

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

Aniruddha Sinha,^{*} Kihwan Kim,[†] Jin-Hwa Kim,[‡] Andrea Serrani,[§] and Mo Samimy[¶]
Ohio State University, Columbus, Ohio 43235

DOI: 10.2514/1.44012

We present and discuss results on 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 6.4×10^5 . Open-loop forcing experimental results are presented, wherein localized arc filament plasma actuators are shown to have two distinct effects on the irrotational near-field pressure, similar to their effects on the far-field acoustics reported earlier. At low forcing Strouhal numbers near the jet preferred mode, a large amplification in the pressure fluctuations is observed. At much higher forcing Strouhal numbers, an attenuation is observed in the pressure fluctuations over a broad range of excitation frequencies, especially in the axisymmetric mode of the pressure fluctuations. Previous experiments have shown that forcing the jet with these low and high Strouhal numbers results in jet mixing enhancement and far-field noise reduction, respectively. Two different gradient-free extremizing feedback control algorithms have been developed and 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. We also show that the far-field acoustic spectrum in steady-state operation of the closed-loop control is quite similar to that observed with optimal open-loop forcing.

Nomenclature

a_1, a_{\min}	= initial and minimum value of step size in a stochastic optimization process
c_1, c_{\min}	= initial and minimum value of perturbation in a stochastic optimization process
D	= nozzle exit diameter
f_F	= forcing frequency
I/O	= input/output, typically referred to a plant in a control system
$MaxS, MinS$	= user-specified simplex size constraints for simplex-based searches
m_F	= azimuthal mode number of forcing
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 the Fourier azimuthal mode m_p
Re_D	= Reynolds number based on jet nozzle exit diameter and exit velocity
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_{rms}	= time duration over which pressure samples are collected for computing the rms
U_j	= nozzle exit velocity
x	= streamwise distance downstream from nozzle exit
λ	= exponential decrement for perturbation levels in a stochastic optimization process

I. Introduction

WE PRESENT and discuss experimental results of feedback control of a high-speed high Reynolds number jet. In particular, we demonstrate control over the irrotational near-field pressure with two different goals: increasing the pressure fluctuations for mixing enhancement or decreasing the fluctuations for potential far-field noise attenuation. The feedback control strategy implemented is shown schematically in Fig. 1. In this section, we will motivate the discussion and establish the background.

The turbulent axisymmetric jet is known to be susceptible to three different instability mechanisms: 1) initial shear layer instability; 2) jet column instability; and 3) 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, thus created near 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 they play a major role in the bulk mixing of the fluids. The jet column 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 a high Reynolds number and Mach 0.9 and 1.3 jets show that $St_D \sim 0.3$ [1,2]. The azimuthal mode instability causes certain helical structures to dominate in the mixing layer. The reader is referred to [2] for a detailed review of jet mixing-layer instabilities.

Since the discovery of large-scale coherent structures over three decades ago, researchers have recognized the importance of their dynamics as potential contributors to the radiated far-field noise. However, in spite of sustained effort ever since, the exact mechanism of influence has not been pinned down as yet. The interested reader is referred to recent review papers [3,4].

Presented as Paper 0849 at the 47th AIAA Aerospace Sciences Meeting and Exhibit, Orlando, FL, 5–8 January 2009; received 24 February 2009; revision received 30 July 2009; accepted for publication 4 August 2009. Copyright © 2009 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0001-1452/10 and \$10.00 in correspondence with the CCC.

^{*}Graduate Student, Gas Dynamics and Turbulence Laboratory, 2300 West Case Road, Student Member AIAA.

[†]Postdoctoral Researcher, Gas Dynamics and Turbulence Laboratory, 2300 West Case Road; currently with POSCO Technical Research Laboratories, Korea.

[‡]Research Associate, Gas Dynamics and Turbulence Laboratory, 2300 West Case Road, Member AIAA.

[§]Associate Professor, Department of Electrical and Computer Engineering, 2300 West Case Road, Senior Member AIAA.

[¶]Howard D. Winbigger Professor of Engineering, Gas Dynamics and Turbulence Laboratory, Department of Mechanical Engineering, 2300 West Case Road; samimy.1@osu.edu. Fellow AIAA (Corresponding Author).




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

Thus, 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.

In general, jet actuation techniques strive to alter the development and dynamics of large-scale structures/instability waves; for feedback control, a real-time measurement of this effect of forcing is needed. Present technology does not allow real-time three-dimensional 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. Here, 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: 1) it is a relatively nonintrusive technique; 2) pressure, being a scalar variable, is considerably easier to measure than the velocity field inside the flow; and 3) unlike some of the velocity measurement tools, like particle image velocimetry (PIV) and laser Doppler velocimetry, pressure sensors are compact and robust and can be feasibly transferred from the laboratory setting to real-world applications [5].

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 fourth power of the wave number [6]. The nature of this rapid roll off indicates that the pressure spectrum will be dominated by the larger turbulent scales; this is termed wave number filtering. It is well known that the near-field pressure consists of two distinct components: hydrodynamic and acoustic [4]. The hydrodynamic pressure fluctuations carry the convective footprint of the underlying turbulence in the mixing layer, whereas the acoustic pressure fluctuations are characterized by acoustic waves propagating with sonic speed [7]. 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. Because of the wave number 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. [8] 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 confirmed the low-dimensional character of the azimuthal pressure spectrum that was originally reported in [9]. In particular, they showed that the instantaneous pressure signal p can be almost completely reconstructed from the sum of the axisymmetric mode $p^{[0]}$ and the first helical mode $p^{[1]}$. The energy of the lower azimuthal pressure modes was found to increase with the downstream distance, up to $x/D \sim 3$, and decrease subsequently, up to $x/D \sim 6$, before increasing slightly again. We have already discussed previously that the stronger the large-scale structures in the mixing layer, the higher the pressure fluctuations in the near field. Because 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. [8] also found that the near-field $p^{[0]}$ signal correlates better with the far-field noise than just the pressure signal by itself. Also, it was reported that the correlation is greater with the far-field microphone at downstream angles compared with the sideline angle, which is expected, as dynamics of large-scale structures are believed to be responsible for the peak of acoustic radiation to shallow angles with respect to the jet axis. 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 deg and the $p^{[0]}$ 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. However, we do show far-field acoustic data collected during the steady-state operation of the feedback control that validates this choice.

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 process [10–13]. Open-loop control is unable to cope with uncertainty in operating conditions because it operates at a preset forcing parameter. On the other hand, feedback control may add robustness to the performance in the presence of gradual variations in the operating conditions.

The reader is referred to [1,2] 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, the background noise, the instability frequencies, and the flow momentum also increase; therefore, actuators must provide excitation signals of much higher amplitude and frequencies. We have recently developed a class of localized arc filament plasma actuators (LAFAPs) that can provide excitation signals of high amplitude and high bandwidth for high-speed and high Reynolds number flow control [1,2,14,15]. 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 rectangular pulse train, for which the frequency, phase, and duty cycle can be controlled independently. The intense localized heating perturbations generated by the LAFAPs can be used to selectively excite the three different instability mechanisms.

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 LAFAPs on the mixing layer [2]. 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 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 centerline velocity decay rate beyond the end of the potential core.

Experiments were also conducted on a Mach 0.9 and $Re_D = 6.4 \times 10^5$ jet to study the effect of LAFAPs on far-field noise [1]; this is also the jet on which we implement closed-loop control in the current work. The far-field jet noise was measured using two microphones located at polar angles of 30 and 90 deg 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 $m_F = 3$. A typical metric for comparing broadband noise levels is the overall sound pressure level (OASPL). Samimy et al. [1] compared the OASPL for the forced jet

to the baseline case (the difference is denoted as ΔOASPL) at the two aforementioned stations, for a range of forcing Strouhal numbers (St_{DF}) and m_F . Several noteworthy features were observed:

1) A well-defined attenuation of noise was evident at high St_{DF} for all m_F at both stations.

2) 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 deg microphone location, a broad minimum in the ΔOASPL of ~ -1.2 dB was found to be centered around $St_{DF} \sim 2.0$; at the 90 deg location, the center of the minimum shifted to $St_{DF} \sim 3.5$, and its value was around -0.6 dB.

3) The highest azimuthal forcing mode employed ($m_F = 3$) was found to be more effective for attenuating noise, compared with the other m_F .

4) 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.

