Sensitivity analysis of limit cycle oscillations

Joshua A. Krakos*, Qiqi Wang, Steven R. Hall, David L. Darmofal

Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, 77 Massachusetts Ave. 33-207, Cambridge MA, 02139, United States

Abstract

Many unsteady problems equilibrate to periodic behavior. For these problems the sensitivity of periodic outputs to system parameters are often desired, and must be estimated from a finite time or frequency domain calculation. Sensitivities computed in the time domain over a finite time span can take excessive time to converge, or fail altogether to converge to the periodic value. We derive a theoretical basis for this error and demonstrate it using two examples: a van der Pol oscillator and vortex shedding from a low Reynolds number airfoil. We show that output windowing enables the accurate computation of periodic output sensitivities, and may allow for decreased simulation time to compute both time-averaged outputs and sensitivities. We classify two distinct window types: *long-time*, over a large, not necessarily integer number of periods; and *short-time*, over a small, integer number of periods. Finally, from these two classes we investigate several examples of window shape and demonstrate their convergence with window size and error in the period approximation, respectively.

Keywords:
sensitivity analysis, unsteady, limit cycles, periodic

1. Introduction

The sensitivity of an objective function to system parameters is an essential component of simulation as a design tool. For problems that exhibit periodic behavior, or reach a quasi-equilibrium through a limit cycle oscillation, a common objective function to optimize is the time-average of some periodic output [1–3]. Let $g(t, \beta)$ be a such a periodic function of time with dependence on a parameter β. We define the objective function, $\mathcal{J}$, as the time-average of g over a period, $T(\beta)$:

$$\mathcal{J}(\beta) \equiv \frac{1}{T(\beta)} \int_0^{T(\beta)} g(t, \beta) dt. \tag{1}$$

*Corresponding author

Email addresses: krakos@mit.edu (Joshua A. Krakos), qiqi@mit.edu (Qiqi Wang), srhall@mit.edu (Steven R. Hall), damofal@mit.edu (David L. Darmofal)

We are interested in calculating $\mathcal{J}'(\beta)$, the derivative of $\mathcal{J}$ with respect to β. A reasonable approach to compute $\mathcal{J}'$ is to approximate the periodic output over a finite time span,

$$\mathcal{J}_s(\beta, M) \equiv \frac{1}{M}\int_0^M g(t,\beta)dt, \tag{2}$$

and compute the sensitivity of this approximate function. For problems where the period of oscillation does not depend on the parameter of interest, especially for forced-period problems, this approach can be effective for both time domain methods [4, 5] and frequency domain methods [6–8]. However, for problems whose period of oscillation depends on the parameter of interest, whereas $\lim_{M\to\infty}\mathcal{J}_s = \mathcal{J}$ in general $\lim_{M\to\infty}\mathcal{J}_s' \neq \mathcal{J}'$, thus additional consideration must be taken to ensure an accurate sensitivity.

Recently, Srinath and Mittal [9] performed adjoint-based shape optimization of a free-periodic unsteady NACA 0012 airfoil at $Re = 500$ using time-averaged outputs. They observed dependence of the robustness and accuracy of their method on the extent of the integration span, M. Wilkins *et al.* [10] investigated the calculation of periodic sensitivities from a finite time simulation of dynamical systems. They note the presence of unbounded terms in the calculated value of the sensitivity, $\mathcal{J}_s'$, as M goes to infinity, a problem previously investigated by Tomovic [11], among others, as well as the error induced by fixed boundary conditions of the unsteady problem (both initial and terminal). Wilkins proposed a method to compute various parameter sensitivities by isolating the sensitivities of the period, the phase, and amplitude, then applied it to a mammalian circadian clock model with 73 states and 231 parameters.

To enable the calculation of accurate periodic sensitivities, we propose a method that is a simple modification of an existing unsteady tangent or adjoint sensitivity implementation. Namely, in the computation of the time-average, we apply a time-dependent weighting *window* to the instantaneous output:

$$\mathcal{J}_w(\beta, M) \equiv \frac{1}{M}\int_0^M g(t,\beta)\, w(t/M)\, dt, \tag{3}$$

where $w(\tau)$ is a time-dependent weighting function with $w(\tau) = 0\, \forall \tau \notin (0,1)$. Similar approaches have been used in the past to improve the output behavior [12] by alleviating phase dependence caused by the sharp endpoints of a straightforward time-average with $w = w_s = 1$, also known as a square or rectangular window. We believe our work is the first use of output windowing to improve the behavior of output sensitivity calculations.

The remainder of the paper is as follows. First, the mathematical foundation for the failure of square output windows is described. Second, drawing from the field of signal processing, *long-time* windows are examined, and window properties that improve output and sensitivity convergence are determined. Next, using the theoretical basis for the error encountered in sensitivity calculation, *short-time* windows are developed to explicitly remove the error terms for a small, integer number of periods, and their convergence with error in approximation to the period is characterized. Finally, the periodic sensitivity is then computed using these windows for two model problems: an ODE van der Pol oscillator and the DG discretization of laminar vortex shedding over a NACA 0012 airfoil.

2. Mathematical formulation

To investigate the behavior of the time-average output and sensitivity, it is useful to define a function h such that

$$h(\tau,\beta) = g(\tau\, T(\beta),\beta), \quad g(t,\beta) = h(t/T(\beta),\beta), \tag{4}$$

then $h(\tau,\beta)$ has a period of 1, independent of β, at the expense of β dependence in both arguments when integrating over t. Because h is periodic with period 1,

$$\int_0^1 h_\tau(\tau,\beta)\, d\tau = 0. \tag{5}$$

With h, the objective function can be represented as

$$\mathcal{J}(\beta) = \int_0^1 h(\tau,\beta)\, d\tau \tag{6}$$

and its derivative as

$$\mathcal{J}'(\beta) = \int_0^1 h_\beta(\tau,\beta)\, d\tau \tag{7}$$

2.1. Breakdown of square windowing

The square window time-average, $\mathcal{J}_s(\beta,M)$, in (2) can be written as

