

Tracking of buffet oscillations in transient simulations of a transonic airfoil using sliding-window Dynamic Mode Decomposition

Sathsara Dias^{*}, Marko Budišić[†], Pat Piperni[‡], Brian Helenbrook[§]

Buffet is a self-sustaining oscillation that can endanger an aircraft in transonic flight. Early automatic detection of the onset of buffet from simulation data is a necessary step in optimization algorithms used to design new aircraft geometries. We use the Dynamic Mode Decomposition (DMD) eigenvalues to detect the onset of buffet from transient simulations of the OAT15A airfoil across a range of angles of attack. As during transients different eigenvalues may be expressed at different times, we process shorter segments of the data (sliding windows) in order to obtain robust, time-varying estimates of frequency of oscillation and decay/growth rates. We demonstrate that the onset of oscillation can be identified by analyzing DMD eigenvalues, and that physically different mechanisms of oscillation can additionally be identified from spatial profiles associated with DMD modes. The dominant buffeting modes can be detected through DMD before they become apparent in the physically-relevant quantities, such as the coefficient of lift. We expect that such early detection of the onset of buffet could reduce the computational burden of aircraft design and expand the feasible design space.

Contents

I	Introduction	2
II	Preliminary analysis of transonic buffet	4
	II.A Qualitative behavior of the simulated flow	4
	II.B Analysis of the lift coefficient time series	4
III	Sliding-window Dynamic Mode Decomposition (DMD) of transient data	8
	III.A Overview of the DMD algorithm	8
	III.B Stabilization of the DMD spectrum using sliding window	9
	III.C The effect of window location on DMD mode profiles	11
IV	Parametric analysis of the buffet using DMD	12
	IV.A The change in spectrum across parameter variation	12
	IV.B Tracking of modes across parameter variation	13
	IV.C Harmonics, secondary mechanisms, and aliasing	14
V	Discussion and conclusions	20
VI	Acknowledgments	20

Nomenclature and acronyms

All quantities are non-dimensional unless indicated otherwise.

α = angle of attack [$^\circ$]

^{*}Graduate Student, Department of Mathematics, diassl@clarkson.edu

[†]Assistant Professor, Department of Mathematics, mbudisic@clarkson.edu

[‡]Associate Professor, Department of Mechanical and Aerospace Engineering, ppiperni@clarkson.edu, AIAA Member

[§]Professor, Department of Mechanical and Aerospace Engineering, ppiperni@clarkson.edu, AIAA Member Grade?

L	=	chord length
ρ	=	fluid density
γ	=	specific heat ratio
Ma	=	Mach number
Re	=	Reynolds number
f	=	frequency of oscillation
St	=	Strouhal number $St = fL/a$
a	=	speed of sound
C_L	=	lift coefficient
T_w	=	size of the data window
FFT	=	Fast Fourier Transform
DMD	=	Dynamic Mode Decomposition
PSD	=	Power Spectral Density
STFT	=	Short-time Fourier transform
C_p	=	pressure coefficient
Δt	=	time step
T	=	duration of the simulation

I. Introduction

Transonic shock buffet is a self-sustained oscillation experienced by some airfoil or wing configurations at transonic Mach numbers and moderate angles of attack (α). The appearance of buffeting limits the aircraft performance envelope due to reduced flight handling qualities and impacts on structural integrity. For these reasons, sustained buffet can be dangerous for commercial and military aircraft, and both are designed to avoid it during normal operations. Determining the buffet limits of an aircraft is a challenging task that often requires computationally intensive non-linear transient simulations. The goal of this work is to see if the dynamic mode decomposition can reduce this computational effort. As an initial exploration, here we examine the onset of buffet on 2D airfoils recognizing that the buffet mechanism on 3D wings is considerably more complex [1–4].

Transonic flow around an airfoil features several structures that interact during a buffet: a supersonic region near the leading edge terminating in a pronounced discontinuity/shock, a separation point downstream of the shock around the mid-chord point, and a backflow in the trailing edge area, as indicated in Figure 1. With increasing α , the supersonic region grows and its trailing shock collides with the separation point, causing a decrease in lift. This collision shrinks the supersonic region, separating the shock and the separation point, restoring the lift, and setting up the next cycle. This periodic increase and decrease in lift is called *transonic buffet* and manifests as a vibration felt by the wing and the rest of the aircraft.

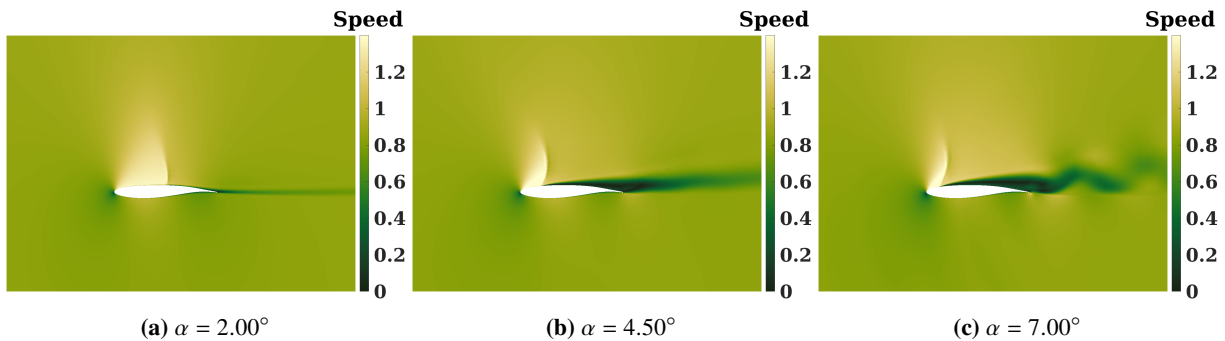


Fig. 1 Local speed around the ONERA OAT15A airfoil during a transonic flight simulation ($t = 100$ in all cases). The three panels correspond, respectively, to an α below the buffet regime, on the verge of buffet, and above the buffet regime. The shock, which oscillates during the buffet, is the sharp vertical line between supersonic (bright) and subsonic speeds.

Despite the critical importance of buffet in aircraft design, there is no single, universally-accepted definition of buffet onset [1–6]. Some examples of empirical detection criteria used are a decrease in lift coefficient C_L with respect

to the Mach number in the transonic regime, the slope of the normal force coefficient dC_N/dC_L with respect to the angle of attack α , and the divergence of the trailing edge pressure. For a historical overview of criteria see [7].

Sartor [8] and Crouch [5, 6] point out that the mathematical mechanism behind buffet onset in airfoils is the Hopf bifurcation, where a linearly-stable steady solution (equilibrium) becomes unstable at a critical parameter value, leading to a birth of a stable limit cycle that becomes the new attractor for the dynamics.

Various airfoil shapes may experience buffet at different combinations of parameters, although the phenomenon itself is the same. We therefore focus on a single configuration, the ONERA OAT15A airfoil flying at a steady speed of $Ma = 0.73$, with the buffet induced by an increase in α . Related configurations has been studied experimentally in wind tunnels [1, 8, 9] with the buffet onset occurring between $3.25^\circ < \alpha < 3.50^\circ$, with oscillations at frequencies of $0.06 < St < 0.09$ [1, 8] where St is the Strouhal number defined as the frequency of oscillation normalized by the flow velocity and chord length. The same range of frequencies is identified by additional analysis of simulations and by linear stability analysis in [1, 8]. For a slightly different Re and Ma configuration, [6] finds the critical α of 3.03° .

Instead of the model-based analysis (linearization), data-driven techniques can be used on both simulated and experimental fields; in particular, proper orthogonal decomposition (POD) [10] and dynamic mode decomposition (DMD) [11, 12] have been investigated in the literature. Both algorithms apply a numerical separation of variables to a snapshot matrix containing the time evolution of data, producing pairs of *modes* (spatial profiles) and *coefficients* (time evolution) that act as elements in the linear decomposition of the original data. In the case of POD, elements in the decomposition are mutually orthogonal, resulting in the highest-possible rate of capture of the ℓ_2 -norm when only a subset of modes is retained. By contrast, DMD modes are computed so that their time evolution is specified by a single (complex-valued) frequency, simplifying the interpretation of the time evolution; on the other hand, the orthogonality of the modes is lost. Both POD and DMD have been given additional theoretical justification which connects them to various aspects of the underlying dynamical models; see [13] for connections between POD and coherent structures in flows, and [14–17] for connections between DMD and Koopman operators of dynamical systems.

Liu et al. [18] investigate transonic shock buffet instability of unsteady flow simulations for the OAT15A and NASA(SC)-0714 airfoils using DMD. The variation of the growth rate of the dominant DMD mode vs. α graph has been used to predict the onset of shock buffet. While this method correctly identified the onset of buffet, it gave an additional false-positive result bringing into question its predictive power.

Poplinger’s [2] DMD analysis showed that the first pair of dominant modes is associated with the buffeting frequency, and their mode profiles represent shock wave/boundary-layer interaction. Other associated DMD modes exhibited more structures in the shock wave traveling region. Using DMD, Poplinger [2] constructed three possible reduced-order models (ROM) for transonic flows.

In other studies, researchers have used the POD to construct reduced-order models (ROM) and identify persistent structures for transonic flows. For example, in Szubert’s analysis [19] of the OAT15A airfoil at $Ma = 0.73$ and $\alpha = 3.5$ deg, the interaction between upstream and downstream von Kármán eddies of the OAT15A airfoil was tracked by wavelet analysis, autoregressive (AR) modeling, and POD. Ohmichi et al. [20] performed numerical simulations of three-dimensional buffet phenomena on a swept wing using incremental POD and incremental DMD based on the Hemati et al. [21] algorithm to consume low memory. This analysis found that spanwise periodic structures arise on the wing’s topside and correlate to previously discovered buffet cells. Kou and Zang [22] applied DMD to build reduced-order models and demonstrated their technique on a transonic NACA0012 airfoil, including the buffeting and wake-oscillation regimes. While they do not primarily detect the onset of the oscillations, they do demonstrate that DMD remains a viable ROM technique across the range of flow behaviors analogous to those studied in our work.

The early DMD algorithms were theoretically well-supported only when applied to long segments of statistically steady data, where the frequency content did not change in time. When applying DMD to segments of trajectories, especially when those trajectories are transient, for example connecting unstable and stable solutions, Page and Kerswell [23] show that a careless choice of trajectory segment may lead to a disorganized and non-reproducible DMD spectrum that does not represent either the dynamics around the “departure” or around the “approach” of the trajectory. While theoretical backing for the Koopman analysis of such time-varying data is in active development [24–26], practical modifications of the DMD involve working with the sliding data window [23, 27]. This is the approach we adopt here.