Later, we will compare the effect of closed-loop forcing on the far-field acoustics to the optimal case of open-loop forcing; the sense of this optimality is to be clarified here because the optimal forcing frequency for far-field noise attenuation has been shown to be a function of the measurement polar angle. It has been discussed before that the larger scales in the jet shear layer dominate the pressure signature in the near field, as well as the acoustic signature at the shallow angles with respect to the jet axis of the far field. On the other hand, the smaller scales of turbulence contribute most to the sideline angles of the far-field noise. Thus, it can be expected that minimizing the near-field pressure fluctuations [actually rms ($p^{[0]}$), as discussed before] may also minimize the far-field noise at downstream angles; the same cannot be claimed for sideline angles. This is also supported by the directivity of the near-field–far-field correlation reported in [8] and mentioned previously. Hence, it is appropriate to compare the far-field spectrum for the feedback-controlled case to the spectrum measured in the open-loop forcing case that resulted in the minimum noise radiated to the 30 deg location. The forcing parameters for the latter have been established to be $m_F = 3$ and any St_{DF} in the broad range from 1.1 to 2.3 [1].

We have discussed previously that the noise attenuation problem can be cast as an optimization problem, wherein we seek the plasma actuator St_{DF} that minimizes 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. A complicating factor in our high Reynolds number application is the stochastic nature of the rms of the near-field pressure when computed over sample sizes that are feasible in a feedback control setting. We have implemented two different model-free feedback controllers for online optimization. This means that on initialization with an arbitrary St_{DF} , the controllers would automatically adjust the St_{DF} to optimize the plant output (see Fig. 1). The same controller can be used for either maximization or minimization by simply switching the sign of the plant output that is fed back. In the past, we have implemented a gradient-based extremum-seeking control strategy [16]; 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 paper.

II. Experimental Setup

A. Flow Facility

All experiments were conducted in the Gas Dynamics and Turbulence Laboratory at Ohio State University. The ambient air was compressed using three five-stage reciprocating compressors, 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 in.) 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 $Re_D = 6.4 \times 10^5$. Further details about the flow facility can be found

in [1]. The boundary layer at the nozzle exit is turbulent with the boundary layer thickness and momentum thickness, estimated to be 1.0 and 0.1 mm, respectively [17].

B. Plasma Actuators and Plasma Generator System

The plasma generator system is detailed elsewhere [1,2,14,15]; here, we briefly describe some of its highlights. Each plasma actuator (LAFPA) consists of a pair of 1-mm-diam 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 the plasma. The switching circuitry for firing the individual LAFPAs is controlled by computer-generated rectangular pulses using a National Instruments (NI) digital-to-analog converter (DAC) card and LabVIEW software. Although the forcing frequency f_F , the azimuthal mode m_F , and the duty cycle can be specified independently, we have obtained the following empirical relation between the duty cycle and f_F (in kHz) based on experience with the current setup:

$$\text{duty cycle (\%)} = \begin{cases} 0.6f_F + 2, & \text{if } f_F \leq 30 \text{ kHz} \\ 0.2857f_F + 11.4286, & \text{if } f_F > 30 \text{ kHz} \end{cases} \quad (1)$$

The jet can be forced in 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 controller 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 analog-to-digital converter (ADC).

C. Near-Field Pressure Measurements

Figure 2 shows a symmetric circular array of eight 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 in.), 10 times larger than the nozzle diameter.

During baseline (unforced) experiments and open-loop forcing cases, the pressure signals were amplified and low-pass filtered at 100 kHz, and they were acquired using an eight channel NI 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 eight channels of output from the




Fig. 2 Schematic of near-field pressure sensing system using a circular array of eight pressure transducers.

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, blocks of data were collected on the eight channels simultaneously at 200 kHz. There were 8192 data points per block, producing a spectral resolution of 24.4 Hz. An average spectrum was obtained by averaging over 96 blocks of data.

D. Far-Field Acoustics Measurements

Far-field sound pressure level (SPL) was measured using two $\frac{1}{4}$ in. B&K microphones, located at 30 and 90 deg polar angles relative to the downstream jet axis; their radial locations were 103 and 48*D*, respectively, measured from the nozzle exit. The far-field acoustics were normalized to a radius of 80*D*. The sampling characteristics for the acoustics were identical to those used for the near-field pressure.

The presence of the near-field pressure measurement rig was found to have no effect on the far-field baseline SPL spectrum at the 30 deg location; however, the spectrum measured at the 90 deg location was found to have an additional 2–3 dB near its peak due to reflection off the rig. So, for proper comparison with the spectra measured with feedback control, the baseline and open-loop spectra reported herein were measured with the near-field rig in place.

III. Baseline Experimental Results

One of the first issues to be addressed 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. [8] 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 with all other streamwise locations [8]. As explained earlier, the pressure sensors should be placed outside (but close to) the jet shear layer to be able to capture the signature of the large-scale structures that are to be controlled. Kim et al. [18] 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 tips of the pressure sensors were located at a radius of 1*D*, with respect to the jet axis. All near-field pressure data presented in this paper were measured at this location.

Figure 3 presents the SPL spectra for the baseline (unforced) jet. In Fig. 3a, we show that the eight 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 spectrum of the eight channels. Figure 3b shows the SPL spectra of the first four azimuthal-mode-filtered pressure signals; this is the most that we can resolve with eight sensors. At this streamwise location, $p^{[1]}$ is the strongest, followed by $p^{[0]}$. As mentioned earlier, Hall et al. [8] have reported that the instantaneous pressure signal can be almost completely reconstructed from the $p^{[0]} + p^{[1]}$ signal. In Fig. 3b, 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 twofold. 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 Strouhal number (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.

All possible forcing azimuthal modes (viz., $m_F = 0, 1, 2, \text{ and } 3$) were explored. The St_{DF} was varied from 0.09 to 5.44 in varying step sizes. The azimuthal-mode-filtered pressure signals and their associated SPL and rms were computed in postprocessing. Figure 4 shows the SPL spectra for two representative open-loop forcing experiments compared with the baseline case. It will be shown that the case with $m_F = 0$ and $St_{DF} = 0.44$ corresponds to a maximum in the near-field pressure fluctuations, whereas the case with $m_F = 3$ and $St_{DF} = 2.72$ 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 also [1]. The high-frequency

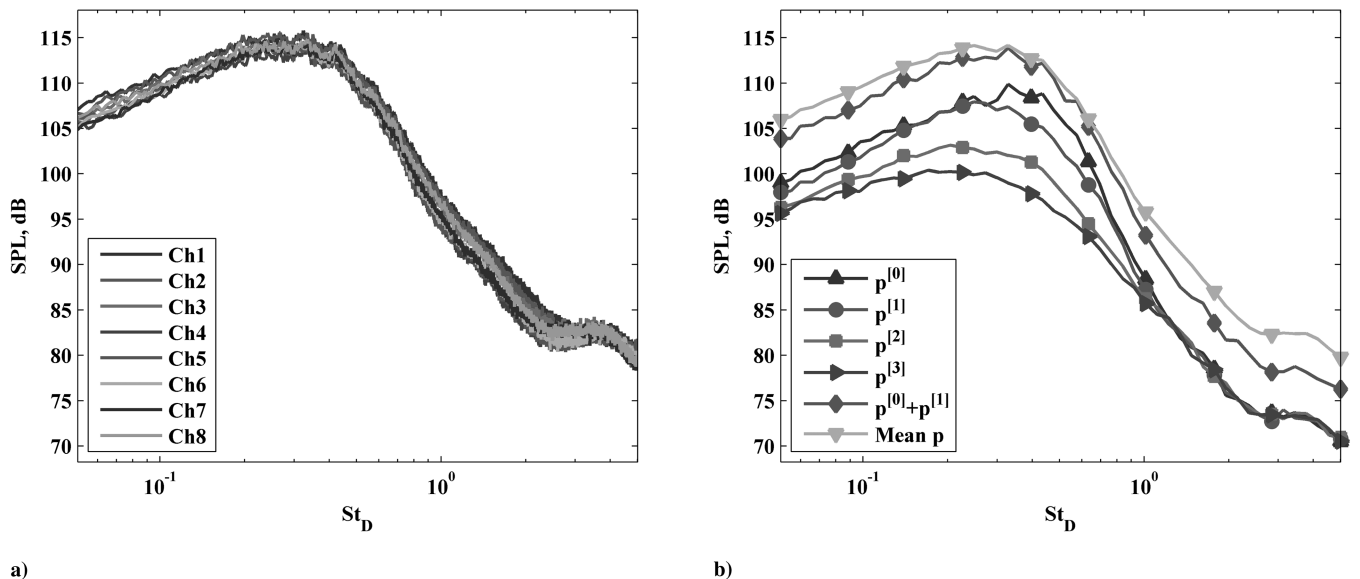


Fig. 3 Baseline near-field pressure spectra for a) the eight individual channels, and b) the various azimuthal modes. The spectra in b) are smoothed for clarity.