$$\mathcal{J}_s(\beta,M) \;=\; \frac{1}{M}\int_0^{NT} g(t,\beta)\, dt - \frac{1}{M}\int_M^{NT} g(t,\beta)\, dt \tag{8}$$

where hereafter $N \equiv \lceil M/T \rceil$. The error in $\mathcal{J}_s$ is then

$$|\mathcal{J}_s(\beta,M) - \mathcal{J}(\beta)| \;\leq\; \frac{T}{M}\max|g| \tag{9}$$

which converges at $\mathcal{O}(M^{-1})$ as M increases. If M is an integer multiple of the T, then $\mathcal{J}_s(\beta,M) = \mathcal{J}(\beta)$ for any N. The sensitivity of $\mathcal{J}_s$ is

$$\mathcal{J}_s'(\beta,M) = \frac{1}{M}\int_0^M g_\beta(t,\beta) dt. \tag{10}$$

Since

$$g_\beta(t,\beta) = \frac{d}{d\beta}h(t/T(\beta),\beta) = -\frac{t\,T'}{T^2}h_\tau(t/T,\beta) + h_\beta(t/T,\beta), \tag{11}$$

integration by parts on the first term yields

$$\mathcal{J}_s'(\beta,M) \;=\; \frac{1}{M}\int_0^M h_\beta(t/T,\beta)\, dt + \frac{T'}{T}\left[\mathcal{J}_s(\beta,M) - h(M/T,\beta)\right]. \tag{12}$$

In the limit as $M \to \infty$, the error in $\mathcal{J}_s'$ is given by

$$\lim_{M\to\infty} \left(\mathcal{J}_s'(\beta, M) - \mathcal{J}'(\beta)\right) = \frac{T'}{T}\left[\mathcal{J}(\beta) - h(M/T, \beta)\right]. \tag{13}$$

In general, the deviation of the instantaneous output from its time-average is nonzero and for problems with nonzero dependence of the period on the parameter, $T'(\beta) \neq 0$, the error is nonzero. Otherwise, the approximation converges to the periodic value, and is exact for M an integer multiple of T.

Transformed into the frequency domain, a periodic function can be expressed as a combination of dirac delta functions at $\omega_0 = 0$ (the time average) and at harmonics of the period at $\omega_i = 2\pi i/T \,\forall i \in \mathbb{N}$ (unsteady terms). An unweighted time average over a fixed span is equivalent to convolving the output, g, with a square window and dividing by the window span, M. To do this, the Fourier transform of the output is multiplied by the Fourier transform of a square function, i.e., a *sinc* function:

$$\widehat{w}_s(t/M) = \frac{M}{\sqrt{2\pi}} sinc\left(\frac{M\omega}{2\pi}\right) = \begin{cases} \frac{1}{\omega}\sqrt{\frac{2}{\pi}} sin\left(\frac{M\omega}{2}\right) & \omega \neq 0 \\ M/\sqrt{2\pi} & \omega = 0 \end{cases} \tag{14}$$

where the hat, $\widehat{(\cdot)}$, signifies the Fourier transform. When the square window spans integer periods, the delta functions of the periodic output signal correspond to zeros of the *sinc* function and only the time-average contribution remains. However, when a parameter of the dynamical system changes the period of oscillation, the location of the harmonics "slide" outwards (inwards) as the period is decreased (increased), see Figure (1). The

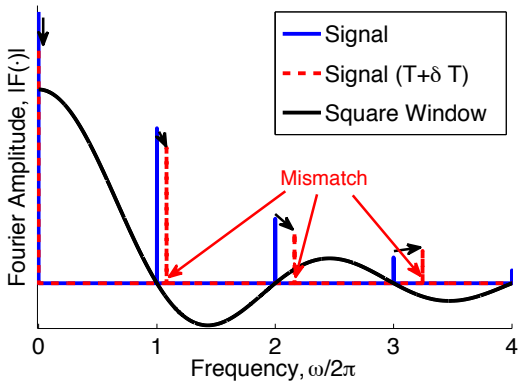

Figure 1: Magnitude of Fourier transform of square window and a periodic signal ($T = 1$)

output sensitivity, $\mathcal{J}'$, is given by the inverse Fourier transform of the derivative of the ω_0 component of the periodic output with respect to β. Convolution of the output with the square window, w_s, includes contributions from the harmonics, ω_i. At the harmonics of the output, the value of w_s is zero, but the derivative of the *sinc* function in (14) with respect to ω is nonzero. This causes the derivative of the square-window convolved output, $d(\widehat{w}_s(\omega_i)\widehat{g}(\omega_i))/d\omega$, to also be nonzero. Through $dT/d\beta \neq 0 \to d\omega_i/d\beta \neq 0$ and the chain rule, this results in a nonzero and nondiminishing erroneous contribution to $\mathcal{J}_s'(\beta)$.

2.2. Windowing

The nonconvergence of the time-average sensitivity can be addressed by applying a time-varying weighting to the instantaneous objective function. This weighting is termed a *window* or *apodization* function and is a commonly used tool in signal processing. The windowing function satisfies

$$w(\tau) = 0 \text{ for } \tau \notin (0,1), \quad \int_0^1 w(\tau)\, d\tau = 1. \tag{15}$$

The definition of the window function outside $(0,1)$ is necessary when M is a noninteger multiple of T. By decomposing the integral in (3) into a summation over multiple periods, the error in $\mathcal{J}_w$ can be expressed as

$$\mathcal{J}_w(\beta, M) - \mathcal{J}(\beta) = \frac{1}{T}\int_0^T g(\theta,\beta)\mathcal{A}(\theta, M) d\theta \tag{16}$$

where

$$\mathcal{A}(\theta, M) = \frac{T}{M}\sum_{i=0}^{N-1} w\left(\frac{\theta + iT}{M}\right) - 1. \tag{17}$$

with $\theta \in [0,1]$. Recall $N \equiv \lceil M/T \rceil$ which allows $t = \theta + (N-1)T > M$. Because $w(\tau) = 0$ for $\tau \notin (0,1)$, the integral is unchanged. Slightly modifying (12), the sensitivity can be approximated with

$$\begin{aligned} \mathcal{J}'_w(\beta, M) &= \frac{1}{M}\int_0^M g_\beta(t,\beta)\, w(t/M)\, dt \\ &= \frac{1}{M}\int_0^M \left\{ -\frac{t\,T'}{T^2} h_\tau(t/T,\beta)\, w(t/M) + h_\beta(t/T,\beta)\, w(t/M) \right\} dt. \end{aligned} \tag{18}$$

