations to recover from such a large excursion. This argues for the imposition of the upper bound.

In summary, three modifications are proposed for the classical NMA to adapt it to our stochastic situation with a possibly time-varying optimum: (i) *Resampling* with *Age-test*, (ii) application of a lower bound on the simplex, and (iii) imposition of an upper bound on the simplex. Although at first glance, the flowchart of the mNMA looks very convoluted (compared to the simple NMA and the even simpler Multi-Directional Search Method), hopefully the above discussion has convinced the reader about the purpose of each of its branches.

E. Modified Kiefer-Wolfowitz Algorithm

The second controller is based on the Kiefer-Wolfowitz algorithm¹⁹ (KWA) for stochastic optimization. We briefly explain the original KWA before presenting the proposed modifications to adapt it to our application.

Consider the stochastic cost function $h(y_n)$, where y_n is the independent variable at the n^{th} iteration of the procedure. It is assumed that the noise is unbiased and that $E\{h(\bullet)\}$ is a convex unimodal function, where $E\{\bullet\}$ is the expectation operator. The goal is to iteratively find the y that maximizes $E\{h(y)\}$ by only evaluating $h(\bullet)$. (For minimization, the negative of $h(\bullet)$ can be passed to the routine.) In the KWA, the derivative $E'\{h(y_n)\}$ is estimated as a finite difference using a positive and negative perturbation. The center y_{n+1} of the next pair of measurements is then computed as

$$y_{n+1} = y_n + a_n \frac{h(y_n+c_n)-h(y_n-c_n)}{c_n}, \quad (6)$$

where $a_n > 0$ controls the step size and $c_n > 0$ is the perturbation level at the n^{th} iteration. The convergence proof of the technique¹⁹ makes the following assumptions on the sequences c_n and a_n :

$$\lim_{n \rightarrow \infty} a_n = 0, \quad \lim_{n \rightarrow \infty} c_n = 0, \quad \sum_{n=1}^{\infty} a_n = \infty, \quad \sum_{n=1}^{\infty} \left(\frac{a_n}{c_n}\right)^2 < \infty. \quad (7)$$

One common choice for a_n is a_1/n , with $a_1 > 0$. Then, a possible choice for c_n is c_1/n^λ , with $0 < \lambda < 0.5$ and $c_1 > 0$.

We propose to modify the KWA by applying lower bounds a_{min} and c_{min} on a_n and c_n , respectively. This would provide a persistent excitation to the process to account for a possible gradual time-variation of the optimal value of the variable y . Note that this would violate the conditions placed on the sequences a_n and c_n . Hence, we cannot expect a convergence to the local optimum; instead, the independent variable can be expected to oscillate in a neighborhood of this optimum.

The KWA is essentially a gradient-descent algorithm, albeit for a noisy function. This is not suitable in situations where the cost function, although convex, has a very small slope in the vicinity of the maximum. Also, this will delay convergence in case the function has inflection points. The solution described in Wilde²⁰ is termed *normalization*. Essentially, instead of the actual value of the slope, only its sign is used in the optimization. Thus, at the end of each iteration, the new center for the next pair of measurements is computed as

$$y_{n+1} = y_n + a_n \operatorname{sign} \left\{ \frac{h(y_n+c_n)-h(y_n-c_n)}{2c_n} \right\} = y_n + a_n \operatorname{sign} \{h(y_n+c_n)-h(y_n-c_n)\}. \quad (8)$$

The last equality utilizes the fact that c_n is positive. This *normalization* transforms the KWA from a gradient-based process to one that uses the rank information of the sampled points only, similar to direct search.

The other important modification is the *acceleration* protocol proposed by Kesten²⁵; he argued that a change in the sign of the slope would indicate that the optimum was crossed over from one side to the other between the last two iterations. Thus such a sign change should be followed by a reduction of the step size, a , for convergence. However, far away from the optimum (signified by a constant sign of the slope), reducing a would only serve to slow down convergence, and hence must be avoided. In the original KWA, a_n and c_n are modified in lock-step with the variable y_n ; the *acceleration* method decouples these two processes. So, although y_k is still updated for every step index k , the index n governing the decrease of a_n is incremented only on the detection of a change in the sign of the estimated slope. Additionally, Kesten stipulated that the perturbation level c_n should be a constant.

We adopted Wilde's *normalization* technique in full, and Kesten's *acceleration* protocol with one modification. In particular, we did not adopt Kesten's stipulation of keeping the perturbation level constant. In simulations preceding the experiments, we obtained the best results by allowing c_n to decrease from an initial higher value down to the lower bound c_{min} imposed on it by the first modification proposed above. The reasoning is as follows. The c_{min} cannot be too large, otherwise the variable y as well as the cost function $h(\bullet)$ will make large excursions from their respective optima, which is undesirable. However, if following Kesten, this relatively small value of perturbation is used from the beginning of the process, then the process is found to be overly sensitive to noise thereby delaying convergence. Thus, in our algorithm, we set the initial perturbation level c_1 suitably high and used the c_1/n^λ formula for decrementing it in lock-step with the step-size decrements. Of course, since this adaptation of c only occurs during the initial stages, we are not constrained anymore by $0 < \lambda < 0.5$ and are free to vary this parameter to improve performance. Note that this does not really violate Kesten's stipulation since the perturbation level becomes constant at c_{min} after a short initial period. All these ideas are incorporated in the mKWA; its pseudo-code is presented in the Appendix.

VI. Results of Extremizing Feedback Control

We performed closed-loop experiments using two different algorithms: the modified Nelder-Mead algorithm (mNMA), and the modified Kiefer-Wolfowitz algorithm (mKWA). Here we present the results and discuss their relative merits. Since the controller implementation was in terms of the forcing frequency instead of the Strouhal number, all closed-loop results will be discussed in frequency terms. For our jet, $f_F = 1$ kHz implies $St_{DF} = 0.091$.

Recall from earlier discussions that the axisymmetric mode of the near-field pressure is best correlated with the far-field noise, and that the suppression of the fluctuations of this pressure mode might mitigate the noise. Thus, we would strive to seek the minimum of RMS ($p^{(0)}$) in real-time by extremizing the forcing frequency for potential noise attenuation. It was shown in Figure 5(b) that at the high forcing frequency range where the minimum is achieved, the effect of different forcing azimuthal modes cannot be distinguished. However, as pointed out in Section I-D, the higher helical modes of forcing are more effective for far-field noise suppression. Since $m_F = 3$ is the highest helical mode achievable with our eight actuators, this was used exclusively in closed-loop experiments aimed at minimization. The details of this particular I/O map obtained in open-loop forcing experiments are shown in Fig. 6(f). We would show convergence to the optimal forcing frequency ($f_F \sim 30$ kHz for RMS ~ 230 Pa) starting from various forcing frequencies in the range of 10 to 50 kHz.