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

narrowband tones appearing in addition to the forcing tone and its harmonics are due to electrical noise from the actuators. These are rendered inconsequential by the filtering detailed next. During feedback control, the signals needed to be bandpass filtered to remove both the mean and the high-frequency noise. Based on the well-defined broadband peaks in the spectra in Fig. 4, a Chebyshev filter was designed with a passband from $St_D = 0.04$ to 0.73 (0.5 to 8 kHz). For accurate comparison, this filter was also applied to the pressure signals captured in open-loop forcing 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 St_{DF} and m_F . The $m_F = 0$ and 3 cases are more pertinent for our application, so that these forcing cases were explored on a finer grid of St_{DF} , hence the lack of markers for these cases. The rms of the respective signals in the baseline case are also shown for ease of comparison. We make the following observations regarding the results:

1) For all forcing azimuthal modes explored, the rms of the pressure as well as the rms 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 instability frequency; previous flow

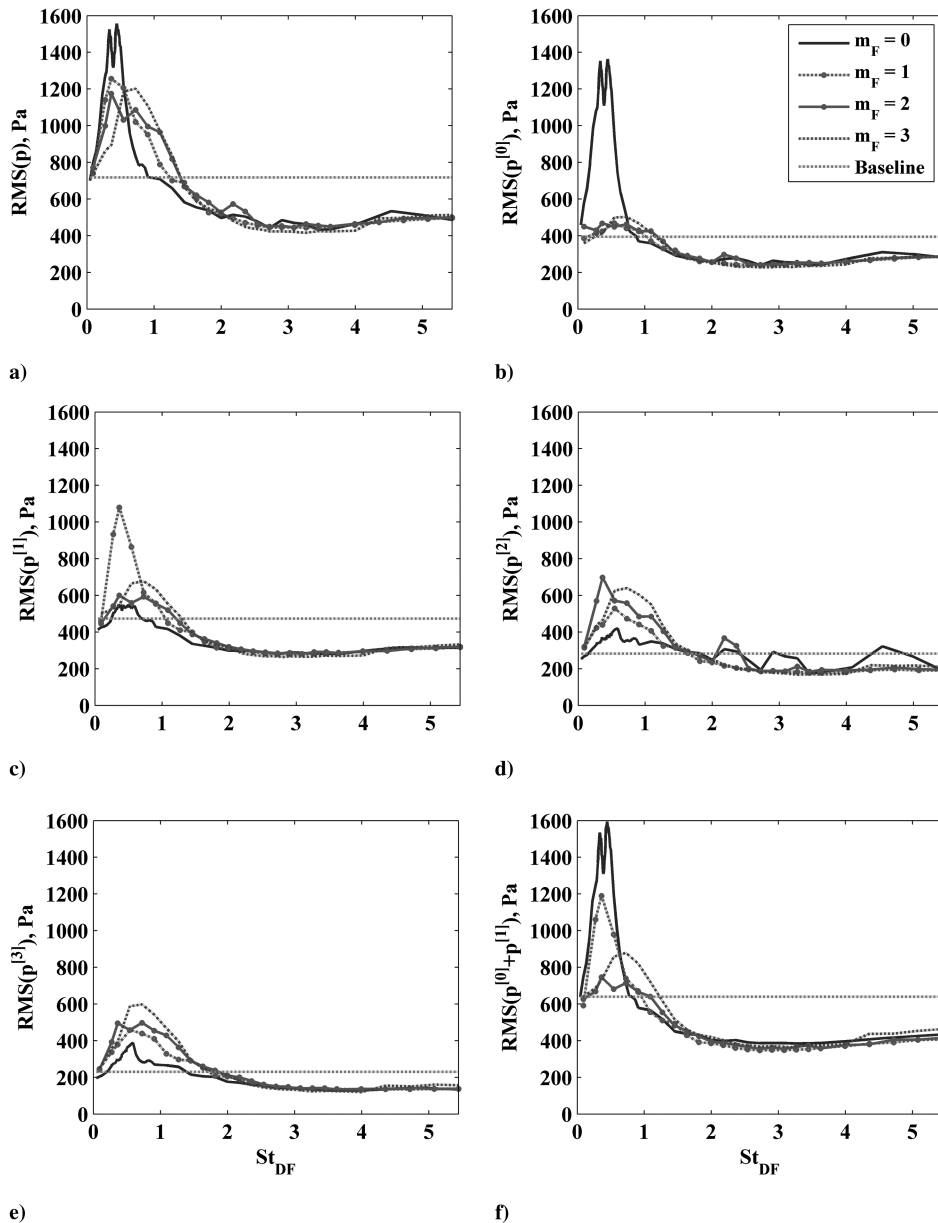


Fig. 5 Effect of open-loop forcing on near-field pressure. RMS of a) pressure, b) $p^{[0]}$, c) $p^{[1]}$, d) $p^{[2]}$, e) $p^{[3]}$, and f) $p^{[0]} + p^{[1]}$ vs St_{DF} .

visualization and PIV results have shown that the jet exhibits the largest response to forcing around this St_{DF} [2]. 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} is thought to be associated with the initial shear layer instability.

2) Fig. 5a shows that the peak in rms (p) is largest for $m_F = 0$ and smallest for $m_F = 2$ and 3.

3) For a particular m_F , the maximum rms is noted in the corresponding pressure azimuthal mode. For example, Fig. 5c shows that the maximum value of rms ($p^{[1]}$) is obtained by forcing the jet at $m_F = 1$.

4) The rms of ($p^{[0]} + p^{[1]}$) shows similar characteristics as the rms of the pressure itself.

5) The switchover from amplification to attenuation (the crossing of the curves with the baseline pressure rms 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} .

6) 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 previously). However, this crossover happens at $St_{DF} \sim 1.8$ for rms ($p^{[3]}$).

7) The effects of different forcing azimuthal modes become indistinguishable at higher forcing Strouhal numbers.

8) The location of the minimum shifts to higher St_{DF} for increasing pressure azimuthal modes. Whereas the minima for rms ($p^{[0]}$) and rms ($p^{[1]}$) occur at $St_{DF} \sim 2.7$, the other modes have their minima at $St_{DF} \sim 3.6$.

9) The minima become flatter for higher pressure azimuthal modes, especially at the high- St_{DF} end.

V. Extremizing Feedback Control Algorithms

Two different algorithms are investigated for the online optimization of the near-field pressure: the modified Nelder–Mead algorithm (MNMA) and the modified Kiefer–Wolfowitz algorithm (MKWA). An overview of the controllers based on these algorithms is presented here, but more details are provided in the Appendix.

A. Modified Nelder–Mead Algorithm

Nelder and Mead’s simplex-based direct search method [19] is a gradient-free optimization algorithm for offline optimization of static deterministic multivariable 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. Barton and Ivey [20] recommended modifications to the original algorithm for robust behavior in the presence of statistical uncertainty. Torczon [21] designed another direct search technique, the multidirectional search method, for provable convergence in a deterministic setting. She also provided empirical evidence of reliable optimization of stochastic cost functions. Our first controller implementation is a model-free direct search algorithm based on all the previously mentioned results, with several additional enhancements for use in online optimization.

Here, we will explain the operation of the algorithm in general terms; the details appear in the Appendix. The only plant input to be optimized in our application is the forcing Strouhal number St_{DF} . The appropriate simplex for this one-dimensional problem is a two-element set of vertices, representing the endpoints of a line on the St_{DF} axis. In our application, the single input/single output (SISO) cost function to be extremized is the rms of the relevant azimuthal mode-filtered near-field pressure signal. The initial choice of the vertices of the simplex (two distinct St_{DF}) must be specified by the user, or some other algorithm. The MNMA iteratively adapts the size and location of the simplex, so that the values of the plant output at both its vertices are extremized. In each iteration of the algorithm, the vertices of the simplex are ranked based on their functional values; it is this use of the mere rank that 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.