Introducing

$$\mathcal{B}(\theta, M) = \sum_{i=0}^{N-1} \frac{\theta + iT}{M} w\left(\frac{\theta + iT}{M}\right) - C(N) \tag{19}$$

with $C(N) \in \mathbb{R}$ arbitrary due to (5), the error in the sensitivity can now be expressed as

$$\mathcal{J}'_w(\beta, M) - \mathcal{J}'(\beta) = \frac{1}{T}\int_0^T h_\beta(\theta/T,\beta)\mathcal{A}(\theta, M) d\theta - \frac{T'}{T^2}\int_0^T h_\tau(\theta/T,\beta)\mathcal{B}(\theta, M) d\theta. \tag{20}$$

Theorem 1.

$$|\mathcal{J}_w(\beta, M) - \mathcal{J}(\beta)| \leq \max|h| \max|\mathcal{A}| \tag{21}$$

$$|\mathcal{J}'_w(\beta, M) - \mathcal{J}'(\beta)| \leq \frac{T'}{T}\max|h_\tau| \max|\mathcal{B}| + \max|h_\beta| \max|\mathcal{A}| \tag{22}$$

Proof. The proof is trivial by bounding the integrals of (16) and (20) with the maxima of their integrands.

The benefit of non-square windows can be further clarified by examining their effect in the frequency domain. Using windowing functions other than a square window allows for two improvements:

1. increased rate at which the Fourier transform of the window converges to zero with increasing frequency
2. increased smoothness near the harmonic frequencies of the output, specifically, increasing the order of lowest order nonzero derivative at the harmonics.

The Fourier transform of the square window converges to zero with $\mathcal{O}(1/\omega)$ as ω increases and has nonzero derivatives at its zeros. Long-time windows specifically address improvement (1) to speed convergence as the number of periods spanned increases, with no attempt to match the period of the output. Short-time windows do attempt to match the period for a small number of periods and are designed to address improvement (2) when the period used to design the window is approximate.

2.3. Long-time windowing

We propose *long-time* windows where M is allowed to become large relative to T. A benefit of taking a long-time window is that there is no need to accurately determine the period.

The convergence of long-time windows with the span of the window is governed by the convergence of $\mathcal{A}$ and $\mathcal{B}$ to the integrals they approximate, which in turn is driven by the smoothness of the function $w(\tau)$ in the interior, $\tau \in (0,1)$, and at the boundaries $\tau = 0$ and $\tau = 1$. The convergence analysis using Fourier transforms is more straightforward over an integer, N, number of periods, so a *periodic extension* of w is defined to fill in the remainder of the period outside of the span, i.e., for $M < t \leq NT$.

Theorem 2. *For a periodic function v of class C^ℓ on the interval $[0,1]$, the trapezoidal rule with N intervals, $I_N(v)$, will converge at*

$$I_N(v) = \int_0^1 v(\tau)\, d\tau + \mathcal{O}(N^{-p}) \tag{23}$$

where

$$p = \begin{cases} 1 & \ell = -1 \\ \ell + 1 & \ell \geq 0, \ell \text{ even} \\ \ell + 2 & \ell > 0, \ell \text{ odd} \end{cases} \tag{24}$$

Proof. The convergence rate of trapezoidal rule quadrature of a periodic function can be determined using the convergence of the Fourier coefficents for a periodic function, see Boyd [13], chapter 2. Note that the convergence rate for an infinitely smooth (analytic or not) periodic function is exponential.

Theorem 3. *For a window $w(\tau)$ of class C^ℓ and with $N = \lceil M/T \rceil$*

$$|\mathcal{J}_w(\beta, M) - \mathcal{J}(\beta)| \quad \leq \quad \max|h| \mathcal{O}(N^{-p}) \tag{25}$$

with p defined by (24)

PROOF. The sum in $\mathcal{A}(\theta, M)$ is a trapezoidal rule quadrature of $w(\tau)$ for $\tau \in (0,1)$. In order to integrate over noninteger numbers of periods with $M \neq NT$, recall from the requirements of a generic window that $w(\tau) = 0 \, \forall \tau \notin (0,1)$. If we treat the extended $w(\tau)$ as periodic with period NT, the rate of convergence of the trapezoidal rule numerical quadrature is given by Theorem (2) and depends on the smoothness of the extended $w(\tau)$. In the asymptotic limit of large M with $\theta \in (0,T)$:

$$\mathcal{A}(\theta, M) = \frac{T}{M} \sum_{i=0}^{N-1} w\left(\frac{\theta + iT}{M}\right) - 1 = \int_0^1 w(\tau)\, d\tau + \mathcal{O}(N^{-p}) - 1 = \mathcal{O}(N^{-p}) \tag{26}$$

Combining this result with (21), proves the theorem.

Theorem 4.

$$|\mathcal{J}'_w(\beta, M) - \mathcal{J}'(\beta)| \quad \leq \quad \left|\frac{T'}{T}\right| \max|h_\tau| \, \mathcal{O}(N^{-(p-1)}) + \max|h_\beta| \, \mathcal{O}(N^{-p}) \tag{27}$$

PROOF. As with $\mathcal{A}(\theta, M)$, the sum in $\mathcal{B}(\theta, M)$ is a periodic trapezoidal quadrature, this time of $\tau\, w(\tau)$ for $\tau \in (0,1)$. For large M (and hence large N) with $\theta \in (0,T)$:

$$\mathcal{B}(\theta, M) \quad = \quad \sum_{i=0}^{N-1} \frac{\theta + iT}{M} w\left(\frac{\theta + iT}{M}\right) - C(N) \tag{28}$$

$$\quad = \quad N\left[\int_0^1 \tau\, w(\tau)\, d\tau + \mathcal{O}(N^{-p})\right] - C(N) \tag{29}$$

The arbitrary C can be selected to cancel the integral, giving

$$\mathcal{B}(\theta, M) \quad = \quad \mathcal{O}(N^{-(p-1)}). \tag{30}$$

Combining the result with (22) gives the first expression on the righthand side of (27). Combining the derivative error (20) with theorem (3) gives the second term on the righthand side of (28), and the theorem is proven.