We now turn to the goal of mixing enhancement by the maximum amplification of the large-scale structures in the jet shear layer. We have discussed before that this is reflected in a broadband increase of the near-field pressure fluctuations. The results from open-loop forcing experiments (see Fig. 5(a)) show that $m_F = 0$ generates the largest amplification of the near-field pressure fluctuations. Hence, all maximum-seeking experiments were performed in this forcing mode. Also, it has been discussed earlier that the near-field pressure is almost fully reconstructed from $p^{(0)} + p^{(1)}$. Hence, the cost function is selected as RMS ($p^{(0)} + p^{(1)}$). A detailed version of the pertinent cases of open-loop forcing is presented in Fig. 7(f). We would try to converge to the optimal forcing frequency ($f_F \sim 4.9$ kHz for RMS ~ 1600 Pa) from various initial forcing frequencies in the range of 1 to 10 kHz; it will be shown that this is not always possible owing to the presence of the other prominent peak at ~ 3.7 kHz. However, since the ordinates of the two peaks are quite similar, achieving either of them would result in significant mixing enhancement.

An important assumption in the discussion of the control algorithms in Section V is the absence of plant dynamics. To verify this, step response experiments were conducted. A delay of no more than 20 ms was noticed between a change in the forcing frequency and its effect on the near-field pressure signal. The possible sources of this delay are (i) the communication lag between the dSpace controller commanding the forcing frequency and the Labview program acquiring it to control the actuators, and (ii) delays internal to the feedback controller and actuator controller owing to the finite processor speeds. To account for this delay in the implementation, the control routine incorporated a wait of 20 ms between directing the actuators to use a particular forcing frequency and starting to record the corresponding pressure data. Once the necessary sample size was collected, its RMS was computed, and this was deemed as the plant output to be used by the controller in deciding the next input. This cycle was then repeated. Such a strategy makes the plant appear effectively static to the controller.

An implicit parameter that sets the pace of our closed-loop experiments is the size of the sampling window for RMS computation. Instead of referring to the window size, we will discuss this in terms of T_{rms} , the duration over which a window of pressure samples is collected. Note that the waiting period of 20 ms described above is included within T_{rms} . Turbulent randomness creates variability in the computed RMS value depending on T_{rms} ; this is seen in the error bars in Figs. 6(f) and 7(f). For fast response, T_{rms} should be as small as possible. However, the variability (or noise) in the output signal also increases with decreasing T_{rms} . Thus a balance is required between an acceptable convergence rate and a signal-to-noise ratio (SNR) that can be handled by the inherent robustness of the implemented algorithms. We define the SNR as the ratio of the mean value of the plant output to its standard deviation when all parameters are held constant. In preliminary experiments of maximum-seeking with either algorithm, $T_{rms} = 0.1$ s proved to be a satisfactory balance. However, this value proved too small for minimum-seeking and we had to resort to $T_{rms} = 0.2$ s to have similar robust performance. In analyzing this finding, we tried to estimate the SNR from the open-loop experiments. The SNR is obviously a function of T_{rms} as well as the forcing frequency. Using $T_{rms} = 0.2$ s, we estimated the point-wise standard deviation for the map to be minimized; this is shown in Fig. 6(f). On the other hand, in Fig. 7(f), the standard deviation is calculated using $T_{rms} = 0.1$ s for the map to be maximized. The average SNRs (calculated over the points indicated in the respective figures) were approximately 64 near the respective optima for both maximum-seeking and minimum-seeking. However, if we used the same $T_{rms} = 0.1$ s for both applications, then the SNR in the optimal range of forcing frequencies for minimum-seeking turned out to be 45. This explains the need for different T_{rms} for the two applications.