In practice, the controller issues a converging sequence of St_{DF} that extremizes the plant. On each such issue, the controller waits until the necessary pressure samples are collected to calculate the rms value, the latter then being used in the computation of the next St_{DF} to be issued. In offline optimization, one is solely concerned with the final optimal St_{DF} ; hence, the simplex is allowed to collapse to a point. However, in an online optimization application, one not only expects the process to converge upon the optimal value but to remain alive to possible subsequent changes in the optimal value itself. This argues for preventing the simplex from collapsing to a point, as such a point simplex would be insensitive to the local gradient of the input/output map. This also corresponds to the notion of persistent excitation used in the adaptive control literature, because it ensures a minimum amplitude of oscillation of the plant input. As detailed in the Appendix, this requirement necessitates several of the modifications incorporated in the MNMA. The remaining modifications address the stochastic nature of the application. Finally, our SISO application afforded an improvement to the original NMA, which was also implemented.

B. Modified Kiefer–Wolfowitz Algorithm

The second controller that we implemented was a modified stochastic optimization technique. Robbins and Monro [22] pioneered the field of stochastic approximation, the study of search techniques that successfully reach a preassigned goal in spite of uncertainty. Kiefer and Wolfowitz [23] adapted the idea to the problem of finding the extremum of a unimodal function in a stochastic setting. Although the original algorithm was gradient based, a later variant was developed for gradient-free stochastic optimization [24]. The general behavior of the MKWA is very similar to the MNMA, in that it also issues a converging sequence of St_{DF} that extremizes the plant. The difference lies in the actual nature of the converging sequence. Basically, the original Kiefer–Wolfowitz algorithm (KWA) employed a gradual convergence to address the stochasticity of the cost function. Some of the modifications in the MKWA hasten this convergence where feasible. Other enhancements keep the controller alive to possible changes in the location of the extremum, as in the MNMA. The implementation details appear in the Appendix.

VI. Results of Extremizing Feedback Control

Here, we present the results and discuss the relative merits of closed-loop experiments performed using the two different algorithms discussed previously (viz., MNMA and MKWA). Because the controller implementation was in terms of the forcing frequency instead of the Strouhal number, all of its parameters were selected in terms of frequencies; they are hereby converted into an equivalent St_D for ease of discussion. The results are shown in Figs. 6–8; the presentation scheme is elucidated here. Figure 6 shows the results of the experiments, wherein the goal was to minimize rms ($p^{[0]}$); Fig. 7 presents the outcome of the experiments, wherein the aim was to maximize rms ($p^{[0]} + p^{[1]}$). Part a) of either figure presents the transients that occurred in the St_{DF} (the plant input) when MNMA was implemented as the feedback controller, and part b) shows the corresponding transients that appear in the relevant pressure signal (the plant output). Similarly, parts c) and d) demonstrate the plant transients (input and output, respectively) with an implementation of MKWA. The St_{DF} in the legends indicate the initial values. Part e) presents the respective pressure spectra once the controllers reached steady-state operation. Finally, part f) of either figure provides the details of the open-loop results (already presented in Fig. 5) that are relevant to these experiments. Error bars are shown with $+/-\sigma$, where σ is the standard deviation with the respective T_{rms} . The spectrum of the optimal case observed in part f) is shown in part e) of the respective figure. Figure 8 shows the far-field acoustic spectra measured during steady-state operation of the controller.

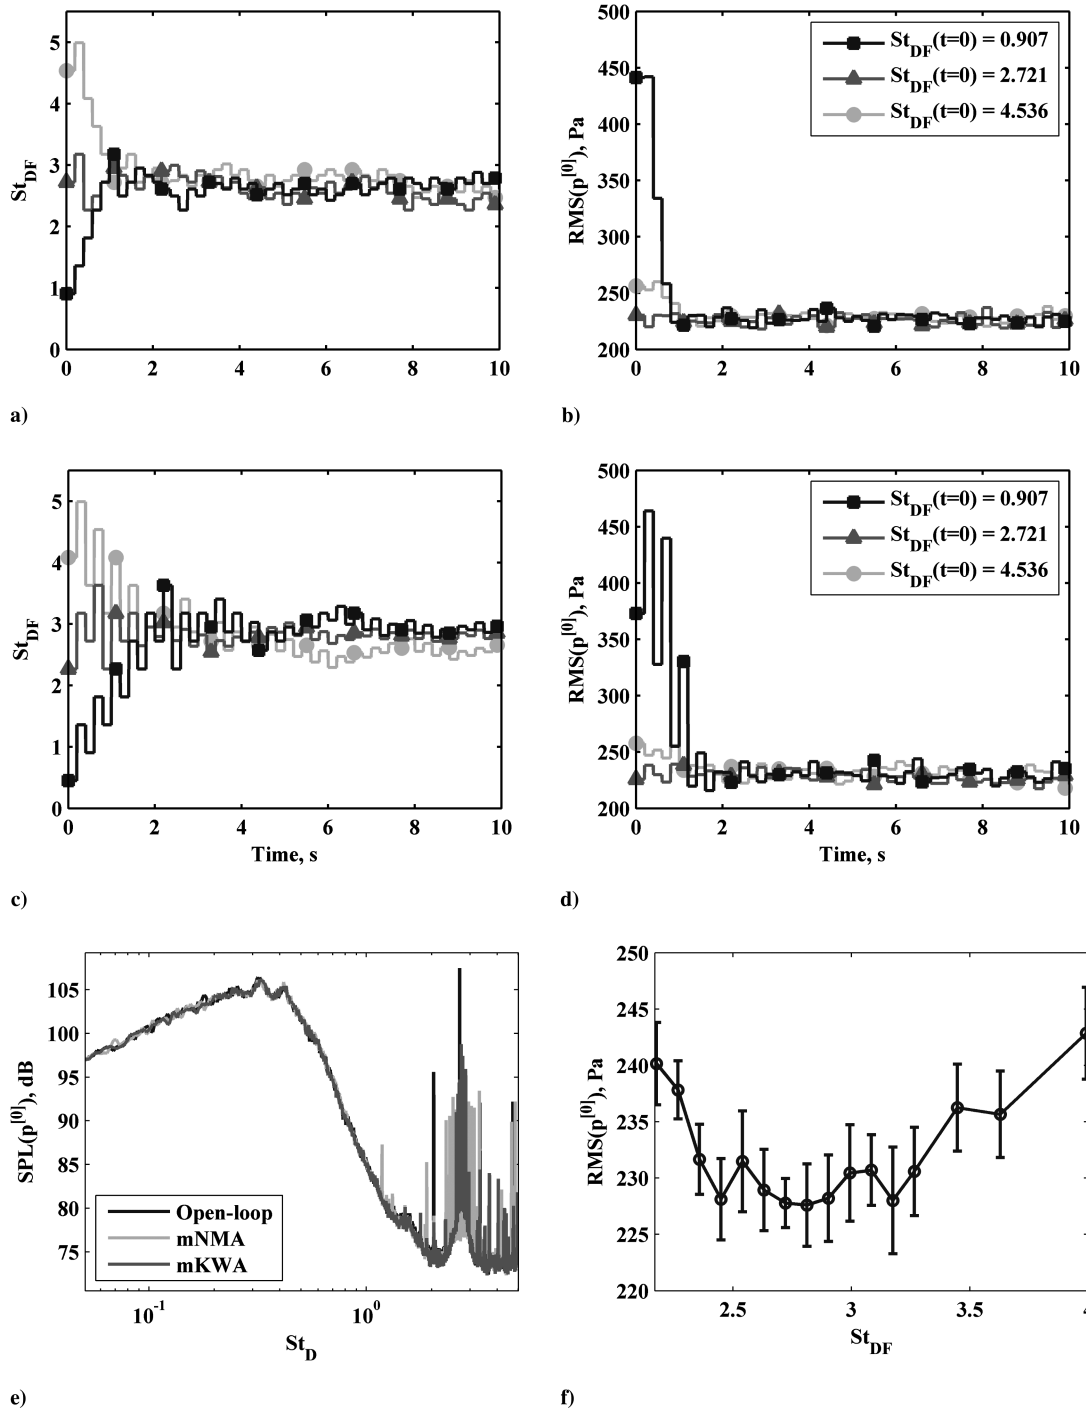


Fig. 6 Transients in minimum-seeking control with: a) and b) MNMA; c) and d) MKWA; e) steady-state near-field spectra; and f) details of pertinent open-loop results with $m_F = 3$.