In this paper, we demonstrate how a sliding window DMD can be coupled with parametric analysis to correctly identify the onset of buffet, calculate the frequency of oscillations, and disambiguate mechanisms leading to oscillations based on their spatial DMD modes. We demonstrate this on a OAT15A airfoil simulated using an unsteady Reynolds-averaged Navier-Stokes (RANS) code [28]. To this end, Section II describes the flow configuration, and uses a Fourier analysis as a baseline calculation of the buffet frequency. Section III describes the DMD algorithm along with the modifications used, and shows how its outputs reflect the transient nature of the flow at a single α value. Section IV

shows how the outputs change when the α is varied, and how the onset of buffet is reflected in them. We conclude in Section V by discussing other flow features reflected in the analysis, as well as prospects for using the outputs in an optimization code.

II. Preliminary analysis of transonic buffet

A. Qualitative behavior of the simulated flow

We study the ONERA OAT15A airfoil at a steady speed, with the onset of buffet controlled by an increase in the angle of attack α . The flow is simulated using the NSU2D numerical code, developed and made available by D. Mavriplis [28–31], which solves the unsteady RANS equations with the Spalart-Allmaras turbulence model and compressibility correction on an unstructured 2D grid, using a nonlinear multigrid solution scheme with matrix dissipation. The results are made nondimensional using the free-stream density and speed of sound and the chord length. The triangular mesh used has a total of 34 521 nodes. The numerical timestep was held constant at 0.01, although in our analysis we used solutions downsampled to the effective timestep of $\Delta t = 1.00$, except where explicitly noted otherwise; all simulations have the total duration of $T = 100$.

The airfoil and flow parameters were set to the following values:

- chord length $L = 1$,
- fluid density $\rho = 1$,
- specific heat ratio $\gamma = 1.4$,
- freestream speed $V = \text{Ma} = 0.73$,
- Reynolds number $\text{Re} = VL/\nu = 3 \times 10^6$.

where ν is the kinematic viscosity. The Prandtl number of the gas was assumed to be that of air. Any oscillation frequency f is reported through its Strouhal number computed as $\text{St} = fL/a$, where a is the speed of sound.

The change in α was effected by the change in the direction of the background flow, without any changes to the geometry of the boundaries at the airflow surface or the associated change in the computational mesh. The initial velocity profile of each simulation was set to the free stream profile. The boundary conditions on the airfoil surface were no slip and no heat transfer (adiabatic assumption).

The output of the code is a time-resolved set of fields associated with the flow. Our analysis uses the horizontal and vertical components of the velocity, and pressure; eddy viscosity and internal energy are also computed by the code, but they were not used in our analysis.

The onset of transonic buffet has been linked to the Hopf bifurcation mechanism [5, 6, 8], schematically shown in Figure 2. In the pre-buffet (low α) regime, the flow is a stable fixed point of the governing differential equation, with its Jacobian featuring a pair of complex eigenvalues that capture the decay rate and frequency of the oscillatory transient. Even though the dynamics evolves in a high-dimensional (computationally) or infinite-dimensional (theoretically) space, the destabilization is associated with the two-dimensional subspace spanned by eigenvectors associated with the mentioned conjugate eigenvalue pair. Components of the flow in this subspace change from stable to unstable when α is increased through the critical value, indicating an oscillatory transient that changes from decaying to growing behavior. At the critical α value, a stable limit cycle detaches from the fixed point and attracts the trajectories. As α further increases, the diameter of the limit cycle grows, resulting in a larger amplitude periodic oscillation; importantly, through the onset, the frequency of oscillation remains constant.

In experimental or closed-box computational studies, analytic bifurcation analysis is not feasible. To benchmark our proposed method, introduced in Section III.A, we estimate the onset of the buffet in two additional ways: peak-to-peak analysis of the oscillations in lift and speed, and Fourier analysis of the lift coefficient time-trace.

B. Analysis of the lift coefficient time series

An airfoil’s lift coefficient is a dimensionless coefficient representing the lift force on an airfoil $C_L = F_L/(\rho V^2 L/2)$ where F_L is the lift force created per unit depth of the airfoil. To determine C_L ’s period of oscillation at a specific α , the peak-to-peak distance between the last two consecutive peaks in the time-trace was measured, and reported for eleven α values in Figure 3. Similarly, the magnitude of peaks was compared to estimate whether the oscillations grow, decay, or stabilize.

We report only a subset of analyzed angles in Figure 3, but enough to notice three general regimes: (i) the pre-buffet regime, where the oscillations decay, (ii) the buffet regime, where the oscillations stabilize in the band $0.06 < \text{St} < 0.09$,


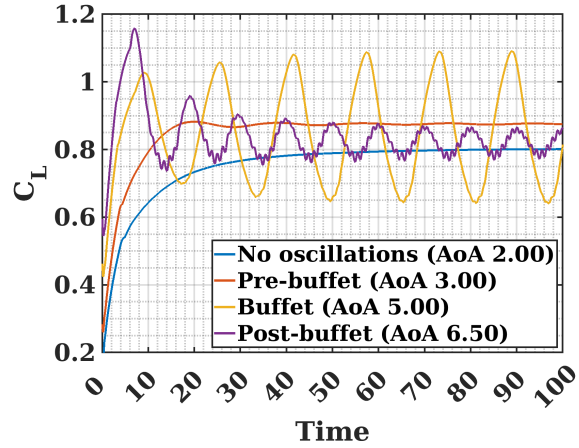


Fig. 2 Sketch of the phase portraits as the flow undergoes a Hopf bifurcation that manifests as the buffet. (a) Below the critical α , the trajectory moves from the initial state to a stable fixed point (steady flow). (b) At the critical α , the limit cycle is born. (c) Above the critical α , the trajectory moves from the initial state to the stable limit cycle (buffet).

α	Frequency (St)	Stability
1.00	N/O	N/O
2.00	N/O	N/O
2.50	N/O	N/O
3.00	0.063	decay
3.25	0.066	steady
3.50	0.067	steady
4.00	0.069	steady
5.00	0.073	steady
6.00	0.081	steady
6.50	0.114	decay
7.00	0.135	decay

(a) Frequency and growth/decay of oscillation measured by peak-to-peak analysis



(b) Time series of lift oscillation

Fig. 3 The lift coefficient C_L for α between $\alpha = 1.00^\circ$ and $\alpha = 7.00^\circ$. N/O indicates a non-oscillatory time trace. The observed buffet frequency is in the band $0.06 < St < 0.09$ which is consistent with similar airfoils described in [8].

and (iii) the post-buffet regime, where the dominant frequency sharply increases and the oscillations again decay. Notice that around the transition, with $3.00^\circ < \alpha < 3.25^\circ$, the oscillation frequency remains constant, which is consistent with the underlying Hopf mechanism.

Oscillations in the lift are closely mirrored by the oscillations in local speed. Figure 4 shows the local speed along an elliptic loop surrounding the wing, where the shock is clearly present as a near-discontinuity in the local speed. Comparing the lift coefficient trace and the trace of the shock location demonstrates that the oscillations in speed mirror the oscillations in the lift coefficient, justifying the use of velocity as the input in our later analysis.

As an alternative to peak-to-peak analysis, we compute the Power Spectral Density (PSD) using a discrete Fourier transform (via Fast Fourier Transform, FFT) of the data. While such analysis is standard and reliable for data in which the transient component is negligible or dominates during only a small portion of the total data set, our data is highly transient. Computing the PSD of the entire timetrace results in the transient severely affecting the “expected” result. For example, at $\alpha = 3.25^\circ$, the PSD, shown in Figure 5, does not show a prominent peak around $St = 0.066$ that was estimated using peak-to-peak analysis.

To remedy this, we use Short-Time Fourier Transform (STFT), where the analysis is computed for a shorter data window, in our case lasting $T_w = 30$. The start of the window is moved along the full input (by one step at a time, producing $(T - T_w)/\Delta t = 70$ PSD traces. When STFT is applied as the primary analysis tool, it is typical to further




Fig. 4 Speed profiles during transient to the buffet at $\alpha = 3.25^\circ$. (a) A time-snapshot of the velocity magnitude contours (b) Four snapshots of the velocity magnitude profiles evaluated along the blue ellipse in (a) measured anti-clockwise starting from the trailing edge. (c) Comparison of the time trace of the lift C_L (red curve) and the velocity magnitude profiles (heatmap) along the blue ellipse in (a).

modify it to minimize artifacts caused by using relatively-short processing windows. For simplicity, we do not pursue those modifications here.

In order to compare the Strouhal number spectrum at each angle, Figure 6 was made by computing STFT for 35 values of α , $\alpha \in [1.00^\circ, 7.00^\circ]$, and splicing them horizontally. Thus, the large scale features from left to right indicate the change in PSD with respect to α , while variation inside each of the smaller vertical bands indicates the change in the PSD corresponding to the movement of the data window inside that specific α simulation.

The overall distribution of PSD indicates that for buffeting α , roughly $\alpha \in [3.20^\circ, 6.30^\circ]$, the power is concentrated around $St \approx 0.1$, which is broadly consistent with peak-to-peak analysis and the experimentally reported data. The shortened window means that these peaks are known only up to the resolution of 0.04 in the St domain.

The stability of oscillation can be estimated by looking inside the vertical bands; around $\alpha \approx 3.00^\circ$ the high values of PSD appear at the start of each simulation, while later, e.g. $\alpha \approx 6.25^\circ$, the high values appear toward the end of each simulation. This is again consistent with the direct peak-to-peak analysis. However, neither Fourier analysis nor peak-to-peak analysis give a reliable way of estimating the growth or decay rates quantitatively.

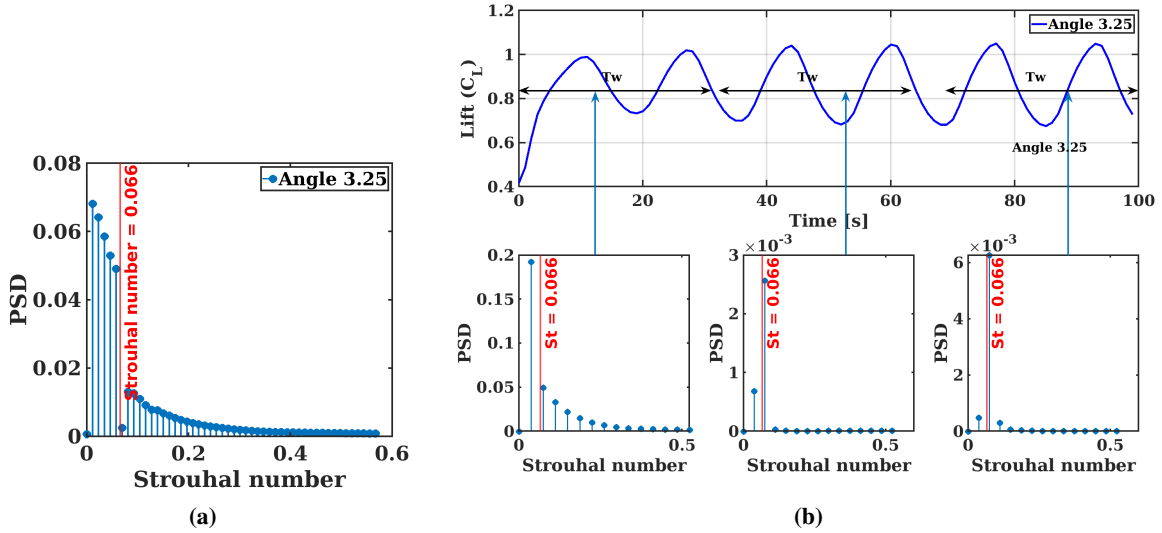


Fig. 5 Power Spectrum Density (PSD) of the lift coefficient C_L mean-removed time series for α of 3.25° : (a) applied to the full time series of duration 100, (b) applied to three segments of duration 30, indicated in the time trace. Notice that the PSD of the full data and of the initial window overlapping the transient do not have a peak at the dominant frequency of oscillation.