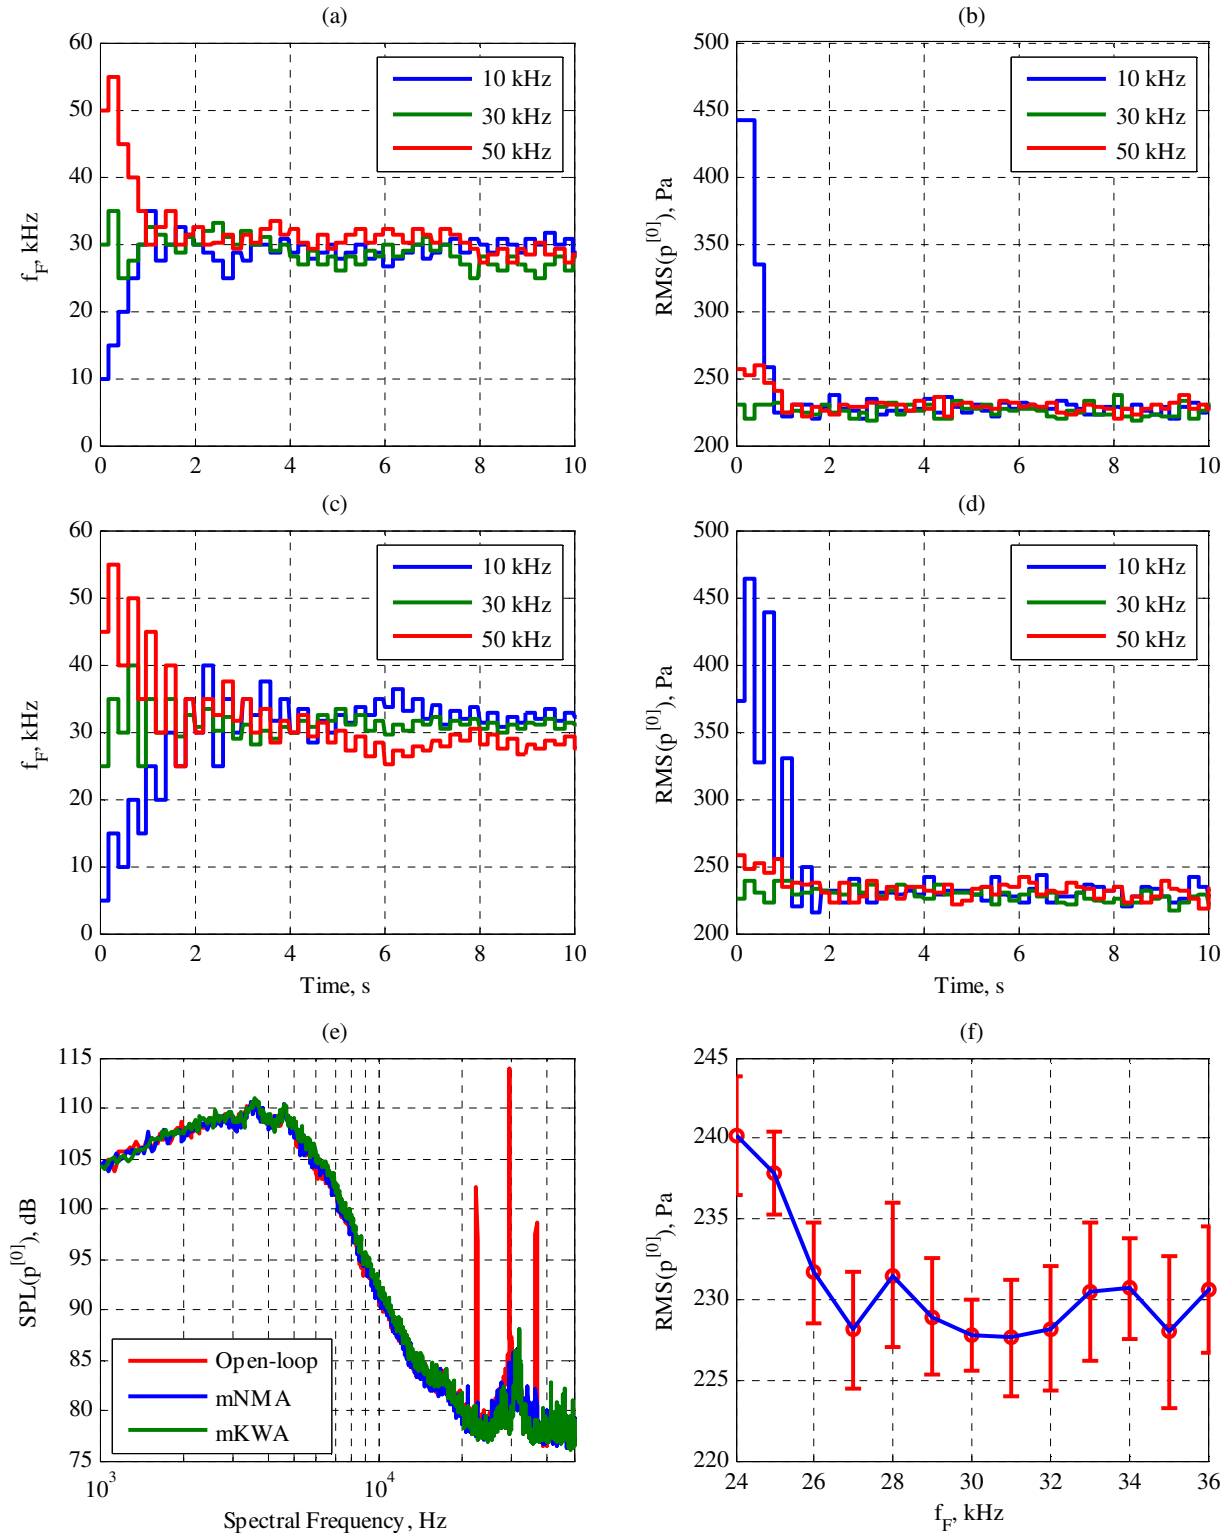


Figure 6. Transients in minimum-seeking control with (a) & (b) mNMA and (c) & (d) mKWA. (e) Steady state spectra in closed-loop with the open-loop spectrum for the optimal case ($f_F = 30$ kHz and $m_F = 3$). (f) Details of pertinent open-loop results with $m_F = 3$. Standard deviations σ were obtained with $T_{rms} = 0.2$ s; the error bars are $\pm \sigma$.