Possible choices of $w(\tau)$ conforming to (15) for use with the windowed objective function (3) include the Hann window ($C^1 \to p = 3$)

$$w(\tau) = 1 - \cos 2\pi\tau \quad \tau \in (0,1) \tag{31}$$

the Hann-squared window ($C^3 \to p = 5$)

$$w(\tau) = \frac{2}{3}\left(1 - \cos 2\pi\tau\right)^2 \quad \tau \in (0,1) \tag{32}$$

and the bump window ($C^\infty \to p = \infty$)

$$w(\tau) = \frac{1}{A} e^{-1/(\tau - \tau^2)} \quad \tau \in (0,1) \tag{33}$$

where A is the appropriate area under the window and infinite p gives exponential convergence. These windows were selected as they demonstrate differing convergence rates,

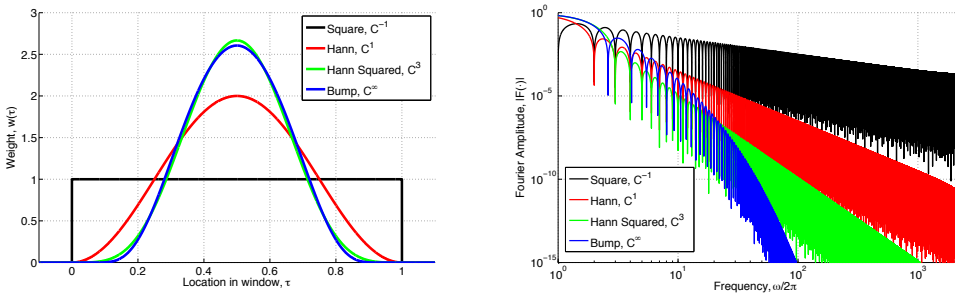

(a) Time domain (b) Frequency domain ($|F(\cdot)| = \widehat{w}(\omega)$)

Figure 2: Long-time windows

all have compact support, and are zero at their endpoints. This facilitates the use of the periodic extension for $M \neq NT$, otherwise there is a C^0 discontinuity at the endpoints as part of the extension. Examination of long-time windows using the frequency domain in Figure (2b) shows the magnitude of $\widehat{w}(\omega)$ decreasing at the expected rates with increasing ω. As the window span is increased, the frequencies that correspond to the harmonics of the output are moved proportionally to the right, are attenuated more, and the errors in the windowed time-average output and sensitivities are decreased accordingly.

The output window of Barth [12] is a modified square window, with one-period sized C^∞ caps on both ends. The span of the window is modified by increasing the span of the enclosed square window. This window is C^∞ smooth, but the convergence rates previously derived are not applicable because the Barth window is not self-similar with respect to τ. Thus, the limits employed in (26) and (28) behave differently and the convergence rate reduces to that of the square window (the limit behavior most similar to the Barth window), albeit with a lower constant factor than the square window.

2.4. Short-time windowing

The second type of output window we classify as a *short-time* window, that is applied over a small, known number of periods. Short-time windows offer the possibility of decreasing the span of integration, as the output weighting is nonzero for only a small number of periods. These windows are designed such that $\mathcal{A} = \mathcal{B} = 0$. For $M = T$, $\mathcal{A} = 0$ requires $w = 1$, the square window. However, $w = 1 \Rightarrow \mathcal{B} \neq 0$, and the error in $\mathcal{J}'_w$ is nonzero. For $M = \{nT : n \in \mathbb{N}\}$, we choose symmetric window functions from period-wise polynomials of degree $n - 1$. Enforcing the conditions $\mathcal{A} = \mathcal{B} = 0$, C^{n-2} continuity between polynomials, and symmetry about $\tau = 1/2$, gives a linear system of coefficients which can be solved for the corresponding window. The resulting short-time windows up through 5 periods are shown in Figure (3a).

We note that the conditions to define the piecewise polynomials are equivalent to repeated convolution with the single-period square window, which has the Fourier transform given in (14) with $M = T$. For an n period window, the resultant Fourier transform is

$$\widehat{w_n}(\tau) = \frac{1}{A(n)} \left(\widehat{w}_s(\tau)\right)^n = \frac{1}{A(n)} \left(\frac{T}{\sqrt{2\pi}}\right)^n sinc\left(\frac{\omega T}{2\pi}\right)^n = sinc\left(\frac{\omega T}{2\pi}\right)^n \quad (34)$$

where $A(n)$ is the renormalization factor necessary to maintain unit area under the window. Demonstrated in the frequency domain in Figure (3b), short-time windows address improvement (2) by increasing the order of the lowest order nonzero derivative of the window with respect to frequency at the harmonics of the periodic signal. In the limit of $n \to \infty$, repeated convolution gives a Gaussian distribution, suggesting a possible long-time window. Because the Gaussian distribution is not compactly supported, we have chosen the similarly exponential 'bump' function for testing as a long-time window.