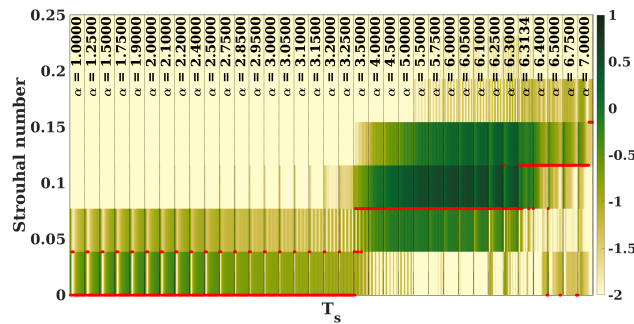


Fig. 6 Heat map of the Power Spectral Density (PSD) of the lift coefficient C_L for various α and for various locations of the data window (duration 30 units). Each vertical strip corresponds to the moving of the data window within the simulation for the indicated α . The red dots indicate the peak of each PSD. Due to the low number of data points within each window, the FFT has a low resolution in the frequency axis, resulting in a pixelated appearance.

III. Sliding-window Dynamic Mode Decomposition (DMD) of transient data

A. Overview of the DMD algorithm

The input into DMD is a spatiotemporal matrix \mathbf{X} whose columns $\mathbf{X}[n]$ represent simulation snapshots, that is evaluations of state variables on a fixed grid containing 34 521 nodes. The variables used are arranged as follows

$$\mathbf{X} = \begin{bmatrix} u/\sqrt{2} \\ v/\sqrt{2} \\ s/\sqrt{2} \\ c_s \end{bmatrix}, \quad (1)$$

where the symbols represent the velocity components (u, v) , the speed $s = \sqrt{u^2 + v^2}$, and the local speed of sound $c_s = \sqrt{\gamma p / \rho}$ computed from the local pressure p and density ρ . The speed s is used explicitly to allow for a direct interpretation of spatial profiles arising from DMD modes, without additional nonlinear transformations in the post processing. Using $[u, v, p]^T$ as the columns of the data matrix would lead to the dimensional incompatibility of the euclidean norm of the columns and the physical energy of the system. For this reason, we use c_s instead of the pressure p and include the scaling factors $\sqrt{2}$. Since the regression below is performed with respect to the euclidean ℓ^2 norm, the scaling of numerical tolerances and errors is more closely connected to the physical energy of the structures that are (numerically) discarded when forming reduced order models of data.

Consequently, the size of each column is $138\,084 = 34\,521 \times 4$ and there are $N = 101$ columns $\mathbf{X}[n]$, $n = 0, 1, \dots, N-1$ representing time instances between 0 to 100, sampled at timestep $\Delta t = 1$.

DMD is the regression of columns $\mathbf{X}[n]$ onto the formula

$$\mathbf{X}[n] = \sum_{k=1}^K \phi_k \lambda_k^n b_k, \quad (2)$$

where ϕ_k are (complex-valued) spatial profiles called DMD modes, λ_k are DMD frequencies, and b_k are coefficients used to combine DMD mode evolution into the data. Complex-valued modes and frequencies always appear as conjugate pairs, as the inputs are all real-valued. In this way, every DMD mode evolves according a single (complex) frequency, which is often written in its ‘‘continuous-time’’ form

$$\lambda_k^n = \exp\{(\sigma_k + i\omega_k)n\Delta t\}, \quad (3)$$

so that coefficients $\sigma_k + i\omega_k \in \mathbb{C}$ can be interpreted in the same way as eigenvalues or characteristic values for linear differential equations: $\sigma_k < 0$ indicate modes whose magnitude decays in time, while $\omega_k \neq 0$ indicates oscillating modes, where the period of oscillation is $2\pi/|\omega_k|$. This DMD algorithm was first introduced to analyze constituent flow patterns in complex flows and to design control and detection strategies [11, 12].

Justification for the DMD regression comes from the representation of dynamics as the evolution of the initial vector by the linear, time-invariant Koopman operator. The DMD frequencies λ_k approximate a subset of eigenvalues of the Koopman operator [15, 32], while modes ϕ_k are projections of simulated variables onto the eigenfunctions of the Koopman operator. Since the interpretation of DMD as regression into (2) suffices for our purposes, we point the interested reader to the recent review [17] for more details on the Koopman operator framework.

While the family of DMD algorithms grows every day, the early *exact* DMD algorithm and its modifications will suffice for the proof-of-principle of the proposed buffet analysis.

$\mathbf{X}[n]$ is organized into two matrices, past $\mathbf{X}_1 = (\mathbf{X}[n])_{n=0}^{N-2}$ and future $\mathbf{X}_2 = (\mathbf{X}[n])_{n=1}^{N-1}$ whose relationship is modeled as

$$\mathbf{X}_2 = \mathbf{K}\mathbf{X}_1 + \mathbf{R}. \quad (4)$$

The DMD matrix \mathbf{K} is a representation of the Koopman operator that maps each snapshot $\mathbf{X}[n]$ into the subsequent one $\mathbf{X}[n+1]$ or, equivalently, each column of the past matrix $\mathbf{X}_1[n]$ into the corresponding column of the future matrix $\mathbf{X}_2[n]$. Matrix \mathbf{R} represents an error term or a residual that is expected to be small.

In principle, \mathbf{K} is computed via a regression, by optimizing the Frobenius norm

$$\min_{\mathbf{K}, \Delta\mathbf{X}_1, \Delta\mathbf{X}_2} \left\| \begin{bmatrix} \Delta\mathbf{X}_1 \\ \Delta\mathbf{X}_2 \end{bmatrix} \right\|_F, \quad \text{subject to } \mathbf{X}_2 + \Delta\mathbf{X}_2 = \mathbf{K}(\mathbf{X}_1 + \Delta\mathbf{X}_1); \quad (5)$$

in this case the residual is $\mathbf{R} = \mathbf{K}\Delta\mathbf{X}_1 - \Delta\mathbf{X}_2$. Such setup corresponds to the total-least-squares (TLSQ) regression, reflecting that \mathbf{X}_1 and \mathbf{X}_2 are equally susceptible to errors. DMD using TLSQ was formulated and studied by Hemati et al. [21].

The matrix \mathbf{K} is not directly of interest, but rather only its eigenvalues and (right) eigenvectors, the DMD eigenvalues and modes. The DMD algorithms approximate this data by performing a regression on a low-dimensional subspace of the full state space. For the total-least-squares DMD [21] the procedure is as follows. First, define

$$\mathbf{Z} := \begin{bmatrix} \mathbf{X}_1 \\ \mathbf{X}_2 \end{bmatrix} \in \mathbb{R}^{2m \times (N-1)}, \quad (6)$$

and then perform the following steps:

- 1) Compute Singular Value Decomposition $\mathbf{Z} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*$, such that the truncation $\mathbf{Z}_k := \mathbf{U}_r\mathbf{\Sigma}_r\mathbf{V}_r^*$ is the best rank- r approximation of \mathbf{Z} ,
- 2) Project inputs onto the range of \mathbf{Z}_k^* : $\tilde{\mathbf{X}}_2 = \mathbf{X}_2\mathbb{P}_{\mathbf{Z}_k^*}$, $\tilde{\mathbf{X}}_1 = \mathbf{X}_1\mathbb{P}_{\mathbf{Z}_k^*}$, where $\mathbb{P}_{\mathbf{Z}_k^*} = \mathbf{V}_r\mathbf{V}_r^*$,
- 3) Compute SVD $\tilde{\mathbf{X}}_1 = \tilde{\mathbf{U}}\tilde{\mathbf{\Sigma}}\tilde{\mathbf{V}}^*$,
- 4) Evaluate ℓ^2 -regression in the low-dimensional subspace $\tilde{\mathbf{K}} := \tilde{\mathbf{U}}^*\tilde{\mathbf{X}}_2\tilde{\mathbf{V}}\tilde{\mathbf{\Sigma}}^{-1} \in \mathbb{R}^{r \times r}$,
- 5) Compute all eigenvalues/-vectors: $\tilde{\mathbf{K}}\tilde{\boldsymbol{\phi}}_k = \lambda_k\tilde{\boldsymbol{\phi}}_k$, and set $\sigma_k + i\omega_k := \ln \lambda_k / \Delta t$,
- 6) Compute DMD modes: $\boldsymbol{\phi} = \tilde{\mathbf{X}}_2\tilde{\mathbf{V}}\tilde{\mathbf{\Sigma}}^{-1}\tilde{\boldsymbol{\phi}}_k$ and normalize to unit norm $\boldsymbol{\phi}_k \leftarrow \boldsymbol{\phi}_k / \|\boldsymbol{\phi}_k\|_2$.

The DMD model reduction process is to choose the value of the parameter r in the first step of the DMD algorithm. However, in this research endeavor, $\tilde{\mathbf{X}}_1$ and $\tilde{\mathbf{X}}_2$ were made using all the right singular vectors.

To fit the library of DMD modes to the data, in order to obtain the decomposition in (2), the combination coefficients b_k are computed by a second regression:

$$\mathbf{b} = \arg \min_{\boldsymbol{\beta} \in \mathbb{C}^r} \sum_{n=0}^{N-1} \|\mathbf{X}[n] - \sum_{k=1}^r \boldsymbol{\phi}_k \lambda_k^n \beta_k\|_2^2 \quad (7)$$

Elements of the vector \mathbf{b} can be interpreted as (complex) contributions of DMD modes to the *initial* condition. Assigning ‘‘importance’’ to each mode based on the value of $|b_k|$ therefore biases decaying modes (that start ‘‘large’’) over unstable modes. Instead, we sort the modes according to their time-averaged ℓ^2 norm

$$\tilde{b}_k^2 := \frac{1}{T} \int_0^T \|b_k e^{\sigma_k t + i\omega_k t} \boldsymbol{\phi}_k\|_2^2 dt = \frac{\exp(2\sigma_k T) - 1}{2\sigma_k T} |b_k|^2 \quad (8)$$

$$\approx [1 + \sigma_k T + \mathcal{O}(\sigma_k^2 T^2)] |b_k|^2, \quad (9)$$

which reduces to $\tilde{b}_k = |b_k|$ for periodic ($\sigma_k = 0$) modes.