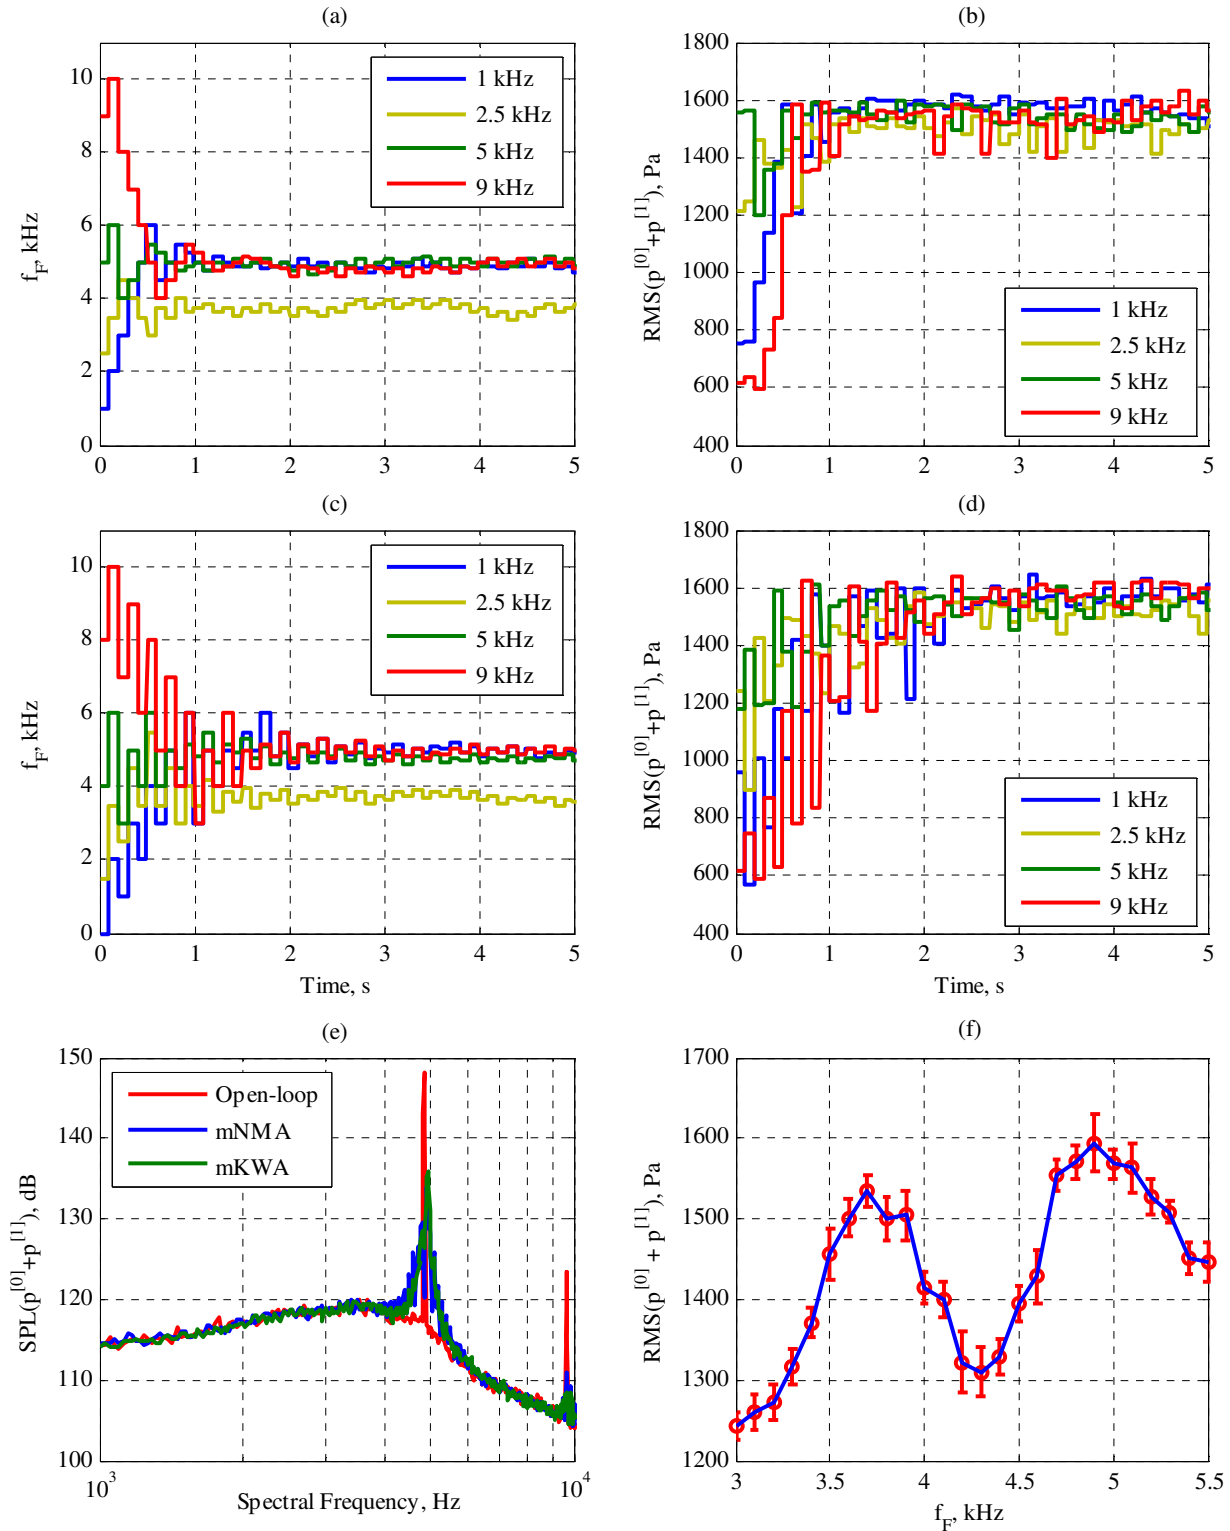


Figure 7. Transients in maximum-seeking control with (a) & (b) mNMA and (c) & (d) mKWA. (e) Steady state spectra in closed-loop with the open-loop spectrum for the optimal case ($f_F = 4.9$ kHz and $m_F = 0$). (f) Details of pertinent open-loop results with $m_F = 0$. Standard deviations σ were obtained with $T_{rms} = 0.1$ s; the error bars are $\pm \sigma$.

The control algorithms have a number of parameters that determine their performance. The initial choice of the parameters was based on an analysis of the open-loop results as well as preliminary simulations. Experimental results were used to guide the selection of the final set of parameters that produced the best performance. Eventually the algorithms were run with these parameter values in at least five cases each to verify their repeatability. We only present representative results from this experimental set for space constraints. For ease of comparison, the transient behaviors of all the minimum-seeking experiments are shown together in Figs. 6(a)-(d). Similarly, all the transient results of maximum-seeking appear in Figs. 7(a)-(d). The steady-state output spectra for closed-loop minimum-seeking are presented in Fig. 6(e); they are compared with the best result obtained in open-loop, viz. the case of $f_F = 30$ kHz and $m_F = 3$. The steady-state output spectra for maximum-seeking applications are presented in Fig. 7(e); here the reference open-loop case had $f_F = 4.9$ kHz and $m_F = 0$.

F. Modified Nelder-Mead Algorithm

We will first discuss the choice of parameters that govern the behavior of the mNMA. We follow previous work¹⁵⁻¹⁷ and always set $\alpha = 1$, $\beta = 0.5$, and $\gamma = 2$. Also, we always choose to perform *Resampling* to correctly address the stochastic nature of the process. In a deterministic process where the goal is to converge to a unique optimum, it makes sense to set a nonzero value for *AgeLimit*. In our stochastic setting there is a persistent excitation, but the optimal forcing frequency should be evaluated as often as feasible so that the average plant output remains close to its optimum. For this, we always set *AgeLimit* to 0. In preliminary experiments with nonzero *AgeLimit*, the output was definitely steadier, but it was generally slightly less optimal than with *AgeLimit* = 0. Search techniques typically require the user to specify an initial guess of the goal. In our implementation, we use this as S_1 , the first vertex of the initial simplex. Of course, the aim of our closed-loop experiments is to show that this initial guess can be chosen in a wide range and still result in optimal plant output within a short time. The other vertex of the initial simplex is computed as $S_2 = S_1 + S_{s0}$. The remaining parameters of the algorithm, viz. S_{s0} , *MaxS*, and *MinS*, are application-dependent and hence detailed separately below.

1. Minimum-Seeking

In this application the mNMA is used to seek the minimum of the I/O map shown in Fig. 6(f); the characteristics of this map dictate the choice of the parameters of the algorithm. The value of S_{s0} must be chosen such that the algorithm can reliably rank its vertices initially; this means that the vertices must be spaced far enough apart so that their functional values have the correct rank in spite of noise. The I/O map shows a very gradual slope in the range from 30 to 60 kHz; this calls for a large S_{s0} . Of course, $S_{s0} \leq \text{MaxS}$. Since we typically desire the fastest convergence, we always specify $S_{s0} = \text{MaxS} = 5$ kHz. With $S_{s0} = \text{MaxS} = 3$ kHz, we could not achieve repeatability in converging to the optimum from an initial forcing frequency of 50 kHz. With $S_{s0} = \text{MaxS} = 10$ kHz, the convergence was repeatable and rapid, but there was a large overshoot that caused delays in settling. (With $S_{s0} = \text{MaxS}$, the algorithm would typically never encounter an *Expansion* step; so to validate this portion of the code we ran several experiments wherein we set $S_{s0} < \text{MaxS}$.) The I/O map also indicates that the minimum is broad; hence an appropriate choice of *MinS* was 1 kHz. Currently, we do not know what would constitute an adequate level of persistent excitation in real applications; we simply wish to show that a meaningful level can be feasibly established by our algorithm. Experiments with *MinS* set to 0 and 2 kHz have also shown satisfactory steady state behavior. Finally, T_{rms} was set to 0.2 s as discussed before. We could show repeatable convergence with $T_{rms} = 0.1$ s when f_F was initiated at 10 kHz, but we could not obtain similar results starting from 50 kHz. This is because the gradual slope in the latter case gets obscured by noise.