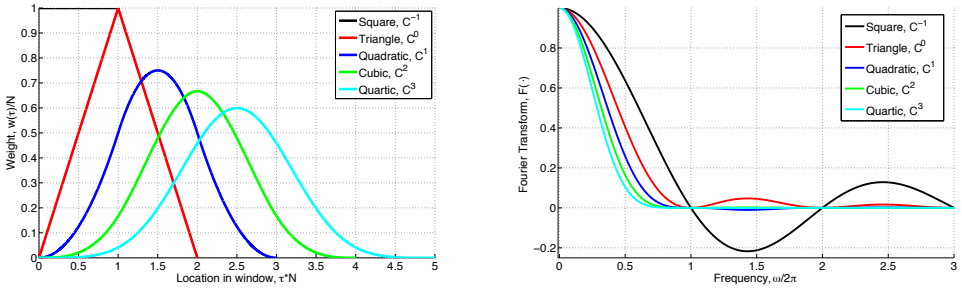

(a) Time domain (b) Frequency domain ($|F(\cdot)| = \widehat{w}(\omega)$)

Figure 3: Short-time windows

2.4.1. Triangular window ($M = 2T$)

For $n = 2$, the resulting piecewise linear function is a triangle window. The objective can be expressed as

$$\mathcal{J}_2(\beta, 2T) = \frac{1}{M} \int_0^M g(t, \beta)\, w_2(t/M)\, dt \tag{35}$$

where the triangular windowing function is

$$w_2(\tau) = \begin{cases} 4\tau & \tau \in (0, 1/2) \\ 4(1-\tau) & \tau \in (1/2, 1) \end{cases} \tag{36}$$

Consistent with the conditions of short-time windows, substituting the triangle window into (19), along with $C = 1$, gives $\mathcal{A} = \mathcal{B} = 0$ and $\mathcal{J}_2 - \mathcal{J} = \mathcal{J}_2' - \mathcal{J}' = 0$.

2.4.2. Piecewise quadratic window ($M = 3T$)

For $n = 3$, the window is piecewise quadratic, with the windowing function given by

$$w_3(\tau) = \frac{9}{2} \begin{cases} 3\tau^2 & \tau \in (0, 1/3) \\ 6\tau(1-\tau) - 1 & \tau \in (1/3, 2/3) \\ 3(1-\tau)^2 & \tau \in (2/3, 1) \end{cases} \tag{37}$$

It can be shown that the piecewise quadratic window meets the condition for zero error, $\mathcal{A} = \mathcal{B} = 0$. Note that to minimize $|\mathcal{B}|$ in (19) let $C = 3/2$.

2.4.3. Cubic polynomial windows and higher

In addition to the first two short-time windows presented, larger span windows can be devised in similar fashion. The increased time span allows for increased smoothness of the window, and decreases dependence on an accurate calculation of the period. To examine this further, demonstrations and comparisons of piecewise cubic ($n = 4$) and quartic ($n = 5$) windows are included.

2.4.4. Convergence with error in period estimate

The period may not be exactly known for short-time windows and an approximate period $T_{est} = T + \delta T$ may be used to set the window span.

Theorem 5. *For an approximate period T_{est} with period error $\delta T = T_{est} - T$ and integration over n periods, the errors in the objective function and sensitivity satisfy:*

$$\begin{aligned} |\mathcal{J}_n(\beta, n(T \pm \delta T)) - \mathcal{J}(\beta)| &\leq \mathcal{O}(\delta T^n) && (38) \\ |\mathcal{J}_n'(\beta, n(T \pm \delta T)) - \mathcal{J}'(\beta)| &\leq \mathcal{O}(\delta T^{n-1}), && (39) \end{aligned}$$

respectively.

Proof. At the estimated harmonics of the output, computed from T_{est}, the multiply-convolved window (34) has zero derivatives with respect to the frequency, up to and including the $(n-1)^{th}$ derivative, with nonzero n^{th} and higher derivatives. To determine the error in the window weight near a harmonic frequency, ω_i, take a Taylor series expansion about the estimated harmonic frequency, $\omega_{i,est} \equiv 2\pi i/T_{est}$, with $\omega_i = \omega_{i,est} + \delta\omega_i$, which reduces to:

$$\begin{aligned} \widehat{w}_n(\omega_i) - \widehat{w}_n(\omega_{i,est}) &= \frac{\widehat{w}_n^{(n)}(\omega_{i,est})}{n!} (\delta\omega_i)^n + \mathcal{O}((\delta\omega_i)^{n+1}) + \dots \\ &= \mathcal{O}((\delta\omega_i)^n) && (40) \end{aligned}$$

In the limit $\delta T \to 0$, $\delta T \propto \delta\omega$. From (40), each harmonic's component is $\mathcal{O}((\delta T)^n)$, and the total contribution to the output error over all harmonics is of the same order, directly giving (38). As previously described, the error in the output sensitivity is proportional to the error introduced into the derivative of the convolved output with respect to the frequency over all $\omega > 0$ harmonics. A Taylor series for the derivative of $d\widehat{w}(\omega_i)/d\omega$ about $\omega_{i,est}$ results in:

$$\begin{aligned} \frac{d\widehat{w}_n(\omega_i)}{d\omega} - \frac{d\widehat{w}_n(\omega_{i,est})}{d\omega} &= \frac{\widehat{w}_n^{(n)}(\omega_{i,est})}{(n-1)!} (\delta\omega_i)^{n-1} + \mathcal{O}((\delta\omega_i)^n) + \dots \\ &= \mathcal{O}((\delta\omega_i)^{n-1}) && (41) \end{aligned}$$

which, similar to the output, translates directly to (39).

3. Periodic sensitivity calculation

The periodic sensitivity may be computed using a tangent or adjoint differentiation of the primal problem. The primal problem is given by:

$$\begin{aligned} \boldsymbol{u}_t + \boldsymbol{f}(\boldsymbol{u}; \beta) &= 0 && (42) \\ \boldsymbol{u}(0) &= \boldsymbol{u}_0 \end{aligned}$$

where $\boldsymbol{u}_0$ is a given initial state. The forward-time simulation of the primal problem must initially be run long enough to equilibrate within some tolerance to periodic oscillation. Once (quasi-)equilibrium has been reached, select a state, $\boldsymbol{u}_{eq}$, on the oscillator at time t_{eq}.