B. Stabilization of the DMD spectrum using sliding window

Even though the lift coefficient develops a clear oscillation at $\alpha \geq 3.15$, applying the DMD to the full transient data set fails to resolve the associated frequency $St \approx 0.07$ as one of the dominant DMD modes. This artifact was explained by Page and Kerswell [23] to be the consequence of fitting a single time-invariant linear model to data produced by a transient (near-heteroclinic in [23]) trajectory of the underlying dynamics, as illustrated in Figure 7, where the trajectory in gray transits from an unstable fixed point to a stable limit cycle.

Since different modes are dominant along different segments of the trajectory, we apply DMD only to a subinterval of time (*data window*) and then *slide* this subinterval from the beginning to the end of each simulation. If DMD was then applied to three different data windows, shown in Figure 7 in red, blue, and purple, three qualitatively different spectra emerge.

The red interval results in the right half-plane dominant eigenvalues, as the dynamics expands away from the unstable equilibrium. The blue interval results in the left half-plane dominant eigenvalues, as the dynamics contracts to the stable limit cycle. The purple interval, that straddles both the repelled (red) and attracted (blue) sections does not yield a numerically-stable calculation of the spectrum.

In simple analytic examples, as demonstrated in [23], the location of the *transient point* separating the repelling and attracting segments can be computed analytically, but we do not attempt to do so here. However, by reducing the size of the data window and positioning it at different places along the trajectory, it is possible to capture the dynamics in




Fig. 7 Sketch of a trajectory (gray) departing the unstable fixed point (steady flow) and approaching the stable limit cycle (buffet). Dominant DMD eigenvalues computed based on the red segment have positive real part, while DMD eigenvalues computed based on the blue part have negative real parts. DMD eigenvalues computed based on the purple segment, which straddles the separation point, are not computed reliably.

each section trajectory by computing a numerically-reproducible eigenvalue spectrum. More precisely, the size of the data window should not be too narrow, in order to correctly resolve the complex frequency of oscillation and should not be too long, so as not to straddle the transient point of the trajectory. In all subsequent results, we use the data window size of $T_W = 30$; we verified that all the results are robust for windows between $20 < T_W < 50$, but we omit this comparison in the interest of space.

In analyzed simulations, all initial conditions were taken to be the free-stream fields, which are typically not in the vicinity of the steady solution (equilibrium). Therefore, we do not expect to capture the modes that govern the departure of solutions from the unstable steady flow in the case of transonic buffet, but only those modes that govern the collapse of the solutions onto the attractor (limit cycle) and those that govern the neutrally-stable behavior along the attractor.

The simulation at $\alpha = 3.25^\circ$ was used to demonstrate how the window's placement affected the DMD eigenvalue spectrum. The data window was placed in locations between $0 - 30$, $30 - 60$, and $70 - 100$ and analyzed the DMD continuous-time eigenvalue spectrum of each location to comprehend the behavior of the dynamics, as shown in Figure 8.

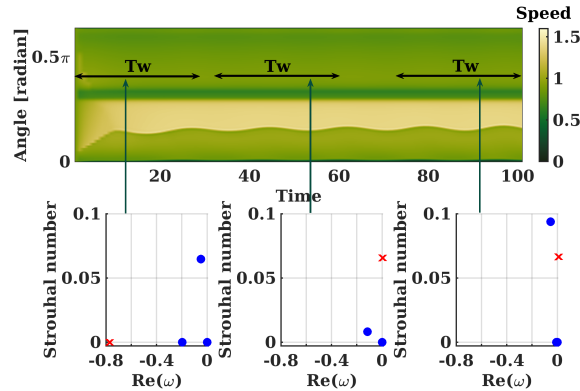


Fig. 8 Eigenvalue spectrum of three different points in the timeline for $\alpha = 3.25^\circ$. The dominant mode eigenvalue (red cross) moves closer to the imaginary axis as the time window moves away from the transient region. This indicates that the flow is transitioning from an unstable fixed point to a stable limit cycle.

When the data window is moved from left to right, as shown in Figure 8, the real component of the first dominant mode shifts from negative to positive, as indicated by the red cross, and the Strouhal number remains close to 0.06 after the first window, which is after transition time. This demonstrates that the NSU2D code's solutions for the high α of 3.25° go to a limit cycle from a stable free stream velocity.

The dominant DMD eigenvalue in the spectrum, shown in Figure 9, tracks the dominant buffet frequency closely, with the other eigenvalues harmonically related to it. Notice that in the early stages of the transient, the eigenvalue has a negative sign, corresponding to the contraction onto the attractor from the initial free-stream profile.

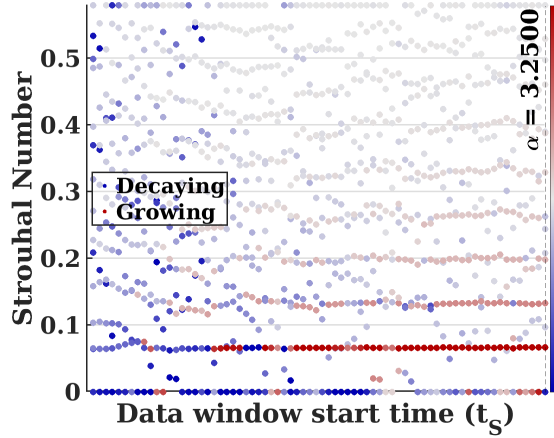


Fig. 9 The change of oscillation frequency of the ten most-dominant DMD eigenvalues with respect to the location of the data window of constant duration $T_W = 30$. Red resp. blue color corresponds to the positive, resp. negative, sign of the real part of eigenvalues; color intensity corresponds to the mean- ℓ^2 norm of the DMD modes. The dominant line around $St = 0.06$ corresponds to the dominant buffet frequency; sub-dominant frequencies are harmonically related to it.

C. The effect of window location on DMD mode profiles

The location of the data window along the transient trajectory potentially affects not only the DMD eigenvalues but also the shape of DMD modes.

Most generally, the mode vector ϕ is a complex-valued vector. The moduli of its elements $|\phi_i|$ can be interpreted as the degree to which the associated (complex) frequency is responsible for the variation in i th measurement variable (observable) in the data. To compare modes ϕ_τ and $\phi_{\tau'}$ computed for data windows starting, respectively, at τ and τ' , we compute the Pearson correlation coefficient $\text{corr}(|\phi_\tau|, |\phi_{\tau'}|)$. Pearson coefficient of two vectors is the cosine of the angle between their variations around their respective means, computed by the dot-product after centering and normalizing,

$$\text{corr}(v, w) := \frac{v - \bar{v}}{\|v - \bar{v}\|} \cdot \frac{w - \bar{w}}{\|w - \bar{w}\|}. \quad (10)$$

In what follows, for each data window associated with a simulation at a single angle α we compare the dominant non-trivial (highest \tilde{b} excluding the trivial $\lambda \neq 1$ mode). The correlation between all pairs of dominant modes chosen in that way can be visualized as a symmetric heatmap plot where the values ≈ 1 indicate profiles of high similarity, and those ≈ 0 indicate profiles that are different from each other, as shown in Figure 10. Groups of consecutive modes similar to each other appear as dark squares.

Figure 10 shows the correlation of spatial profiles of dominant DMD eigenvectors for α below the buffeting regime, in the pre-buffet regime, inside the buffet regime, and the post-buffet regime. For all angle, the initial transient period where there is no correlation between dominant modes is clearly visible; in all cases, after the transient, the dominant mode is computed robustly with respect to the choice of the data window.

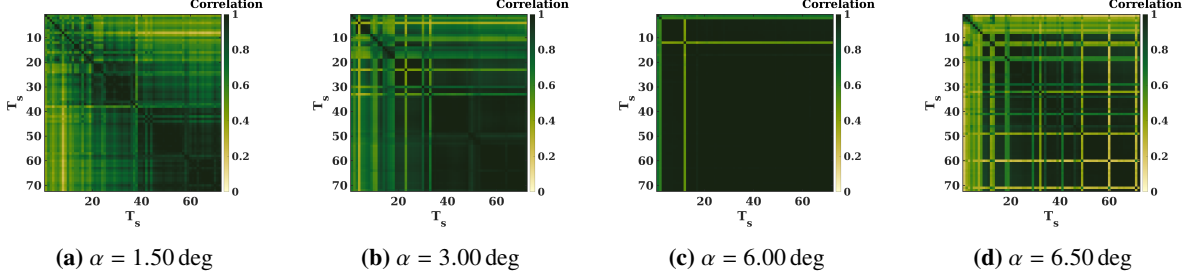


Fig. 10 Pearson correlation (10) between speed portion of the modulus of dominant DMD modes for all data windows in simulations with four values of α . Axes correspond to start times for the two windows compared. High values (dark) indicate closely-aligned modes, low values (light) indicate almost-orthogonal modes. All panels demonstrate that after the initial transient, the dominant modes remain robust with respect to the choice of the data window; regular vertical lines in $\alpha = 6.50^\circ$ panel are outliers, associated with non-buffeting profiles concentrated in the wake, which appear due to insufficiently high simulation step (aliasing).

IV. Parametric analysis of the buffet using DMD

We now move from tracking the change in DMD spectrum and modes *within* a single α simulation to tracking them *between* variations in α . In Section II.A we mentioned that the onset of buffet in airfoils can be explained by a Hopf bifurcation mechanism. If we were analyzing the spectrum and eigenvectors of the flow linearized around the steady profile (Jacobian) instead of the DMD matrix, we would expect to see a pair of oscillating eigenvalues whose stability changes from decaying to growing, and whose eigenvectors span the plane of oscillation of the trajectory. In this section we show the effect of the onset on the spectrum and eigenvalues of the DMD, as well as artifacts that arise from applying DMD to a realistic simulation dataset.

A. The change in spectrum across parameter variation

We analyze a dataset containing simulations for 35 different α in the interval $1.00^\circ \leq \alpha \leq 7.00^\circ$. Other parameters remain the same as before: all simulations were initialized to the free-stream profile, the time step is always $\Delta t = 1.00$, data window size is $T_W = 30$, and modes were sorted in importance according to the mean ℓ^2 norm (8).

First, we visualize the change in oscillation frequency of DMD modes. To simultaneously show the effect of the location of the data window within each simulation, and the effect of changing α between simulation, we extend the Figure 9 by horizontally stacking visualizations of the imaginary parts all dominant eigenvalues for the computed angle, resulting in Figure 11. For comparison, we overlay the eigenvalues over the corresponding power spectral density plots of the lift coefficient C_L computed using the STFT, and the range of frequency at which the buffet oscillates, as determined by Sartor [8].