For reasons stated previously, experiments aimed at potential noise attenuation were designed to minimize rms ($p^{[0]}$) in real time by optimizing the St_{DF} ; these will be called minimum-seeking experiments. The effect of different m_F cannot be distinguished in the high- St_{DF} range in which the minimum is achieved (see Fig. 5b). However, as pointed out in Sec. I, $m_F = 3$ is most effective for far-field noise suppression. Thus, this was used exclusively in closed-loop experiments aimed at minimizing noise. The details of this particular I/O map, obtained in open-loop forcing experiments, are shown in Fig. 6f. We will show convergence to the optimal St_{DF} of ~ 2.7 , starting from various initial values from 0.9 to 4.5. In addition to this near-field result, we will also show that the maximal far-field noise attenuation observed in open-loop experiments was well-

replicated in steady-state operation of the closed-loop control system.

We now turn to the goal of mixing enhancement by the maximum amplification of the large-scale structures in the jet shear layer. It has been discussed before that this amplification is reflected in a broadband increase of fluctuations of $p^{[0]} + p^{[1]}$ in the near field. Figure 5f shows that $m_F = 0$ is most effective in this regard. Hence, in closed-loop experiments focusing on mixing enhancement, $m_F = 0$ was used exclusively and rms ($p^{[0]} + p^{[1]}$) was maximized in real time by optimizing the St_{DF} ; these will be referred to as maximum-seeking experiments. A detailed version of the pertinent cases of open-loop forcing is presented in Fig. 7f. We sought to converge to the optimal St_{DF} of 0.45, starting from various initial values in the

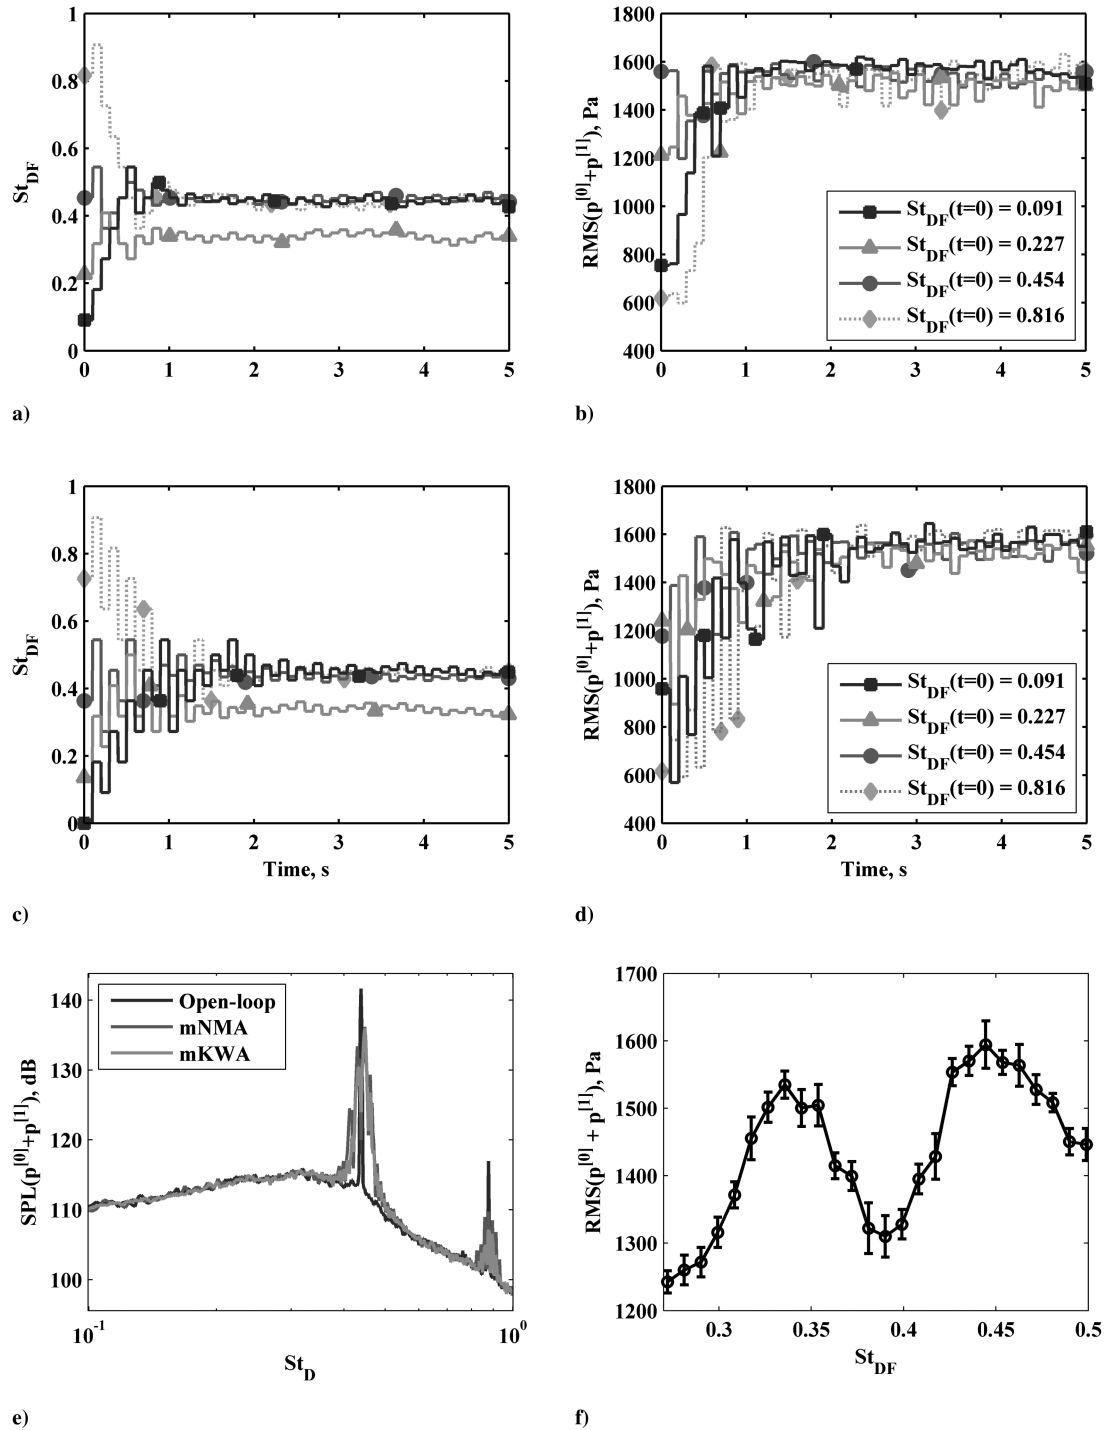


Fig. 7 Transients in maximum-seeking control with: a) and b) MNMA, c) and d) MKWA; e) steady-state near-field spectra; and f) details of pertinent open-loop results with $m_F = 0$.

range of 0.1 to 0.8; it will be shown that this was not always possible, owing to the presence of the other prominent peak at ~ 0.34 . However, because 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 Sec. V is the absence of plant dynamics. Step response experiments were conducted to verify this. A delay of no more than 20 ms was noticed between a change in St_{DF} and its effect on the near-field pressure signal. The possible sources of this delay are 1) the communication lag between the dSpace controller commanding the forcing frequency and the LabVIEW program acquiring it to control the actuators, and 2) the 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 determining 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 the closed-loop experiments is T_{rms} , the duration over which pressure samples are collected for computing the rms. Note that the waiting period of 20 ms described previously 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. 6f and 7f. For fast

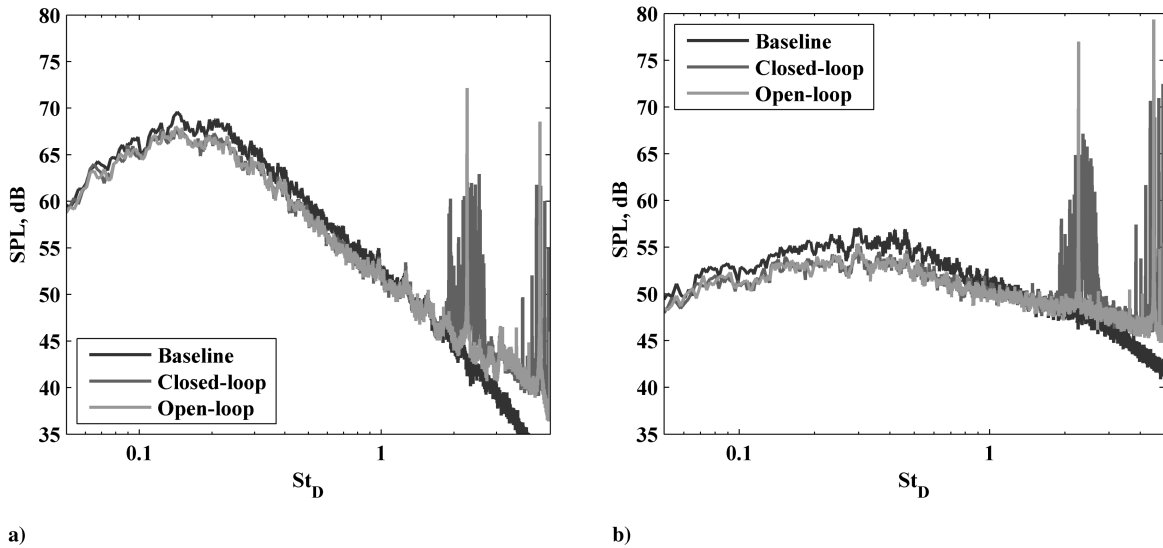


Fig. 8 Selected far-field acoustic spectra at a) 30 deg and b) 90 deg polar angles.