The transient behavior is demonstrated with the process initialized at three different forcing frequencies, viz. 10, 30, and 50 kHz. Figure 6(b) shows that in all three cases, the plant output reaches and settles around 230 Pa within $5T_{rms}$, i.e. using just 5 evaluations of the plant I/O map. This convergence rate represents an order of magnitude improvement compared to our results with the gradient-based extremum-seeking controller¹³. Figure 6(a) indicates that f_F reaches ~ 30 kHz within the same time but it varies within 25 to 35 kHz subsequently, even though *MinS* = 1 kHz. This can be explained by the broadness of the minimum. The three cases have similar steady state behavior, hence we only show the case with initial $f_F = 50$ kHz in Fig. 6(e). The spectrum of the output is quite similar to the optimal open-loop spectrum, especially in the high amplitude region of interest. The forcing tone in the feedback controlled case is smaller due to the persistent perturbations in the forcing frequency. Thus the best open-loop result in minimization of the near-field pressure fluctuations can be replicated in closed-loop by the mNMA with a very short transient phase.

2. Maximum-Seeking

We now focus on the online maximization of the I/O map shown in Fig. 7(f) using the mNMA. It has already been discussed that the relevant portion of this map has two distinct maxima – a 1600 Pa peak at 4.9 kHz and a 1530 Pa peak at 3.7 kHz. Direct search normally cannot guarantee the convergence to a specific peak starting from any

point in an interval containing multiple peaks. However, repeated experimentation can lead us to a choice of a parameter set that might do so. This is not the goal of model-free optimization, wherein we specifically avoid using any significant knowledge of the plant. Instead, through logical arguments and a few experiments, we did arrive at a set of parameters that gave *repeatable* convergence to the higher peak at 4.9 kHz from a *majority* of the initial conditions within 1 and 9 kHz. Since we could not get this result for all initial conditions in the range, we cannot make any claim for the basin of attraction of the individual peaks. Instead, our motive here is to explore the limitation of the mNMA for this application.

We chose $Ss_0 = MaxS$ following the reasoning laid out in our discussion regarding minimum-seeking. Owing to the smaller range of initial conditions and the sharpness of the peaks, we selected $Ss_0 = MaxS = 1$ kHz. Values of 800 and 2000 Hz were also tried with similar results, attesting to the relative robustness of the technique to the specific choice of this parameter. Further reduction of Ss_0 led to delays in convergence due to corruption by noise. The sharpness of the peak led us to choose $MinS = 100$ Hz. A value of 200 Hz did not lead to substantial degradation in performance. However, with $MinS = 500$ Hz, large oscillations were observed in the forcing frequency and the average output was reduced since the larger perturbations were keeping the output away from its optimum. As in minimum-seeking, we do not know a suitable level of persistent excitation; hence our goal here has been to demonstrate the flexibility of the algorithm. Finally, T_{rms} was chosen as 0.1 s as discussed before. In experiments with $T_{rms} = 0.05$ s, we found one case where the process failed to converge when initialized at 9 kHz.

Figure 7(a) shows that the plant input settles within $10T_{rms}$ from a variety of initial values. On initialization from 1, 5, and 9 kHz, the final mean value of the input was around 4.9 kHz; starting from 2.5 kHz, the steady state average was around 3.7 kHz. Actually, we tested the algorithm with all initial forcing frequencies in the range from 1 to 9 kHz in increments of 0.5 kHz. The process repeatedly converged to the peak at 3.7 kHz when initiated from 2.5 and 3.5 kHz; with all other initial conditions, it repeatedly reached the peak at 4.9 kHz. In Fig. 7(b) we note that the output settles around 1500 Pa within $10T_{rms}$ of initialization for all cases. In steady state, the output is seen to have significant oscillations compared to the minimum-seeking case. Also, the mean output is somewhat lower than the optimal value found in open-loop experiments (~ 1600 Pa). This can be explained by three factors: (i) from open-loop experiments, we already noted a greater variability in the output for maximum-seeking compared to that for minimum-seeking (see Figs. 6(f) and 7(f)); (ii) the shorter T_{rms} (0.1 s in maximum-seeking compared to 0.2 s in minimum-seeking) further increases the noise content in the processed output; and (iii) the sharpness of the peak dictates that minor perturbations from the optimal input create substantial degradations in the output. Figure 6(e) presents the steady state output spectrum for the experiment with initial $f_F = 1$ kHz. (The other initial forcing frequencies, apart from 2.5 and 3.5 kHz of course, produced very similar steady state behavior, and hence SPL spectra.) Within the high-amplitude range of interest, this spectrum closely resembles the optimal open-loop spectrum, except around the forcing tone and its harmonics. The disparity is obviously due to the persistent excitation of the forcing frequency in closed-loop control. In summary, the mNMA meets expectations in maximum-seeking: there are two distinct peaks (with similar ordinate values) within the set of possible initial conditions of forcing frequency, and the process successfully converges to either one of them in an adequately short time.

G. Modified Kiefer-Wolfowitz Algorithm

Before presenting the experimental results, we explain the choice of parameters governing the mKWA. Preliminary investigation showed that the *normalization* and *acceleration* modifications always produce faster convergence without any negative effect on the steady-state performance; hence, these were always performed. Next, we set $\lambda = 1.0$. No degradation of performance was observed (in fact, there was an improvement) even when using $\lambda = 5.0$. With this value of λ , the process uses c_1 as the perturbation level until the first zero-crossing of the estimated slope. Subsequently, as n becomes 2, the perturbation level would become c_{min} since we always had $c_{min} > c_1/2^5$. The improved performance was obviously due to lower oscillations of the output after the initial rise (or fall) period. Although we never noticed it in experiments, a single noise-corrupted observation might cause an erroneous zero-crossing of the estimated slope causing the perturbation level to fall to c_{min} from the beginning itself; this would defeat the purpose of implementing the gradual decrement mechanism. At the other extreme, a small value of λ would result in a large settling time for the perturbation level during which the plant output will continue to make large excursions from its optimum. As a compromise, we chose $\lambda = 1.0$.