The tangent problem corresponding to integrating the windowed output over a span of length M is:

$$\begin{aligned} \mathbf{v}_t + \boldsymbol{f}_\beta(\boldsymbol{u};\beta) + \boldsymbol{f}_{\boldsymbol{u}}(\boldsymbol{u};\beta)\mathbf{v} &= 0 \qquad \forall t \in [t_{eq}, t_r + M] \\ \mathbf{v}(t_{eq}) &= 0 \end{aligned} \tag{43}$$

where $\mathbf{v} \equiv \boldsymbol{u}_\beta$ is the tangent variable. The primal state is required in (43), and is typically solved for simultaneously with the tangent, using $\boldsymbol{u}_{eq}$ as a fixed initial condition for the primal problem. t_r is a relaxation time defined such that $t \in (t_{eq}, t_r)$ is sufficient time that (quasi-)periodicity can be recovered given a perturbation to the equilibrium solution at time t_{eq}. The relaxation time allows the tangent variable to equilibrate, and is required to remove the effects of the fixed initial condition, $\mathbf{v}(t_{eq}) = 0$. Rather than beginning calculation of the tangent state at time t_{eq}, the primal and tangent states may also be equilibrated simultaneously: the strict requirement is that both the primal state, $\boldsymbol{u}$, and the tangent state, $\mathbf{v}$, have reached (quasi-)equilibrium at time t_r. Once the tangent state is computed, the output sensitivity to a parameter is then:

$$\frac{d\mathcal{J}_w}{d\beta} = \int_{t_r}^{t_r+M} w((t-t_r)/M)\left\{g_\beta + g_{\boldsymbol{u}}\mathbf{v}\right\}\, dt \tag{44}$$

where $g(\boldsymbol{u};\beta)$ is the instantaneous objective function with the time dependence coming through $\boldsymbol{u} \equiv \boldsymbol{u}(t,\beta)$. If the unwindowed integrands of (44) are stored for each timestep, the window weighting may be applied as a postprocessing step, allowing computation with several different windows using the same primal and tangent solutions.

In summary, the process to apply a window to a tangent solution is:

1. Solve primal problem forward in time to quasi-equilibrium (periodic oscillation) for $\boldsymbol{u}_{eq}$;
2. Solve primal and tangent problems forward in time to t_r to equilibrate the tangent solution;
3. Continue primal and tangent problems forward in time the span of the window to $t_r + M$;
   - Simultaneously integrate the sensitivity, or
   - Store the unweighted sensitivity integrand of (44) for later sensitivity calculations.

The adjoint problem to solve in backward time is:

$$\boldsymbol{\psi}(t_r + M) = 0 \tag{45}$$

$$-\boldsymbol{\psi}_t + \boldsymbol{f}_{\boldsymbol{u}}\boldsymbol{u} = w((t-t_r)/M)g_{\boldsymbol{u}} \qquad \forall t \in [t_r, t_r + M] \tag{46}$$

$$-\boldsymbol{\psi}_t + \boldsymbol{f}_{\boldsymbol{u}}\boldsymbol{u} = 0 \qquad \forall t \in [t_{eq}, t_r] \tag{47}$$

Equation (47) is solved backward in time from t_r to allow the adjoint state to equilibrate after the forcing applied by the windowed output in (46). This is analogous to the

equilibration of the tangent state before the nonzero window can be applied. The output sensitivity to a parameter is:

$$\frac{d\mathcal{J}_w}{d\beta} = \int_{t_r}^{t_r+M} w((t-t_r)/M) g_\beta \, dt - \int_{t_0}^{t_r+M} \boldsymbol{\psi} \boldsymbol{f}_\beta \, dt$$

Unlike the tangent state, the calculation of the adjoint state is not separable from a particular window weighting. For $t < t_{eq}$, the equilibrated adjoint no longer contributes to the sensitivity calculation and need not be solved. The process to apply a window to an adjoint solution is:

1. Solve primal problem forward in time to quasi-equilibrium (periodic oscillation) for $\boldsymbol{u}_{eq}$;
2. Continue primal problem forward in time to $t_r + M$;
3. Solve adjoint backward in time from $t_r + M$ to t_r with nonzero forcing term from the windowed output, integrating the sensitivity as you go;
4. Continue adjoint backwards solve with no forcing term to allow it to equilibrate at t_{eq}, continuing to calculate the sensitivity.

The choice between a tangent and adjoint sensitivity is problem dependent, typically determined by the number of inputs relative to the number of outputs. The tangent sensitivity was used for the test cases in this paper, as it allowed the evaluation of multiple windows for a single primal and tangent forward time solution. To maintain accuracy, the time integration of the output sensitivity in (44) is carried out using the same timestepping scheme as the primal and tangent solves.

4. van der Pol oscillator

4.1. Model problem

The van der Pol oscillator is a nonlinear modification of the linear oscillator:

$$u_{tt} + \beta(u^2 - 1)u_t + u = 0 \tag{48}$$

which reduced to a first order system becomes:

$$\begin{bmatrix} u_1 \\ u_2 \end{bmatrix}_t + \begin{bmatrix} -u_2 \\ \beta(u_1^2 - 1)u_2 + u_1 \end{bmatrix} = 0 \tag{49}$$

with two states, $u_1 \equiv u$ and $u_2 \equiv u_t$. For a positive nonlinear parameter, $\beta > 0$, the oscillator exhibits a limit cycle oscillation. Moreover, the period of the limit cycle depends on the nonlinear parameter, leading to failure of square windowing calculations. The instantaneous output is given by the square of the first state, $g(t) = u(t)^2$, for which the time-average output and sensitivity with respect to β are of interested. The resulting problem is solved using a fourth order ESDIRK timestepping scheme. The sensitivity is computed with the tangent derivative method, using the tangent of the discrete primal system. The primal solution for the first 10 periods is shown in Figure (4). For simulations presented here, the initial condition is $\boldsymbol{u} = (1, 0)$, with $t_r = 50T$ to equilibrate the tangent derivative.