The most prominent feature in Figure 9 is the near-continuous line of eigenvalues appearing roughly for $2.75^\circ \leq \alpha \leq 6.3134^\circ$. The stability of eigenvalues changes from decaying (negative real part, in blue) to growing (positive real part, in red) around $\alpha = 3.10^\circ$ along this line of eigenvalues, indicating that the buffet mechanism appears as a decaying mode first, before switching the stability to growing. We interpret the change in stability of DMD eigenvalues with the corresponding change in the stability of the tangent plane around the steady solution (Hopf bifurcation).

This is consistent with the behavior that the C_L time series exhibit in time traces such as Figure 3. Since the PSD reliably detects only steady-state oscillations, the PSD of C_L does not show a peak in the Sartor band until the DMD eigenvalues become (significantly) close to the neutral stability region. Additionally, the PSD by itself does not indicate the stability of the detected oscillations.

In Figure 12 the distribution of the real parts of DMD eigenvalues, i.e., rates of exponential growth and decay, is visualized by showing the location of dominant DMD eigenvalues (left panel) and by showing the distribution of mean ℓ^2 norm across the DMD eigenvalues (right panel).

To visualize the distribution of the mean ℓ^2 norm, we perform a kernel density estimate (KDE) of the eigenvalues weighted by the mean ℓ^2 norm squared. The resulting distribution is given by

$$p_h(s) \propto \sum_k \tilde{b}_k^2 \exp \frac{-|s - \sigma_k|^2}{h^2}, \quad (11)$$

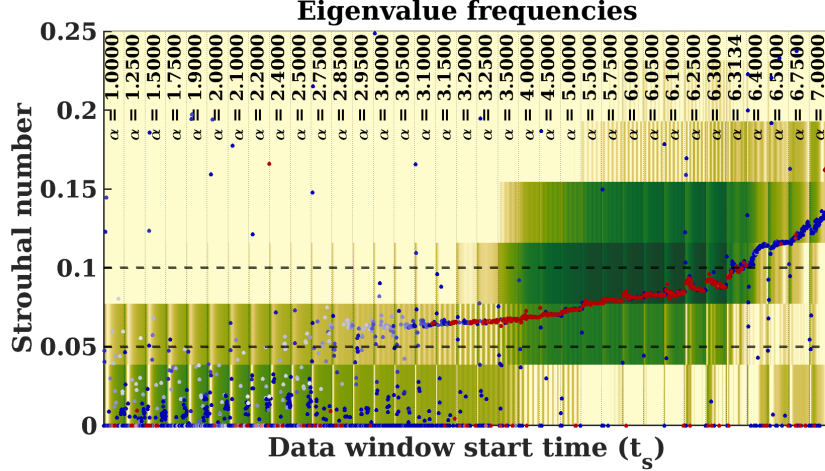


Fig. 11 Oscillation frequency (imaginary part) of the first dominant non-trivial (non-constant) DMD mode, for 35 simulations with varying α , and all data windows within each of those simulations. The color indicates the stability (sign of the real part), that is growing (positive sign) or decaying (negative sign). For comparison, the horizontal lines indicate the buffeting frequency band $0.05 < St < 0.10$, as determined by Sartor [8], and the heatmap is the value of PSD of the lift coefficient C_L calculated using the STFT. The Dominant mode oscillates at a frequency in Sartor’s band starting at $\alpha \approx 2.75^\circ$. The dominant mode decays for $2.75^\circ \leq \alpha \leq 3.05^\circ$, which we refer to as the *pre-buffet* regime. The dominant mode grows (weakly) for $3.10^\circ \leq \alpha \leq 6.3134^\circ$, which we refer to as the *buffet* regime. For $\alpha > 6.3134^\circ$ the dominant mode again decays.

where \tilde{b}_k are mean ℓ^2 norms (8) of eigenvalues, σ_k are real parts of eigenvalues, h is the width of the gaussian kernel, the sum ranges across all the non-trivial DMD eigenvalues, and \propto indicates that the distribution is scaled to an integral equal to 1. The kernel width h is a tuneable hyperparameter that produces a more-or-less smooth distribution; we used $h = 6 \times 10^{-3}$.

The visualization shows that within the buffet region ($3.10^\circ \leq \alpha \leq 6.3134^\circ$) the eigenvalues concentrate around the neutrally-stable zero-axis, which is consistent with the sustained oscillation. The stability in the pre-buffet region ($2.75^\circ \leq \alpha \leq 3.05^\circ$), where the *imaginary* parts of DMD eigenvalues already start assembling into a line that is continued through the buffet, is not significantly different from the stability in the region where no oscillations are present ($\alpha \leq 2.5^\circ$).

B. Tracking of modes across parameter variation

We now examine the change in the spatial fields associated with DMD modes as α is changed. In all visualizations we use heatmaps to visualize the real and imaginary parts of the speed segment of the DMD eigenvector associated with a particular mode; modes in the pre-buffet and buffet are shown in Figure 13, and post-buffet in Figure 14.

The signature of the buffet mechanism is clearly visible in the dominant modes in Figure 13 as the band perpendicular to the top of the airfoil, growing wider as α is increased. The band indicates the oscillation of the shock that separates the super- and subsonic flow; the shock line in this case stays perpendicular to the airfoil and its amplitude of oscillation grows as the angle is increased. At the same time, region of separated flow broadens. The separated flow region above the airfoil shows little change during the cycle of oscillations (both real and imaginary parts take the same “color”) in that region, while the wake around half a chordal length behind the wing starts oscillating as α increases further into the zone of buffet.

After the disappearance of the buffet ($\alpha > 6.31$ in Figure 14), the structure of the dominant mode significantly changes, as the dynamics in the wake becomes dominant. The supersonic bubble at the top of the airfoil shrinks (see also Figure 1 for a snapshot), so that the shock at its end moves closer to the leading edge. The separated flow above the airfoil develops a longitudinal oscillation which eventually gets overshadowed by the dynamics in the wake, clearly visible at $\alpha = 7.00^\circ$. Sartor et al. [8] attribute the dynamics in the wake to the roll-up of the shear layer (Kelvin–Helmholtz instability) which should have a broadband frequency spectrum. We leave for future work the analysis of this phenomenon, including the manner in which the DMD eigenvalues distribute in presence of irregular,

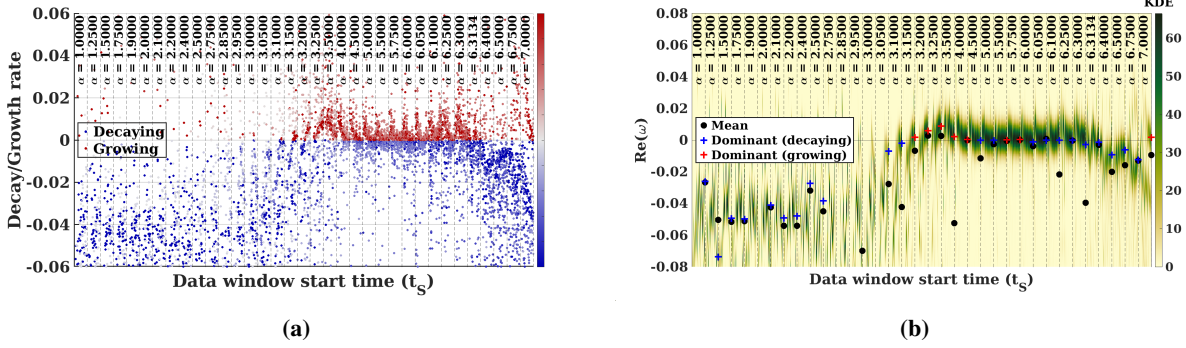


Fig. 12 Growth/decay rate (real part) of the first dominant non-trivial (non-constant) DMD mode, for 35 simulations with varying α , and all data windows of length $T_W = 30$ within each of those simulations. Panel (a) shows the real part of the dominant 10 DMD eigenvalues, excluding the trivial invariant mode, with stability indicated in color. Panel (b) shows the heatmap of the kernel density estimate (11) of the distribution of real parts of all computed non-trivial DMD eigenvalues, weighted by the mean ℓ^2 norm (8). Additionally, the mean of the KDE distribution and the first dominant DMD eigenvalue are shown for the last data window for each α . Both panels show that within the buffet region $3.10^\circ \leq \alpha \leq 6.3134^\circ$ the eigenvalues concentrate around the neutrally-stable region (weakly decaying or weakly growing).

e.g., chaotic, dynamics.

The similarity between dominant DMD modes can again be quantified using the Pearson correlation (10), displayed as the similarity matrix in Figure 15. Although the α was incremented nonuniformly (to give more resolution around the onset and end of buffet), the correlation matrix demonstrates clearly that the dominant mode changes very slowly, if at all, for angles associated with the buffet, and especially $4.00^\circ \leq \alpha \leq 6.31^\circ$.

The appearance of a dark square after $\alpha > 6.30^\circ$ indicates that the structure of the mechanism post-buffet changes discontinuously after the buffet ends. On the other hand, at the onset of the buffet ($\alpha \approx 3.10^\circ$)

The dark square for $2.40^\circ \leq \alpha \leq 3.25^\circ$ indicates that the dominant mode whose eigenvalue changes the stability around $\alpha = 3.15^\circ$, also changes minimally through the critical α , especially for $3.00^\circ \leq \alpha \leq 3.25^\circ$. This is analogous to the structure of the Jacobian of the fixed point associated with the Hopf bifurcation.

A closer inspection of the square $2.40^\circ \leq \alpha \leq 3.25^\circ$ shows that the similarity decreases along lines parallel to the diagonal, indicating that the mode change does exist through this region, but that it is continuous, i.e., modes closer in α are more similar. In contrast, changing α between $\alpha = 3.25^\circ$ and $\alpha = 3.50^\circ$ results in a significant change in the structure of the dominant DMD mode, associated with the appearance of the structure in the wake of the airfoil.

Finally, we show the relationship between dominant vectors by a principal component analysis (PCA)/multidimensional scaling (MDS) by computing the first three singular vectors of the Pearson correlation matrix for the non-trivial dominant modes in the last data window of the simulation for each α . Visualizing simulations by using these vectors as three spatial coordinates leads to the Figure 16. As it is clear, certain simulations tend to be mapped close to each other, forming clusters; we applied agglomerative hierarchical tree clustering (`clusterdata` in MATLAB 2022b) to *all* 35 PCA vectors to assign colors to each of the clusters, and distinguish structures not immediately visible from the first three vectors. Visualization using the first two PCA vectors show that there are two distinct spatial clusters associated with the dominant mode: one for approximately $1.75^\circ \leq \alpha \leq 3.5^\circ$, where the transition to buffet occurs, and one for approximately $4.00^\circ \leq \alpha \leq 6.3^\circ$, characterizing the well-developed buffet. Visualization using the second and third PCA vectors additionally shows that the internal structure of each of these clusters is ordered with respect to the change in α .

C. Harmonics, secondary mechanisms, and aliasing

Sub-dominant DMD modes at times can be associated with harmonics of the dominant oscillations, and at other times reflect additional mechanisms contributing to the phenomenon.