The timing parameter T_{rms} was chosen as in the mNMA experiments; i.e. $T_{rms} = 0.2$ s for minimum-seeking and 0.1 s for maximum-seeking. In the mKWA, each decision to issue a new plant input is based on two observations of the output – positive and negative perturbations from the central value. On the other hand, in the mNMA, each such decision is based on only one observation of the output. This makes the former inherently more robust. Correspondingly, we found no issues with convergence even when the values of T_{rms} were halved. However, to maintain a decent margin of safety, we settled for the values mentioned above, in line with the mNMA experiments.

This leaves four other parameters: c_1 , a_1 , c_{min} , and a_{min} . These had to be chosen differently for the minimum-seeking and maximum-seeking applications, and will be discussed below.

3. Minimum-Seeking

Here we discuss experiments in minimization of the I/O map shown in Fig. 6(f) using mKWA. This map possesses a gradual slope at the high-frequency end. If the process noise is higher than the difference in output with the positive and negative perturbations introduced by the mKWA, the convergence (if it happens at all) would be very slow indeed. To address this, c_1 was chosen as 5 kHz. The value of a_1 was selected as 5 kHz also to hasten convergence. This is similar, but not equivalent, to selecting Ss_0 and $MaxS$ as 5 kHz in the mNMA. The parameters c_{min} and a_{min} together play a role similar to $MinS$ in the mNMA implementation. Following the reasoning for the latter, we selected $c_{min} = a_{min} = 0.5$ kHz. Here a_{min} keeps the optimization process “alive” in the long run whereas c_{min} ensures that the estimation of the sign of the local gradient is relatively insensitive to noise. In this respect, the mKWA offers more flexibility than the mNMA since in the latter both these effects are controlled by $MinS$.

Figure 6(c) demonstrates that starting from various initial conditions, the plant input reaches around 30 kHz within $10T_{rms}$, but it varies within 25 and 35 kHz subsequently. Figure 6(d) shows that the output reaches and settles around 230 Pa within the same interval in all three cases. Except for the doubled settling time, all other characteristics are very similar to the mNMA results, including the representative steady-state output spectrum shown in Fig. 6(e). The settling time is doubled because the mKWA makes two observations of the output for every observation made by the mNMA. While this makes the mKWA more robust, a price is paid in the delayed convergence. Of course, we can obtain the same real convergence time as in the mNMA by halving the T_{rms} to 0.1 s. As mentioned before, brief experimentation with this parameter in the mKWA has not revealed any adverse effect unlike that found in the mNMA, but we retained $T_{rms} = 0.2$ s to have some margin of safety.

4. Maximum-Seeking

The pertinent plant I/O map for our maximum-seeking application is shown in Fig. 7(f). The parameter c_1 must be large enough so that noise does not obscure the proper estimation of the sign of the local slope using the positive and negative excursions. However, if c_1 is too large then it will take longer to decrement to c_{min} , thereby delaying the settling of the output. As a compromise, we selected $c_1 = 1$ kHz. In a few experiments with $c_1 = 0.8$ kHz, we did not notice any erroneous convergence behavior. Also with $c_1 = 2$ kHz, the oscillations indeed took longer to settle down. The parameter a_1 must be large for fast convergence, but an upper bound is imposed by the same desire for quick settling; we selected $a_1 = 1$ kHz. As in minimum-seeking, we mention that this choice of c_1 and a_1 is similar, but not equivalent, to selecting $MaxS = 1$ kHz in maximum-seeking with the mNMA. The parameters c_{min} and a_{min} were both selected as 50 Hz, exploiting their relationship to $MinS$ in the mNMA; the value of the latter has been justified already. Lastly, we chose $T_{rms} = 0.1$ s as explained before.

Figure 7(c) shows that plant input settles within $22T_{rms}$ from a variety of initial values. Starting from 2.5 kHz, the steady state value of f_F was around 3.7 kHz. However, the input settled around 4.9 kHz when initialized from all other values in the range from 1 to 9 kHz in increments of 0.5 kHz (three of these cases are shown in the figure). The reader will recall that similar results were obtained with the mNMA too, except that in the latter the initial f_F of 3.5 kHz also resulted in convergence to the peak at 3.7 kHz. Too much significance should not be read into this disparity, given that neither algorithm is designed to optimize multi-modal functions. Figures 7(d) & (e) demonstrate that the time-domain as well as frequency-domain behavior of the steady-state plant output with the mKWA are very similar to those observed with the mNMA, except for the doubled convergence time in the former case. The explanation given for this observation in the case of minimum-seeking also applies to the current application.

VII. Conclusion

In this paper we have presented results of development and application of extremizing feedback control to axisymmetric jets. In particular, we have shown control authority on the near-field pressure of a Mach 0.9 jet with a Reynolds number based on nozzle exit diameter of 7.6×10^5 . Open-loop forcing using localized arc filament plasma actuators is shown to have two distinct effects on the near-field: a large amplification occurs in the pressure fluctuations at lower forcing frequencies near the jet column mode, whereas a broad attenuation is found at higher forcing frequencies. The peak corresponds to maximum bulk mixing in the jet shear layer; the minimum is important for its potential link to far-field noise attenuation. In this paper we focused exclusively on the near-field; thus this link was not investigated.

Previous researchers have shown that the axisymmetric mode of the near-field pressure is best correlated with the far-field noise. Thus one application of feedback control is in minimizing this signal for potential noise mitigation by tuning the forcing frequency. On the other hand, it is also known that the instantaneous near-field pressure signal is well-reconstructed from the sum of its axisymmetric and first helical modal components only.

Therefore, the second application of feedback control is in maximizing the latter sum to obtain mixing enhancement. Both applications can thus be regarded as online optimization problems involving single-variable cost functions. An important characteristic of both applications is the high level of noise introduced in the corresponding cost functions by turbulence.

Two different gradient-free offline optimization algorithms were investigated – the Nelder-Mead direct search method and the Kiefer-Wolfowitz stochastic optimization routine. The former is designed for efficient optimization of deterministic cost functions whereas the latter can also handle noisy functions. Several modifications were incorporated in the original algorithms before they could be implemented in feedback controllers for online optimization. In closed-loop control experiments for minimum-seeking and maximum-seeking, both controllers were able to replicate the optimal open-loop control results after very brief transient phases. The modified Kiefer-Wolfowitz method was found to be more robust but the convergence speed was typically halved in comparison to the modified Nelder-Mead controller. One shortcoming of the former is the absence of any mechanism to increase the perturbation and increment levels; this might be a disadvantage if the desire is to quickly account for changes in the location of the optimum during operation. The *Expansion* mechanism in the modified Nelder-Mead algorithm would be very useful in this regard.

In conclusion, both feedback controllers investigated in this article demonstrate significant promise for the online optimization of the near-field pressure of axisymmetric jets.