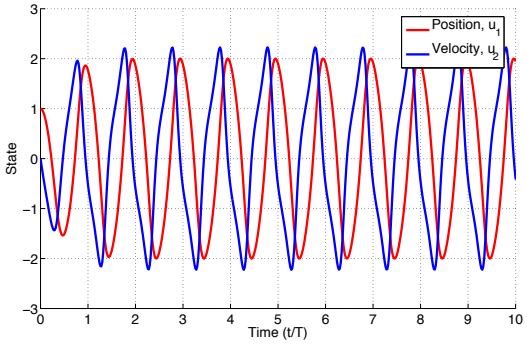

Figure 4: van der Pol oscillator primal solution ($\beta = 0.5$)

4.2. Results

For a finite-span square window, the resulting time-average slowly converges, as seen by the output error in Figure (5a). The rate of convergence of the output is shown in Figure (5b) for several long-time windows to demonstrate the improved rate of convergence of the output versus window span as the smoothness of the output window is increased. All the computed time-averaged outputs are plotted as disconnected points in the figure. Because the error in the output is highly dependent on the phase of the window, i.e. where the window overlaps with the oscillation, there can be large variations in the error with small changes in the window span. In order to show error bounds, each point used to define the solid curves in Figure (5b) is the maximum over a one period range of window sizes. Because the exact output is not known, the error is computed relative to the output with the bump window at $M = 800T$. Consistent with theory, the output error is

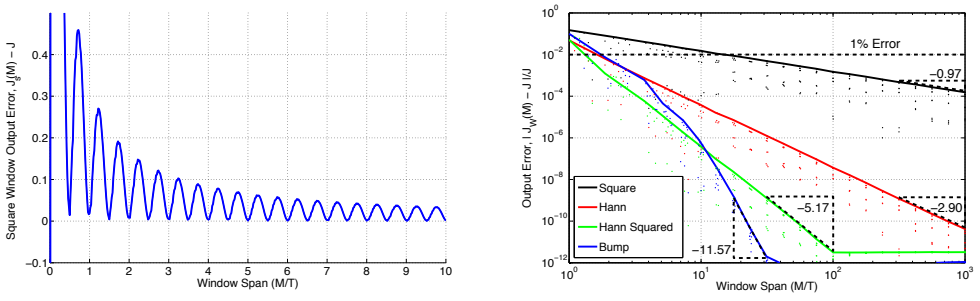

(a) Output error vs. window size (square window)

(b) Output error vs. window size

Figure 5: Output behavior vs. M for $\beta = 0.5$

convergent with increasing window span for all windows. The square window converges at $\mathcal{O}(M^{-1})$, and the Hann and Hann-squared at $\mathcal{O}(M^{-3})$ and $\mathcal{O}(M^{-5})$, respectively. Finally, the bump appears to exhibit the expected exponential convergence.

To gain insight into the behavior of the output sensitivity, we first look at the variation of $\mathcal{J}$ and $\mathcal{J}_s$ for $0 < \beta \leq 2$ in Figure (6a). the time-average smoothly increases, but the

square window with $M = 100T_{\beta=0.5}$ oscillates with approximately 1% error. Plotting the output error for three different window spans in Figure (6b) shows not only a decrease in the error magnitude, but a coincident increase in the frequency of the output error oscillation. This is caused by the sensitivity of the period to changes in β. A change in β then causes the phase of the oscillation to shift at the window span boundaries. As the window span increases, the span contains a larger number of periods of output oscillation and changes in β cause larger shifts in phase at the window endpoints. As previously demonstrated, increasing M causes the square windowed output to converge at $\mathcal{O}(M^{-1})$, but the phase shift with changing β simultaneously increases at $\mathcal{O}(M)$, resulting in a nondiminishing output sensitivity error.

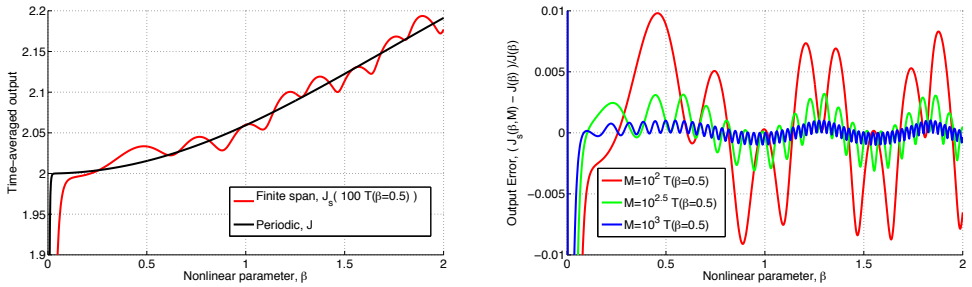

(a) Output vs. β (b) Finite span output error vs. β

Figure 6: Output behavior vs. nonlinear parameter β

The output sensitivity with respect to β was investigated in Figure (7) with the same long-time windows as for the output. For shorter span windows, the convergence of the square and Hann windows match the convergence rate of the output, i.e., the error in the sensitivity is dominated by the $\mathcal{A}$ term. For larger span (smaller error) square and Hann windows, the convergence rate decreases by one order as predicted by Theorem (4). The Hann-squared window matches the output convergence rate, but levels off due to numerical precision with the expectation that it would have otherwise had similar behavior. The convergence of the bump window for the sensitivity integration retains the same exponential shape as for the output before leveling off due to precision.

Short-time windows were investigated with the square (1 period), triangle, piecewise quadratic, cubic, and quartic windows for both positive and negative period error. Figure (8a) confirms that for an approximate period T_{est} with period error $\delta T = T_{est} - T$, the error in the output converges at the expected rate. Figure (8b) shows that the output sensitivity follows the same convergence rate as the output for $\delta T/T > 2e-3$, indicating that in this range the error is bounded by the error associated with $\mathcal{A}$. For smaller errors, the error is dominated by the $\mathcal{B}$ term and converges at the expected slower rate.