The Figure 17 shows the frequencies of the five most-dominant DMD modes with the integer-multiples of the most dominant mode indicated in green. This visualization shows that the sub-dominant DMD eigenvalues are exactly harmonics (integer-multiples) of the dominant mode frequency for angles between, approximately, $3.50^\circ \leq \alpha \leq 6.30^\circ$.

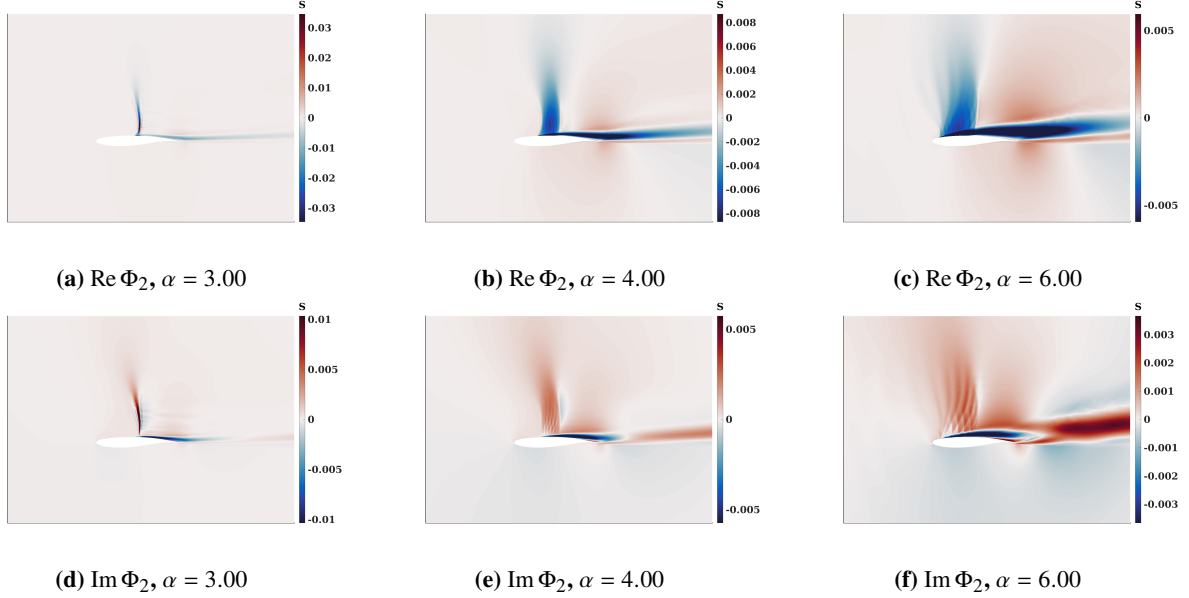


Fig. 13 The dominant DMD eigenvector for $\alpha = 3.00$ (pre-buffet), 4.00 (buffet), 6.00 (buffet), respectively in each column. The color corresponds to the speed components of the complex eigenvectors; top row is the real, and bottom the imaginary part. The widening band on the top side of the airfoil (white shape) corresponds to the path of the oscillation of the shock line separating the supersonic and subsonic flow. The widening band in the wake corresponds to the separated flow.

The structure of the speed profiles associated with the harmonic DMD modes is shown in Figure 18. The increasing number of bands perpendicular to the airfoil are consistent with the translation of the shock line during each cycle of the buffet. The structure of modes in the zone of separated flow is more complicated, starting with a longitudinal oscillation, but not indicating a pattern that can be associated with a simple rotating or translating wave.

Outside of this *strong buffet* region, the harmonics are either faint, or are not listed among the most dominant modes. For $\alpha \leq 6.50^\circ$, additional structured line of eigenvalues at $St \approx 0.15$ and $St \approx 0.3$ are visible in Figure 17. Since there are no known competing oscillation mechanisms in the given frequency range, this line of eigenvalues was suspected of being an alias of higher-frequency behavior visible around $St \approx 0.2$ as a consequence of relatively-low time resolution. Given the sampling rate of $f_s = 1/\Delta t$, the behavior that is oscillating at frequency f larger than the Nyquist frequency $f_N = f_s/2$ will appear to be oscillating at a reduced frequency $f_s - f$. Reducing the time step, from $\Delta t = 1$ to $\Delta t = 0.5$ and $\Delta t = 0.25$, as shown in Figure 19, shows that the corresponding eigenvalue lines stabilize in the range $0.8 < St < 1.2$, indicating that this is the physical range of oscillations in the wake.

To confirm that no other properties of DMD modes are affected, we compared the spatial profiles and mean ℓ^2 norms of the associated modes computed at two different time steps. As Figure 20 shows, the computed spatial profiles are largely unaffected by the change of the timestep. While this behavior is not directly related to the buffet, it offers an opportunity to conclude that even in the case of undersampled data, the DMD profiles may accurately reflect the distribution of activity in the flow even if the time signature of that activity is calculated inaccurately.

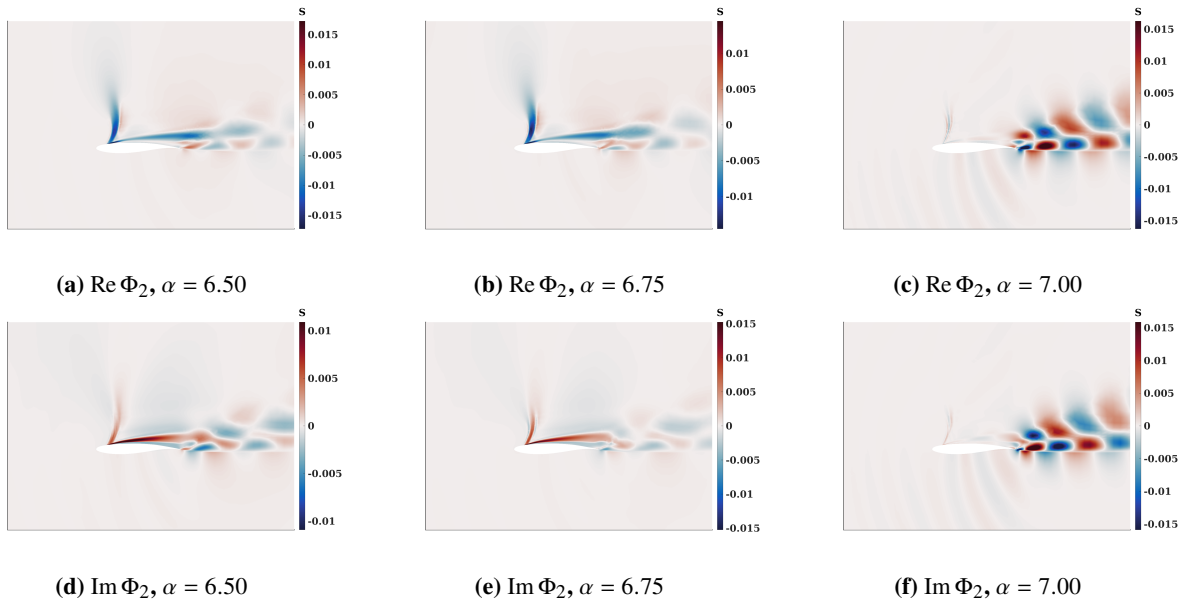


Fig. 14 The dominant DMD eigenvector for α in the region after the buffet disappears. The color corresponds to the speed components of the complex eigenvectors; top row is the real, and bottom the imaginary part. The structure in the wake corresponds to the shear layer roll-up (Kelvin–Helmholtz instability [8]).

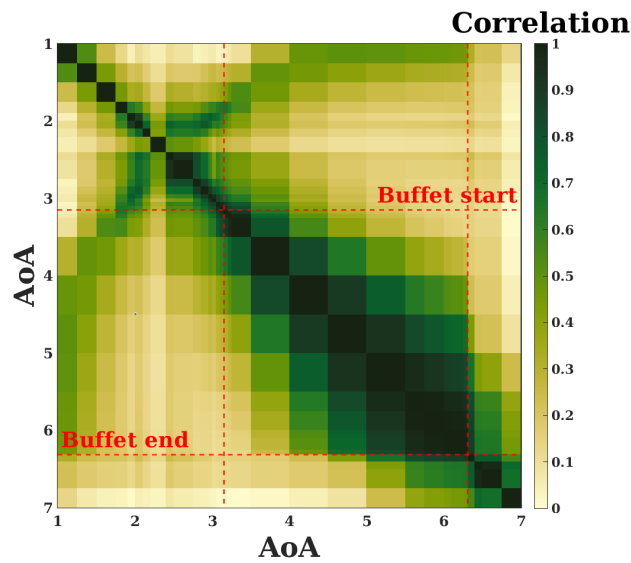


Fig. 15 Pearson correlation (10) between speed portion of the modulus of dominant DMD modes for the last data window in each simulation across the range of α . α resolution was changed non-uniformly, showing more detail at the onset and end of the buffet. High values (dark) indicate closely-aligned modes, low values (light) indicate almost-orthogonal modes. The dashed horizontal and vertical guidelines indicate the onset and disappearance of buffet, based on the DMD eigenvalue frequency analysis.

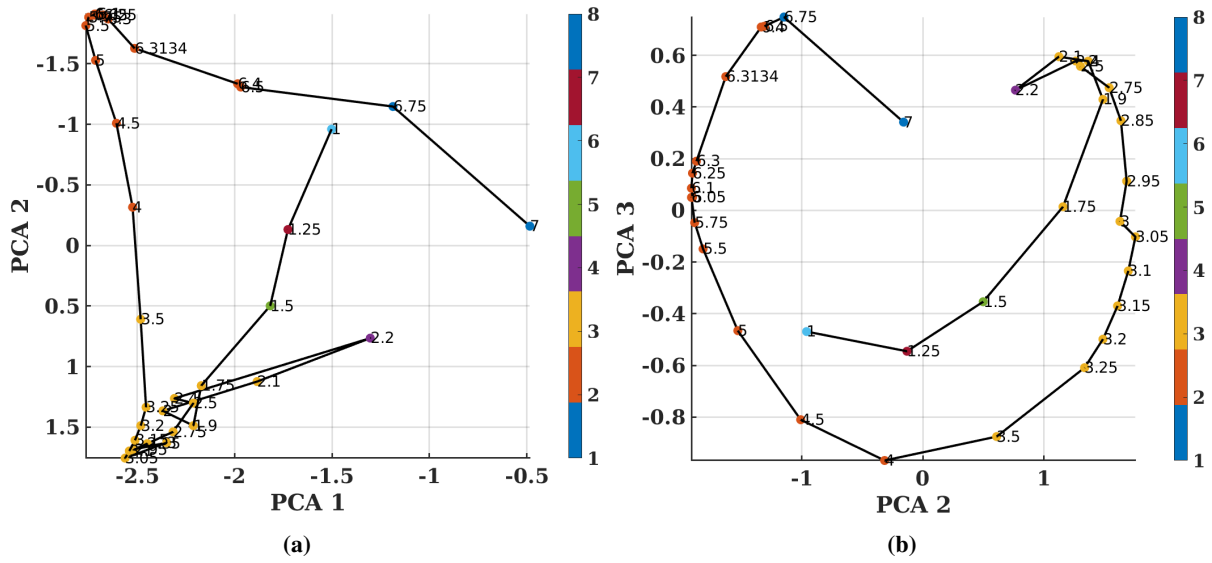


Fig. 16 Embedding of simulations for 35 values of α into the first three principal component (PCA) vectors, computed using the singular value decomposition of the Pearson correlation matrix between speed component of the dominant DMD modes for the last data window in each simulation. Labels next to points are values of α ; the lines connect points that have consecutive α s. Panel (a) is the projection onto the first and second PCA vectors, panel (b) onto the second and third PCA vectors. Colors correspond to clusters computed using hierarchical agglomerative clustering based on the geometric embedding into the all 35 PCA vectors.