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. The SNR is defined as the ratio of the mean value of the plant output to its standard deviation when all parameters are held constant; it depends on the St_{DF} as well as the T_{rms} . In preliminary experiments of maximum seeking with either algorithm, $T_{\text{rms}} = 0.1$ s proved to be a satisfactory balance; the average SNR (calculated over the open-loop experimental data points indicated in Fig. 7f) was ~ 64 . However, for minimum seeking, a minimum $T_{\text{rms}} = 0.2$ s was required to have similar robust performance; the corresponding SNR (calculated from Fig. 6f) was also ~ 64 . Using $T_{\text{rms}} = 0.1$ s for minimum seeking, the SNR turned out to be 45. This explains the need for different T_{rms} for the two applications.

The control algorithms have a number of parameters that determine their performance. Some of the parameters can be selected independent of the specific application and are detailed in the Appendix. The initial choice of the remaining parameters was based on an analysis of the open-loop results, as well as the preliminary simulations. Experimental results were used to guide the selection of the final set of parameters that produced the best performance, and these are presented next. 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 due to space constraints.

A. Modified Nelder–Mead Algorithm

1. Minimum Seeking

In this application, the MNMA is used to seek the minimum of the I/O map shown in Fig. 6f; the characteristics of this map dictate the choice of the parameters of the algorithm. The map shows a very gradual slope in the range of St_{DF} from 3 to 5; thus, a large value of $MaxS = 0.45$ was selected. With $MaxS = 0.27$, we could not achieve repeatability in converging to the optimum from an initial St_{DF} of 4.5. With $MaxS = 0.9$, the convergence was repeatable and rapid, but there was a large overshoot that caused delays in settling. The I/O map also indicates that the minimum is broad; hence, an appropriate choice of $MinS$ was 0.09. 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 0.18 have also shown satisfactory steady-state behavior.

The transient behavior is demonstrated with the process initialized at three different St_{DF} . Figure 6b shows that in all three cases, the

plant output reaches and settles around 230 Pa within $5T_{\text{rms}}$ (i.e., using just five evaluations of the plant I/O map). This convergence rate represents an order of magnitude improvement compared with our results with the gradient-based extremum-seeking controller [16]. Figure 6a indicates that St_{DF} reaches ~ 2.7 within the same time, but it varies within 2.4 to 3.0 subsequently, even though $MinS = 0.09$; this can be explained by the broadness of the minimum. The three cases have similar steady-state behavior, hence we only show the near-field spectrum for an initial $St_{DF} = 4.5$ in Fig. 6e. The spectrum with MNMA is quite similar to the optimal open-loop spectrum as established from Fig. 6f, especially in the high-amplitude region of interest. The forcing tone in the feedback-controlled case is smaller but broadened due to the persistent perturbations in the St_{DF} . 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.

In Fig. 8, we compare the far-field acoustics at two different polar angles, with the optimal open-loop forcing and the steady state of closed-loop forcing; this optimality has already been explained in Sec. I. In steady state, the closed-loop controller using the MNMA algorithm is seen to create a far-field acoustic signature that is quite similar to the signature obtained in open-loop experiments. Owing to the persistent excitation, the closed-loop process results in a broad peak around the optimal St_{DF} and its harmonics, compared with the sharper but stronger peak obtained in open-loop experiments that are run at a single St_{DF} . This supports the choice of rms ($p^{[0]}$) as the quantity to minimize in the near field for far-field noise attenuation. It also provides succinct evidence of the effectiveness of closed-loop control.

2. Maximum Seeking

We now focus on the online maximization of the I/O map shown in Fig. 7f using the MNMA. It has already been discussed that the relevant portion of this map has two distinct maxima: a 1600 Pa peak at $St_{DF} = 0.45$ and a 1530 Pa peak at 0.33. 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 $St_{DF} = 0.45$ from a majority of the initial conditions within 0.1 and 0.8. Because 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.

Compared with minimum seeking, the range of initial conditions is smaller and the extremum is sharper; and so we selected $MaxS = 0.09$. Values of 0.07 and 0.18 (not shown here) were also tried with similar results, attesting to the relative robustness of the technique to the specific choice of this parameter. Further reduction of $MaxS$ led to delays in convergence due to increased uncertainty. The sharpness of the peak led us to choose $MinS = 0.009$. A value of 0.018 did not lead to substantial degradation in performance. However, with $MinS = 0.045$, large oscillations were observed in the forcing frequency, and the average output was reduced because 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.

Figure 7a shows that the plant input settles within $10T_{rms}$ from a variety of initial values. On initialization from $St_{DF} = 0.09, 0.45,$ and 0.82 , the final mean value of the input was around 0.45; starting from $St_{DF} = 0.23$, the steady-state average was around 0.33. Actually, we tested the algorithm with all initial forcing Strouhal numbers in the range from 0.09 to 0.82 in increments of 0.045. The process repeatedly converged to the peak at $St_{DF} = 0.33$ when initiated from 0.23 and 0.32; with all other initial conditions, it repeatedly reached the peak at 0.45. In Fig. 7b 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 with the minimum-seeking case. Also, the mean output is somewhat lower than the optimal value found in open-loop experiments (~ 1600 Pa). Both these effects can be explained by the sharpness of the maximum. Figure 7e presents the steady-state output spectrum for the experiment with initial $St_{DF} = 0.09$. (The other initial St_{DF} , apart from 0.23 and 0.32 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 St_{DF} 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.

B. Modified Kiefer–Wolfowitz Algorithm

The decrement exponent parameter λ was chosen as 1.0 as a compromise between rapid convergence and reliability. 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 MKWA 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 previously, in line with the MNMA experiments. The choice of the remaining parameters (viz., c_1, a_1, c_{min} , and a_{min}) is discussed next.

1. Minimum Seeking

Here we discuss experiments in minimization of the I/O map shown in Fig. 6f using MKWA. This map possesses a gradual slope at the high- St_{DF} end. If the process uncertainty 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 0.45. The value of a_1 was selected as 0.45 also to hasten convergence. This is similar, but not equivalent, to selecting the initial simplex size and $MaxS$ as 0.45 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 $c_{min} = a_{min} = 0.045$. 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 stochasticity. In this respect, the MKWA offers more flexibility than the MNMA, because in the latter, both of these effects are controlled by $MinS$.

Figure 6c demonstrates that, starting from various initial conditions, the plant input St_{DF} reaches around 2.7 within $10T_{rms}$, but it subsequently varies within 0.24 and 0.30. Figure 6d 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. 6e. We also expect the effect on the far-field acoustics to be similar. The settling time is doubled, because the MKWA makes two observations of the output for every observation made by the MNMA. Although 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.

2. Maximum Seeking

The pertinent plant I/O map for our maximum-seeking application is shown in Fig. 7f. Following the arguments laid out previously for minimum seeking, we selected $c_1 = 0.09$. In a few experiments with $c_1 = 0.07$, we did not notice any erroneous convergence behavior; however, with $c_1 = 0.18$, the oscillations took longer to settle down. Similar considerations led to the choice of $a_1 = 0.09$. The parameters c_{min} and a_{min} were both selected as 0.0045, exploiting their relationship to $MinS$ in the MNMA; the value of the latter has been justified already.