Appendix

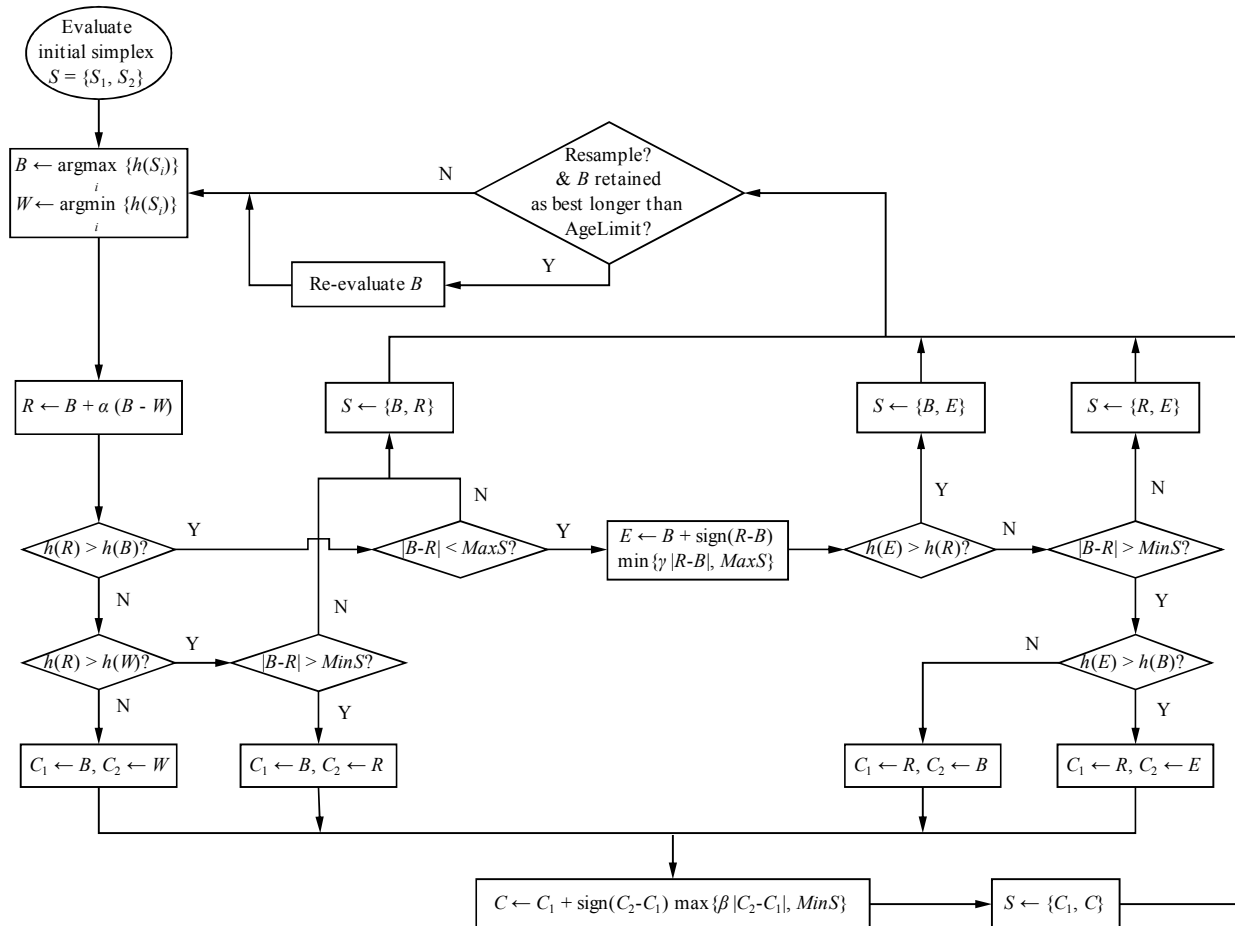


Figure 8. Flowchart for the modified Nelder-Mead algorithm for function maximization. To minimize the cost function $h(\bullet)$, its negative is maximized.

Pseudo-code for the Modified Nelder-Mead Algorithm for Function Maximazation

Given an initial simplex, S , with vertices $\{S_1, S_2\}$. Also, given $MaxS$, $MinS$, and $AgeLimit$.

Calculate $h(S_1)$ and $h(S_2)$. $Age \leftarrow 1$.

while (not manually overridden)

$ContractFlag \leftarrow FALSE$, $BFlag \leftarrow FALSE$, $B \leftarrow \arg \max_i \{h(S_i)\}$, $W \leftarrow \arg \min_i \{h(S_i)\}$.

$R \leftarrow B + \alpha (B - W)$. Calculate $h(R)$. /* Reflection

if ($h(R) > h(B)$) **then**

if ($|B - R| < MaxS$) **then**

$E \leftarrow B + \text{sign}(R - B) \min\{\gamma |R - B|, MaxS\}$. Calculate $h(E)$. /* Expansion

if ($h(E) > h(R)$) **then**

$S \leftarrow \{B, E\}$

else

if ($|B - R| > MinS$) **then**

if ($h(E) > h(B)$) **then**

$C_1 \leftarrow R$, $C_2 \leftarrow E$, $ContractFlag \leftarrow TRUE$

else

$C_1 \leftarrow R$, $C_2 \leftarrow B$, $ContractFlag \leftarrow TRUE$

endif

else

$S \leftarrow \{R, E\}$

endif

endif

else

$S \leftarrow \{R, B\}$

endif

else if ($h(R) > h(W)$) **then**

if ($|B - R| > MinS$) **then**

$C_1 \leftarrow B$, $C_2 \leftarrow R$, $ContractFlag \leftarrow TRUE$, $BFlag \leftarrow TRUE$

else

$S \leftarrow \{B, R\}$, $BFlag \leftarrow TRUE$

endif

else

$C_1 \leftarrow B$, $C_2 \leftarrow W$, $ContractFlag \leftarrow TRUE$, $BFlag \leftarrow TRUE$

endif

if ($ContractFlag == TRUE$) **then**

$C \leftarrow C_1 + \text{sign}(C_2 - C_1) \max\{\beta |C_2 - C_1|, MinS\}$. Calculate $h(C)$. /* Contraction

$S \leftarrow \{C_1, C\}$

endif

if ($BFlag == TRUE$) **then**

$Age \leftarrow Age + 1$

else

$Age \leftarrow 1$

end

if ($Resampling$ is desired AND $BFlag == TRUE$ AND $Age > AgeLimit$) **then**

Recalculate $h(B)$, $Age \leftarrow 1$.

endif

endwhile

Pseudo-code for the Modified Kiefer-Wolfowitz Algorithm for Function Maximazation

Given an initial value y_i . Also given the parameters c_1 , a_1 , c_{min} , a_{min} , and λ .