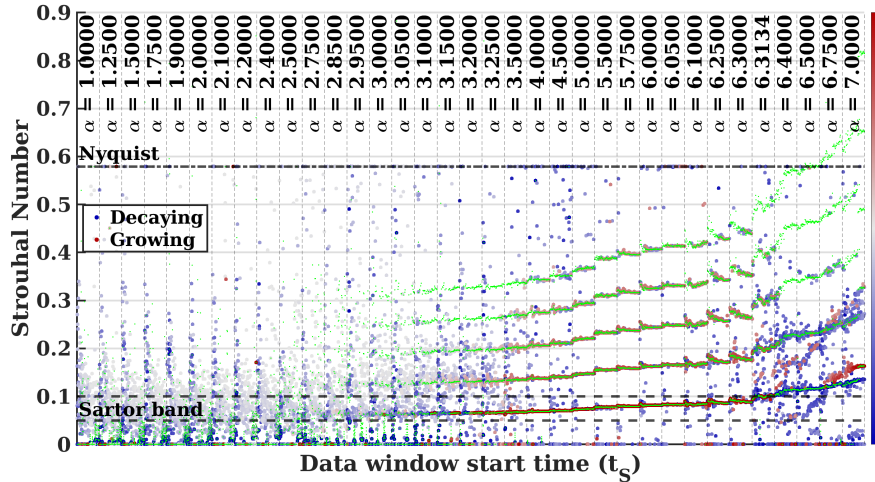


Fig. 17 Imaginary component (St number) of DMD eigenvalues for 35 simulations with varying α , and all data windows within each of those simulations. The buffet occurs approximately between $3.10^\circ \leq \alpha \leq 6.3134^\circ$. Top five dominant eigenvalues ranked by mean ℓ^2 norm (8), excluding the trivial mode at the origin, are shown for each data window (vertical slice); color indicates the sign of the real part, indicating growth (red, positive) or decay (blue, negative). The green dots are the first dominant eigenvalue along with its integer multiples, demonstrating that, during the buffet, the dominant eigenvalues are harmonically-related.

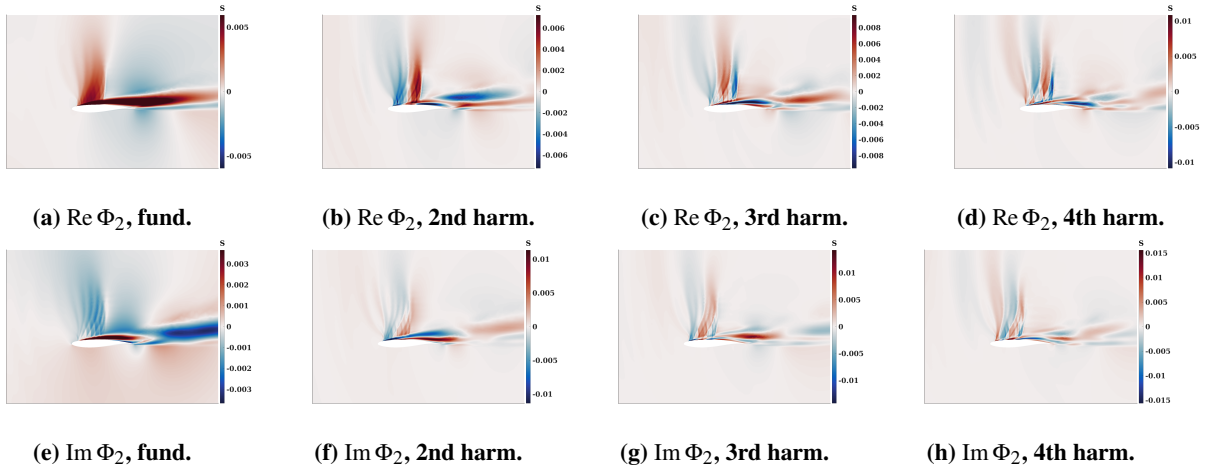


Fig. 18 The modes for the dominant DMD eigenvalue and its three harmonically-related eigenvalues for $\alpha = 6.00^\circ$. The increase in striation as the harmonic number is increased on the top of the airfoil is consistent with the shock line traveling perpendicularly to the airfoil.

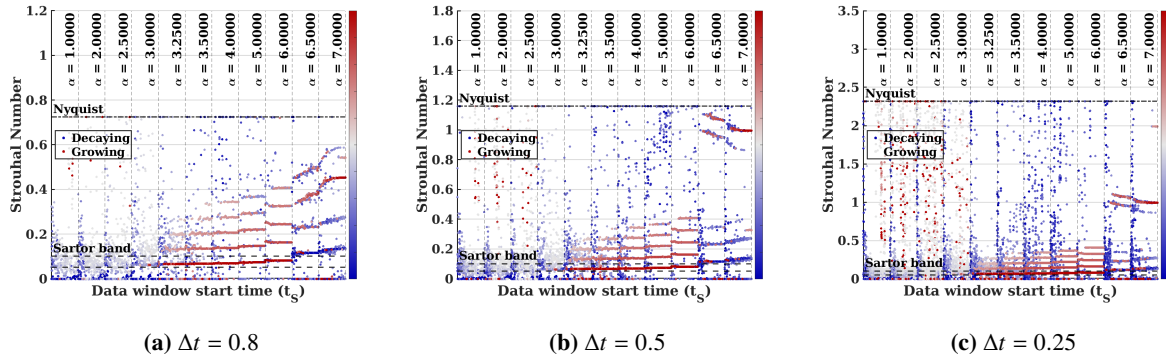


Fig. 19 Imaginary component (St number) of DMD eigenvalues for $\alpha = 1.00^\circ - 7.00^\circ$, and all data windows within each of those simulations. Notice that the lower buffeting harmonics are identified appropriately regardless of the sampling frequency. The frequencies of wake modes (which are dominant in $\alpha > 6.50^\circ$) change depending on the sampling rate, indicating that the DMD analysis for low sampling rates detects their alias mirrored across the Nyquist frequency.




Fig. 20 Spatial profiles of DMD modes associated with the wake oscillations for $\alpha = 7.00^\circ$ at time steps $\Delta t = 1.00$ and $\Delta t = 0.50$, corresponding to the pair of eigenvalue lines in the upper right corner of Figure 19(b). The pair of modes has an artificially low frequency of oscillations for $\Delta t = 1.00$ due to mirroring across the Nyquist frequency $St_N = 0.579$ (aliasing).

V. Discussion and conclusions

The goal of this paper was to clarify all the steps involved in a parametric analysis of the unsteady transonic flow using the Dynamic Mode Decomposition (DMD) algorithm. We focused on the variation of angle-of-attack (α) that leads to the development of buffet oscillations. For most α values, the dominant complex-conjugate pair of DMD modes was sufficient to identify the physical frequency of oscillations of the buffet, consistent with the published analysis of the same airfoil configuration, and demonstrate that the spatial profiles correspond to locations of the shock and separation point whose dynamical interaction leads to the buffet.

The primary lesson of the paper is that, when dealing with strongly-transient data, the windowing of the DMD is essential to obtaining the clear spectral characterization. Of course, this is not *a priori* surprising as suggested by [23]. Our results demonstrate the effect of the position of the windowed data on the DMD eigenvalues and modes, and clarify how to choose the window location within the data set. In particular, for the analyzed problem it appears that the “steady state” window can be detected by monitoring the variation in the dominant DMD eigenvector, as suggested by Figure 10.

Analyzing the variation in DMD spectrum across the data windows and across parameter variation, as described in Section IV.A, suggests several indicators for detecting a sustained oscillation in the flow: (a) the appearance of the DMD mode oscillating at the corresponding frequency, (b) the near-neutral stability of the mode, (c) the concentration of the remaining DMD eigenvalues around the neutral-stability line, (d) the appearance of the harmonically-related eigenvalues.

The dominant DMD eigenvectors can be additionally used to spatially identify mechanisms involved in the oscillation. In particular, we were able to distinguish between the two types of oscillations present in our simulated flows, the buffet and the wake oscillations, by means of the Pearson correlation (essentially, angle) between the associated spatial profiles as shown in Figure 15 and Figure 16. In particular, such quantitative comparison (as opposed to visually comparing the plots) can be used to detect finer variation between the spatial profiles as the parameter is changed and further identify sub-regimes of the phenomenon by means of clustering or geometric embedding such as multidimensional scaling or principal component analysis.

Our results additionally show that while the DMD is affected by aliasing when the snapshots are undersampled in time, the spatial profiles of the dominant oscillating DMD modes remain unchanged, even though the associated eigenvalues are mirrored across the Nyquist frequency. This is an additional argument for supplementing the DMD eigenvalue analysis with the analysis of spatial profiles of DMD modes in detecting and tracking parameter regimes.

We initially interpreted the buffet onset as a Hopf bifurcation of the steady equilibrium. At the coarse level, this connection was confirmed: DMD eigenvalues do show similar behavior to eigenvalues of the Jacobian in that their imaginary parts (frequency of oscillation) do not change through the onset, and the subspace of oscillation (dominant DMD modes) vary continuously, and not by a large amount. The more direct connections were harder to establish: the DMD is computed based on a finite segment of a non-equilibrium trajectory, and while guidance of what can be expected was given in [23], the conclusive prediction of the change in DMD eigenvalues with respect to the change in data window and the parameter is still out of reach.

Comparison of the “short-time” DMD used here with the naive short-time Fourier analysis shows that DMD computes the frequency of oscillations with higher precision than the FFT with the same number of samples, in addition to also providing the information about the stability of the oscillations. The precision of the short-time Fourier analysis could be improved using weighted windows, as is common in signal processing, but we did not pursue the effects of such weighting on DMD here.