Figure 7c shows that plant input settles within $22T_{rms}$ from a variety of initial values. Starting from 0.23, the steady-state value of St_{DF} was around 0.33. However, the input settled around 0.45 when initialized from all other values in the range from 0.09 to 0.82 kHz in increments of 0.045 (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 St_{DF} of 0.32 also resulted in convergence to the peak at 0.33. Too much significance should not be read into this disparity, given that neither algorithm is designed to optimize multimodal functions. Figure 7d demonstrates that the time-domain behavior of the steady-state plant output with MKWA is very similar to that observed with 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. Figure 7e shows that the steady-state frequency-domain behavior of MKWA is similar to that of MNMA.

VII. Conclusions

In this paper, we have presented and discussed the results of the development and application of extremizing feedback control to axisymmetric jets. In particular, we have shown control authority on the irrotational near-field pressure of a Mach 0.9 jet with a Reynolds number based on jet diameter of 6.4×10^5 . Open-loop forcing using localized arc filament plasma actuators is shown to have two distinct effects on the near-field pressure; a large amplification occurs in the pressure fluctuations at an St_{DF} near the jet column mode, whereas a broad attenuation is found at higher values. 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.

We developed a feedback control strategy that minimizes the rms of the axisymmetric mode of the near-field pressure in real time by automatically tuning the St_{DF} of the LAFPA; this was motivated by the correlation of this signal with the far-field noise. Another feedback controller was designed to maximize the rms of the sum of the axisymmetric and the first helical modal components of the near-field pressure by tuning the St_{DF} ; the motivation was the correlation of this signal with the bulk mixing in the shear layer. Thus, both the far-field noise mitigation and the bulk mixing enhancement

applications can be regarded as online extremization problems involving single-variable cost functions. An important characteristic of both applications is the high level of uncertainty 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 stochastic functions. Several novel enhancements were incorporated into 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 effects on the near-field pressure after brief transient phases. We have also demonstrated that feedback control can attenuate the far-field noise to a degree similar to the maximal observed in previous open-loop forcing experiments. The modified Kiefer–Wolfowitz method was found to be more robust, but the convergence rate was typically halved in comparison with 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 MNMA would be very useful in this regard.

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

Appendix

I. Modified Nelder–Mead Algorithm

For brevity, we will discuss only the important changes that were made to the original NMA [19] in developing the MNMA; the relevant steps are serialized in the flowchart in Fig. A1. The first modification is the imposition of user-defined constraints on the

simplex size (viz., $MaxS$ and $MinS$). The lower bound $MinS$ addresses the requirement of persistent excitation, as discussed in Sec. V. 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 (see Fig. A1). 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 each iteration of the MNMA, the simplex size constraint must be explicitly addressed; this accounts for the decision steps marked 1, 3, and 4 in Fig. A1; these are absent in the original NMA. This requirement also makes the computations in the expansion and contraction steps slightly more involved.

The iteration in MNMA starts with the reflection step shown in Fig. A1; in proceeding through the iteration, the first noteworthy modification is the inclusion of the decision box 1, for reasons stated previously. The algorithm reaches this point if the reflection produces a new best vertex; the original NMA follows this up with an expansion step. However, if the simplex has already attained its maximum allowable size, then it cannot be expanded further. Instead, on the N branch of decision box 1, the new simplex is set as $\{B, R\}$. Then, in the ensuing reflection, B would be reflected through R to give the E vertex that an expansion step would have evaluated next.

A major difference between the MNMA and NMA is at point 2 in Fig. A1. At this step, the original NMA would have set the new simplex as $\{B, R\}$ before proceeding to the next reflection. There is an issue with this scenario that is peculiar to a SISO problem. Using the typical values of $\alpha = 1$ and $\gamma = 2$, the next vertex to be evaluated on reflection of B through R would be E again. Because 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 subinterval. An additional complication is introduced by the lower bound $MinS$ on the simplex size. Let us

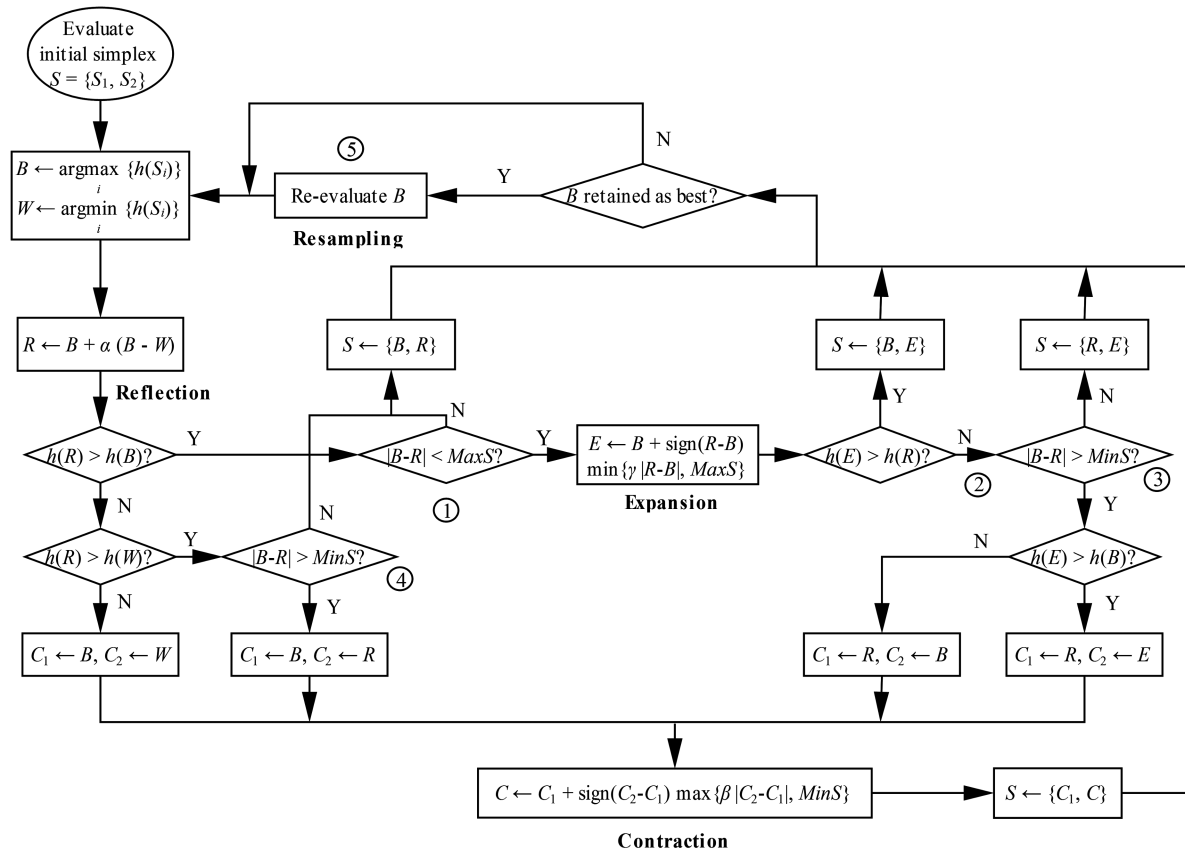


Fig. A1 Flowchart of the MNMA for maximizing the cost function $h(\bullet)$. To minimize $h(\bullet)$, its negative is maximized. The major points of departure from the original NMA are serialized.

consider the logical step in case the simplex cannot be contracted any further (i.e., if $|B - R| = \text{Min}S$, then R , being the best vertex, should be retained in the simplex). E has been newly evaluated, and hence should not be reevaluated. And so, the only option is to reevaluate B . This can be achieved by setting the simplex to $\{R, E\}$ and looping back to a reflection. If, on the other hand, $|B - R| > \text{Min}S$, then we borrow from Nelder and Mead's idea of the inside/outside contraction [19]. In particular, if $h(E) > h(B)$, then it is logical to contract the $\{R, E\}$ simplex toward R ; otherwise, the $\{R, B\}$ simplex is to be contracted toward R .

The next modification is the addition of decision step 4 in Fig. A1. At this point, if the simplex size is the smallest that it can be, then the ensuing 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.

The original NMA has a shrink step, which, for a SISO problem, duplicates the contraction step, albeit with a possibly different coefficient; it is therefore not included in the proposed algorithm.