$k \leftarrow 1, n \leftarrow 1$.

while (not manually overridden)

$sgn_k \leftarrow \text{sign} \{h(y_k + c_n) - h(y_k - c_n)\}$

if (Acceleration is desired) **then**

if ($k > 1$ AND $sgn_k \neq sgn_{k-1}$) **then**

$n \leftarrow n + 1, c_n \leftarrow \max\{c_1/n^\lambda, c_{min}\}, a_n \leftarrow \max\{a_1/n, a_{min}\}$.

endif

else

$n \leftarrow n + 1, c_n \leftarrow \max\{c_1/n^\lambda, c_{min}\}, a_n \leftarrow \max\{a_1/n, a_{min}\}$.

endif

if (Normalization is desired) **then**

$y_{k+1} \leftarrow y_k + a_n sgn_k$

else

$y_{k+1} \leftarrow y_k + a_n \{h(y_k + c_n) - h(y_k - c_n)\}/2c_n$

endif

$k \leftarrow k + 1$

endwhile

Acknowledgments

The support of this research by the Air Force Office of Scientific Research with Drs. Rhett Jefferies and John Schmisser is greatly appreciated. The authors would like to thank Martin Kearney-Fischer for his help on this project.

References

- ¹Samimy, M., Kim, J.-H., Kastner, J., Adamovich, I., and Utkin, Y., Active control of high-speed and high-Reynolds-number jets using plasma actuators, *Journal of Fluid Mechanics*, Vol. 578, 2007, pp. 305–330.
- ²Samimy, M., Kim, J.-H., Kastner, J., Adamovich, I., and Utkin, Y., Active control of a Mach 0.9 jet for noise mitigation using plasma actuators, *AIAA Journal*, Vol. 45, No. 4, 2007, pp. 890–901.
- ³Jordan, P. and Gervais, Y., Subsonic jet aeroacoustics: associating experiment, modelling and simulation, *Experiments in Fluids*, Vol. 44, No. 1, 2008, pp. 1–21.
- ⁴Tam, C. K. W., Jet noise: Since 1952, *Theoretical and Computational Fluid Dynamics*, Vol. 10, No. 1-4, 1998, pp. 393–405.
- ⁵George, W. K., Beuther, P. D., and Arndt, R. E. A., Pressure spectra in turbulent free shear flows, *Journal of Fluid Mechanics*, Vol. 148, No. 1, 1984, pp. 155–191.
- ⁶Hall, J., Pinier, J., Hall, A. M., and Glauser, M. N., Two-point correlations of the near-field and far-field pressure in a transonic jet, *Proceedings of the Fluids Engineering Summer Meeting, Miami, FL*, No. FEDSM2006-98458, ASME, 2006.
- ⁷King, R., editor, *Active Flow Control – Papers Contributed to the Conference “Active Flow Control 2006,” Berlin, Germany, September 27 to 29, 2006*, Springer, 2007.
- ⁸Gad-el Hak, M., *Flow Control: Passive, Active, and Reactive Flow Management*, Cambridge Univ. Press, 2004.
- ⁹Aamo, O. M. and Krstic, M., *Flow Control by Feedback: Stabilization and Mixing*, Springer, 2003.
- ¹⁰Samimy, M., Debiasi, M., Caraballo, E., Serrani, A., Yuan, X., Little, J., and Myatt, J. H., Feedback control of subsonic cavity flows using reduced-order models, *Journal of Fluid Mechanics*, Vol. 579, 2007, pp. 315–346.
- ¹¹Samimy, M., Adamovich, I., Webb, B., Kastner, J., Hileman, J., Keshav, S., and Palm, P., Development and characterization of plasma actuators for high-speed jet control, *Experiments in Fluids*, Vol. 37, No. 4, 2004, pp. 577–588.
- ¹²Utkin, Y. G., Keshav, S., Kim, J.-H., Kastner, J., Adamovich, I. V., and Samimy, M., Development and use of localized arc filament plasma actuators for high-speed flow control, *Journal of Physics D*, Vol. 40, No. 3, 2007, pp. 685–694.
- ¹³Sinha, A., Kim, K., Kim, J.-H., Serrani, A., and Samimy, M., Towards Feedback Control of High-Speed and High-Reynolds Number Jets, *4th AIAA Flow Control Conference, AIAA Paper 2008-3862*, 2008.
- ¹⁴Teel, A. R. and Popovic, D., Solving smooth and nonsmooth multivariable extremum seeking problems by the methods of nonlinear programming, *Proc. of the 2001 American Control Conference*, Vol. 3, 2001, pp. 2394–2399.

- ¹⁵Nelder, J. A. and Mead, R., A Simplex Method for Function Minimization, *The Computer Journal*, Vol. 7, 1965, pp. 308–313.
- ¹⁶Barton, R. R. and Ivey, J. S., Nelder-Mead Simplex Modifications for Simulation Optimization, *Management Science*, Vol. 42, No. 7, 1996, pp. 954–973.
- ¹⁷Torczon, V. J., *Multi-Directional Search: A direct search algorithm for parallel machines*, Ph.D. thesis, Rice University, 1989.
- ¹⁸Robbins, H. and Monro, S., A Stochastic Approximation Method, *Annals of Mathematical Statistics*, Vol. 22, No. 3, 1951, pp. 400–407.
- ¹⁹Kiefer, J. and Wolfowitz, J., Stochastic Estimation of the Maximum of a Regression Function, *The Annals of Mathematical Statistics*, Vol. 23, No. 3, 1952, pp. 462–466.
- ²⁰Wilde, D. J., *Optimum seeking methods*, Prentice Hall Inc., 1964.
- ²¹Wellstead, P. E. and Scotson, P. G., Self-tuning extremum control, *IEE Proceedings D*, Vol. 137, No. 3, 1990, pp. 165–175.
- ²²Zhang, Y., Stability and Performance Tradeoff with Discrete Time Triangular Search Minimum Seeking, *Proc. of the 2000 American Control Conference*, Vol. 1, No. 6, 2000, pp. 423–427.
- ²³Kastner, J., Hileman, J., and Samimy, M., Exploring High-Speed Axisymmetric Jet Noise Control Using Hartmann Tube Fluidic Actuators, *42nd AIAA Aerospace Sciences Meeting and Exhibit, AIAA Paper 2004-186*, 2004.
- ²⁴Kim, J.-H., Kastner, J., Adamovich, I., and Samimy, M., Active Control of High Subsonic Jets, *45th AIAA Aerospace Sciences Meeting and Exhibit, AIAA Paper 2007-320*, 2007.
- ²⁵Kesten, H., Accelerated Stochastic Approximation, *The Annals of Mathematical Statistics*, Vol. 29, No. 1, 1958, pp. 41–59.