Based on these results, we expect that DMD analysis could be used inside aircraft design optimization loops. In particular, since the spatial profile associated with the dominant DMD mode remains relatively unchanged through the onset of the buffet as the parameter is varied, it could be used as a target for design procedures. For example, a design procedure may modify the airfoil in a way that pushes the dominant DMD profile away from the profile present at the onset time, corresponding to moving the design “up” on the right branch of the spatial profile space shown in Figure 16(b). In other contexts where the buffet needs to be avoided, detecting that the design is close to the buffet region may be done by detecting decaying DMD eigenvalues of the correct frequency, i.e., the initial blue segment of the eigenvalue line in Figure 11, without ever simulating designs with parameters in the buffet region.

VI. Acknowledgments

We thank D. Mavriplis of University of Wyoming for sharing the NSU2D code with us and with his help with the setup of the code. Some of the computing for this paper was performed on Clarkson University’s ACRES cluster; we are thankful to Clarkson University and the Office of Information Technology for providing computational resources

and support that contributed to these research results. Additional computational resources for this grant were provided by the National Science Foundation under Grant No. 1925596.

References

- [1] Jacquin, L., Brion, V., Molton, P., Sipp, D., Dandois, J., Deck, S., Sartor, F., Coustols, E., and Caruana, D., “Testing in Aerodynamics Research at ONERA: The Example of the Transonic Buffet,” *Aerospace Lab*, , No. 12, 2016, pp. pages 1–11. <https://doi.org/10.12762/2016.AL12-06>.
- [2] Poplinger, L., Raveh, D. E., and Dowell, E. H., “Modal Analysis of Transonic Shock Buffet on 2D Airfoil,” *AIAA Journal*, Vol. 57, No. 7, 2019, pp. 2851–2866. <https://doi.org/10.2514/1.J057893>.
- [3] Zhao, Y., Dai, Z., Tian, Y., and Xiong, Y., “Flow Characteristics around Airfoils near Transonic Buffet Onset Conditions,” *Chinese Journal of Aeronautics*, Vol. 33, No. 5, 2020, pp. 1405–1420. <https://doi.org/10.1016/j.cja.2019.12.022>.
- [4] Dandois, J., “Experimental Study of Transonic Buffet Phenomenon on a 3D Swept Wing,” *Physics of Fluids*, Vol. 28, No. 1, 2016, p. 016101. <https://doi.org/10.1063/1.4937426>.
- [5] Crouch, J. D., Garbaruk, A., Magidov, D., and Travin, A., “Origin of Transonic Buffet on Aerofoils,” *Journal of Fluid Mechanics*, Vol. 628, 2009, pp. 357–369. <https://doi.org/10.1017/S0022112009006673>.
- [6] Crouch, J. D., Garbaruk, A., and Magidov, D., “Predicting the Onset of Flow Unsteadiness Based on Global Instability,” *Journal of Computational Physics*, Vol. 224, No. 2, 2007, pp. 924–940. <https://doi.org/10.1016/j.jcp.2006.10.035>.
- [7] Lee, B. H. K., “Self-Sustained Shock Oscillations on Airfoils at Transonic Speeds,” *Progress in Aerospace Sciences*, Vol. 37, No. 2, 2001, pp. 147–196. [https://doi.org/10.1016/S0376-0421\(01\)00003-3](https://doi.org/10.1016/S0376-0421(01)00003-3).
- [8] Sartor, F., Mettot, C., and Sipp, D., “Stability, Receptivity, and Sensitivity Analyses of Buffeting Transonic Flow over a Profile,” *AIAA Journal*, Vol. 53, No. 7, 2015, pp. 1980–1993. <https://doi.org/10.2514/1.J053588>.
- [9] Bhamidipati, K. K., Reasor, D. A., and Pasiliao, C. L., “Unstructured Grid Simulations of Transonic Shockwave-Boundary Layer Interaction-Induced Oscillations,” *22nd AIAA Computational Fluid Dynamics Conference*, American Institute of Aeronautics and Astronautics, 2015. <https://doi.org/10.2514/6.2015-2287>.
- [10] Berkooz, G., Holmes, P., and Lumley, J. L., “The Proper Orthogonal Decomposition in the Analysis of Turbulent Flows,” *Annual Review of Fluid Mechanics*, Vol. 25, No. 1, 1993, pp. 539–575. <https://doi.org/10.1146/annurev.fl.25.010193.002543>.
- [11] Schmid, P. J., “Dynamic Mode Decomposition of Numerical and Experimental Data,” *Journal of Fluid Mechanics*, Vol. 656, 2010, pp. 5–28. <https://doi.org/10.1017/S0022112010001217>.
- [12] Rowley, C. W., Mezić, I., Bagheri, S., Schlatter, P., and Henningson, D. S., “Spectral analysis of nonlinear flows,” *Journal of fluid mechanics*, Vol. 641, 2009, pp. 115–127.
- [13] Holmes, P., Lumley, J. L., and Berkooz, G., *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, Cambridge Monographs on Mechanics, Cambridge University Press, Cambridge, 1996. <https://doi.org/10.1017/CBO9780511622700>.
- [14] Budišić, M., Mohr, R., and Mezić, I., “Applied koopmanism,” *Chaos: An Interdisciplinary Journal of Nonlinear Science*, Vol. 22, No. 4, 2012, p. 047510.
- [15] Mezić, I., “Analysis of Fluid Flows via Spectral Properties of the Koopman Operator,” *Annual Review of Fluid Mechanics*, Vol. 45, No. 1, 2013, pp. 357–378. <https://doi.org/10/gcvjw4>.
- [16] Mezić, I., “Spectrum of the Koopman Operator, Spectral Expansions in Functional Spaces, and State-Space Geometry,” *Journal of Nonlinear Science*, 2019. <https://doi.org/10.1007/s00332-019-09598-5>.
- [17] Brunton, S. L., Budišić, M., Kaiser, E., and Kutz, J. N., “Modern Koopman Theory for Dynamical Systems,” *SIAM Review*, Vol. 64, No. 2, 2022, pp. 229–340. <https://doi.org/10.1137/21M1401243>.
- [18] Liu, J., and Yang, Z., “Numerical study on transonic shock oscillation suppression and buffet load alleviation for a supercritical airfoil using a microtab,” *Engineering Applications of Computational Fluid Mechanics*, Vol. 10, 2016, pp. 529 – 544.
- [19] Szubert, D., Grossi, F., Jimenez Garcia, A., Hoarau, Y., Hunt, J. C., and Braza, M., “Shock-vortex shear-layer interaction in the transonic flow around a supercritical airfoil at high Reynolds number in buffet conditions,” *Journal of Fluids and Structures*, Vol. 55, 2015, pp. 276–302. <https://doi.org/https://doi.org/10.1016/j.jfluidstructs.2015.03.005>, URL <https://www.sciencedirect.com/science/article/pii/S0889974615000602>.

- [20] Ohmichi, Y., Ishida, T., and Hashimoto, A., “Numerical investigation of transonic buffet on a three-dimensional wing using incremental mode decomposition,” *55th AIAA Aerospace Sciences Meeting*, 2017, p. 1436.
- [21] Hemati, M. S., Rowley, C. W., Deem, E. A., and Cattafesta, L. N., “De-Biasing the Dynamic Mode Decomposition for Applied Koopman Spectral Analysis of Noisy Datasets,” *Theoretical and Computational Fluid Dynamics*, Vol. 31, No. 4, 2017, pp. 349–368. <https://doi.org/10/gbqfcj>.
- [22] Kou, J., and Zhang, W., “An improved criterion to select dominant modes from dynamic mode decomposition,” *European Journal of Mechanics - B/Fluids*, Vol. 62, 2017, pp. 109–129. <https://doi.org/https://doi.org/10.1016/j.euromechflu.2016.11.015>, URL <https://www.sciencedirect.com/science/article/pii/S0997754616302990>.
- [23] Page, J., and Kerswell, R. R., “Koopman Mode Expansions between Simple Invariant Solutions,” *Journal of Fluid Mechanics*, Vol. 879, 2019, pp. 1–27. <https://doi.org/10.1017/jfm.2019.686>.
- [24] Črnjarić-Žic, N., Maćešić, S., and Mezić, I., “Koopman Operator Spectrum for Random Dynamical Systems,” *Journal of Nonlinear Science*, Vol. 30, No. 5, 2020, pp. 2007–2056. <https://doi.org/10.1007/s00332-019-09582-z>.
- [25] Maćešić, S., and Črnjarić-Žic, N., “Koopman Operator Theory for Nonautonomous and Stochastic Systems,” *The Koopman Operator in Systems and Control: Concepts, Methodologies, and Applications*, edited by A. Mauroy, I. Mezić, and Y. Susuki, Lecture Notes in Control and Information Sciences, Springer International Publishing, Cham, 2020, pp. 131–160. https://doi.org/10.1007/978-3-030-35713-9_6.
- [26] Maćešić, S., Črnjarić-Žic, N., and Mezić, I., “Koopman Operator Family Spectrum for Nonautonomous Systems,” *SIAM Journal on Applied Dynamical Systems*, Vol. 17, No. 4, 2018, pp. 2478–2515. <https://doi.org/10.1137/17M1133610>.
- [27] Brunton, B. W., Johnson, L. A., Ojemann, J. G., and Kutz, J. N., “Extracting Spatial–Temporal Coherent Patterns in Large-Scale Neural Recordings Using Dynamic Mode Decomposition,” *Journal of Neuroscience Methods*, Vol. 258, 2016, pp. 1–15. <https://doi.org/10.1016/j.jneumeth.2015.10.010>.
- [28] Nambu, T., Mavriplis, D. J., and Mani, K., “Adjoint-based Shape Optimization of High-lift Airfoil using the NSU2D Unstructured Mesh Solver,” *52nd aerospace sciences meeting*, 2014, p. 0554.
- [29] Brazell, M. J., Mavriplis, D. J., and Yang, Z., “Mesh-Resolved Airfoil Simulations Using Finite Volume and Discontinuous Galerkin Solvers,” *AIAA Journal*, Vol. 54, No. 9, 2016, pp. 2659–2670.
- [30] Valarezo, W. O., and Mavriplis, D. J., “Navier-Stokes applications to high-lift airfoil analysis,” *Journal of aircraft*, Vol. 32, No. 3, 1995, pp. 618–624.
- [31] Ahrabi, B. R., and Mavriplis, D. J., “An Implicit Block ILU Smoother for Preconditioning of Newton–Krylov Solvers with Application in High-Order Stabilized Finite-Element Methods,” *Computer Methods in Applied Mechanics and Engineering*, Vol. 358, 2020, p. 112637. <https://doi.org/10.1016/j.cma.2019.112637>.
- [32] Mezić, I., “Spectral Properties of Dynamical Systems, Model Reduction and Decompositions,” *Nonlinear Dynamics. An International Journal of Nonlinear Dynamics and Chaos in Engineering Systems*, Vol. 41, No. 1-3, 2005, pp. 309–325. <https://doi.org/10.1007/s11071-005-2824-x>.