The final modification is the resampling step adopted from Barton and Ivey [20] to address the stochasticity in our application. 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. And so, at the end of an iteration (before the next reflection step), if B is retained as the best vertex, then it is reevaluated, and its old functional value is replaced with the new one in subsequent rank determinations.

Following previous work [19,21], the transformation parameters were set at $\alpha = 1$, $\beta = 0.5$, and $\gamma = 2$. Search techniques typically require the user to specify an initial guess of the goal. In our implementation, this guess was used 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 $S_2 = S_1 + Ss_0$. The value of Ss_0 must be chosen, such that the algorithm can reliably rank its vertices initially; this means that S_1 and S_2 must be spaced far enough apart that their functional values have the correct rank in spite of uncertainty. For fastest convergence, one should always specify $Ss_0 = \text{Max}S$. The choice of $\text{Max}S$ and $\text{Min}S$ are detailed in the main text.

II. Modified Kiefer–Wolfowitz Algorithm

We will discuss only the important changes that were made to the original Kiefer–Wolfowitz algorithm (KWA) [23] in developing the MKWA; its pseudocode appears at the end. 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 uncertainty 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 $h(\bullet)$. In the KWA, the gradient $E'\{h(y_n)\}$ is estimated as a finite difference using a positive and negative perturbation to y_n . However, in our application, the gradient of the plant I/O map may become very small in the vicinity of the extremum, thereby slowing down convergence if the gradient descent method is used. Hence, in the MKWA, we incorporated Wilde's [24] idea of using only the sign of the estimated gradient. With this, at the end of each iteration, the new center for the next pair of perturbations is computed as

$$y_{n+1} = y_n + a_n \text{sign}\{h(y_n + c_n) - h(y_n - c_n)\} \quad (\text{A1})$$

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 technique [23] makes the following assumptions on the sequences c_n and a_n :

$$\lim_{n \rightarrow \infty} a_n = 0, \quad \lim_{n \rightarrow \infty} c_n = 0, \quad \sum_{n=1}^{\infty} a_n = \infty, \quad \sum_{n=1}^{\infty} \left(\frac{a_n}{c_n}\right)^2 < \infty \quad (\text{A2})$$

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 $c_1 > 0$.

The next modification is the application of lower bounds a_{\min} and c_{\min} on a_n and c_n , respectively, to provide the persistent excitation mentioned in Sec. V. 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 final modification incorporates the acceleration protocol proposed by Kesten [25]. With this, y_k is still updated with the iteration index k , but 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; however, this was found to be counterproductive for our online optimization application. 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} ; the reasoning is as follows. The c_{\min} cannot be too large, otherwise the variable y as well as the output $h(\bullet)$ 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 stochastic uncertainties, 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 lockstep with the step-size decrements. Of course, because 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, because the perturbation level becomes constant at c_{\min} after a few iterations.

The pseudocode of the MKWA for function maximization is as follows:

```

Given an initial value  $y_i$ . Also given the parameters  $c_1, a_1, c_{\min}, a_{\min}$ , and  $\lambda$ . Initialize  $k$  and  $n$  with 1.
WHILE (not manually overridden)
   $\text{sgn}_k \leftarrow \text{sign}\{h(y_k + c_n) - h(y_k - c_n)\}$ 
  IF ( $k > 1$  AND  $\text{sgn}_k \neq \text{sgn}_{k-1}$ ), 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
   $y_{k+1} \leftarrow y_k + a_n \text{sgn}_k$ 
   $k \leftarrow k + 1$ 
ENDWHILE.

```

Acknowledgments

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

References

- [1] Samimy, 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. doi:10.2514/1.27499
- [2] Samimy, 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, May 2007, pp. 305–330. doi:10.1017/S0022112007004867
- [3] Tam, C. K. W., "Jet Noise: Since 1952," *Theoretical and Computational Fluid Dynamics*, Vol. 10, Nos. 1–4, 1998, pp. 393–405. doi:10.1007/s001620050072
- [4] Jordan, P., and Gervais, Y., "Subsonic Jet Aeroacoustics: Associating Experiment, Modelling and Simulation," *Experiments in Fluids*, Vol. 44, No. 1, 2008, pp. 1–21. doi:10.1007/s00348-007-0395-y
- [5] Tinney, C. E., Ukeiley, L. S., and Glauser, M. N., "Low-Dimensional Characteristics of a Transonic Jet. Part 2. Estimate and Far-Field Prediction," *Journal of Fluid Mechanics*, Vol. 615, Nov. 2008, pp. 53–92. doi:10.1017/S0022112008003601

- [6] George, W. K., Beuther, P. D., and Arndt, R. E. A., "Pressure Spectra in Turbulent Free Shear Flows," *Journal of Fluid Mechanics*, Vol. 148, Nov. 1984, pp. 155–191.
doi:10.1017/S0022112084002299
- [7] Tinney, C. E., and Jordan, P., "The Near Pressure Field of Co-Axial Subsonic Jets," *Journal of Fluid Mechanics*, Vol. 611, Sept. 2008, pp. 175–204.
doi:10.1017/S0022112008001833
- [8] Hall, J., Pinier, J., Hall, A. M., and Glauser, M. N., "Two-Point Correlations of the Near and Far-Field Pressure in a Transonic Jet," *Fluids Engineering Summer Meeting*, American Society of Mechanical Engineers Paper FEDSM2006-98458, Fairfield, NJ, 2006.
- [9] Michalke, A., and Fuchs, H. V., "On Turbulence and Noise of an Axisymmetric Shear Flow," *Journal of Fluid Mechanics*, Vol. 70, No. 1, 1975, pp. 179–205.
doi:10.1017/S0022112075001966
- [10] Aamo, O. M., and Krstic, M., *Flow Control by Feedback: Stabilization and Mixing*, Springer, New York, 2003.
- [11] Gad-el Hak, M., *Flow Control: Passive, Active, and Reactive Flow Management*, Cambridge Univ. Press, New York, 2000.
- [12] King, R. (ed.), *Active Flow Control: Papers Contributed to the Conference "Active Flow Control 2006"*, Springer, Berlin, Aug. 2007.
- [13] Samimy, 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, May 2007, pp. 315–346.
doi:10.1017/S0022112007005204
- [14] Samimy, 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.
doi:10.1007/s00348-004-0854-7
- [15] Utkin, 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: Applied Physics*, Vol. 40, No. 3, 2007, pp. 685–694.
doi:10.1088/0022-3727/40/3/S06
- [16] Sinha, 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.
- [17] Kearney-Fischer, M., Kim, J.-H., and Samimy, M., "Control of a High Reynolds Number Mach 0.9 Heated Jet Using Plasma Actuators," *Physics of Fluids*, Vol. 21, No. 9, 2009, pp. 1–14.
doi:10.1063/1.3210771
- [18] Kim, J.-H., Kastner, J., and Samimy, M., "Active Control of a High Reynolds Number Mach 0.9 Axisymmetric Jet," *AIAA Journal*, Vol. 47, No. 1, 2009, pp. 116–128.
doi:10.2514/1.36801
- [19] Nelder, J. A., and Mead, R., "A Simplex Method for Function Minimization," *Computer Journal*, Vol. 7, No. 4, 1965, pp. 308–313.
doi:10.1093/comjnl/7.4.308
- [20] Barton, R. R., and Ivey, J. S., "Nelder–Mead Simplex Modifications for Simulation Optimization," *Management Science*, Vol. 42, No. 7, 1996, pp. 954–973.
doi:10.1287/mnsc.42.7.954
- [21] Torczon, V. J., "Multi-Directional Search: A Direct Search Algorithm for Parallel Machines," Ph.D. Thesis, Rice Univ., Houston, TX, 1989.
- [22] Robbins, H., and Monro, S., "A Stochastic Approximation Method," *Annals of Mathematical Statistics*, Vol. 22, No. 3, 1951, pp. 400–407.
doi:10.1214/aoms/117729586
- [23] Kiefer, J., and Wolfowitz, J., "Stochastic Estimation of the Maximum of a Regression Function," *Annals of Mathematical Statistics*, Vol. 23, No. 3, 1952, pp. 462–466.
doi:10.1214/aoms/117729392
- [24] Wilde, D. J., *Optimum Seeking Methods*, Prentice–Hall, Upper Saddle River, NJ, 1964.
- [25] Kesten, H., "Accelerated Stochastic Approximation," *Annals of Mathematical Statistics*, Vol. 29, No. 1, 1958, pp. 41–59.
doi:10.1214/aoms/1177706705

M. Glauser
Associate Editor