These convergence rates are highly dependent on the overlap of the timesteps with the window. In the case where the timestep does not line up with the window discontinuities, an error (dependent on the smoothness of the window) can be introduced which decreases the convergence rate. This is demonstrated in Figure (9) for timesteps corresponding to 10, 100, 1000, and 10000 steps per period and the error defined relative to the piecewise

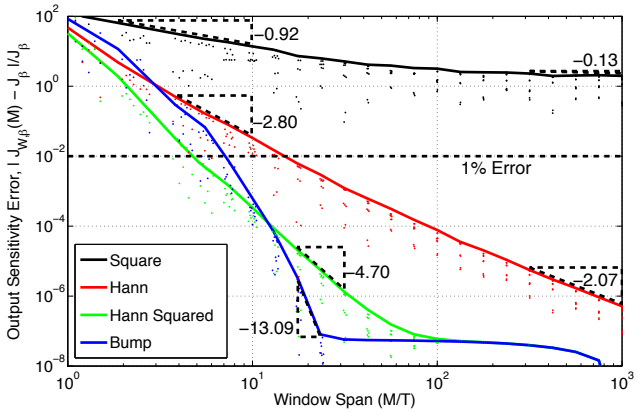

Figure 7: Output sensitivity error vs. window size

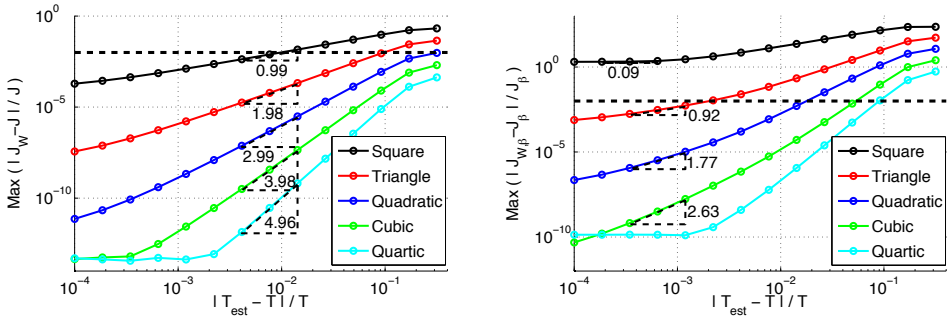

(a) Maximum output error

(b) Maximum sensitivity error

Figure 8: Short-time window time-average error vs. $|\delta T|/T$ for $\beta = 0.5$ (dash line at 1% error)

quartic window at the best calculated period for that timestep. Whereas the convergence rates with respect to window span for long time windows are relatively insensitive to the timestep, short time windows have proven to be extremely sensitive.

For many engineering applications, 1% error is an acceptable error threshold. To achieve this error level in the output requires $\approx$ 10 periods with square window and 2 or fewer for any of the smoother windows. For the sensitivity the square window levels off before achieving this level accuracy. The Hann window requires almost 15 periods; the Hann-squared, 5 periods; and the bump, 7. It is only for lower error that the bump window requires a shorter window span than the Hann-squared window. Any of the non-square short-time windows, achieve 1% error in the output with error in the period as high as 10%, and the quartic window achieves 1% error in the sensitivity with period approximation 9%. Compared to long-time windows, short-time windows offer a modest decrease in the number of periods that must be simulated to compute the output sensitivity, even using relatively inaccurate period approximations. Relative to the size of

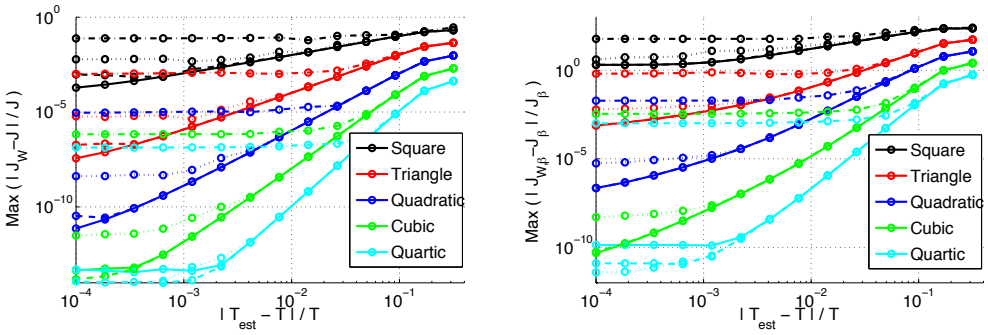

(a) Maximum output error

(b) Maximum sensitivity error

Figure 9: Timestep study of short-time window time-average error vs. $|\delta T|/T$ for $\beta = 0.5$ [$log_{10}(T/\Delta t)$: solid = 4; dash = 3; dot = 2; dash-dot = 1]

the lead-in, this savings may not justify the need to approximate the period and modify the output window accordingly.

5. Airfoil at high angle of attack

5.1. Model problem

Subsonic, laminar flow about an airfoil is used to demonstrate the impact of output windowing on sensitivity analysis of an aerodynamics application. We use a discontinuous Galerkin (DG) finite element discretization to solve the Navier-Stokes equations [14–17] Specifically, we consider the $Ma = 0.5$, $Re = 1500$ flow about a two-dimensional NACA 0012 airfoil at 9° angles of attack. This problem was chosen because the flowfield exhibits an unstable stationary point leading to cyclic unsteady behavior. In the computational mesh, the airfoil leading edge is located at $(0, 0)$ with a unit chord. The domain is square and extends $[-10, 10] \times [-10, 10]$.

The mesh refinement was determined using adjoint-based adaptation, driving the error in lift coefficient, C_L, to less than 0.0005 for a $p = 3$ stationary point solution at integer angles of attack from $\alpha = 0°$ to $\alpha = 13°$. For the problems in this work, a mesh was generated with 2158 elements, seen in Figure (10a). The resulting airfoil surface is defined by 101 nodes, connected by cubic polynomial edges ($q = 3$), with both the endpoint and edge nodes projected to the analytic geometry.

In this work, the flow and adjoint states are represented in each element by a polynomial of degree 3 ($p = 3$). For time-accurate unsteady problems, a timestep of 0.1 with units based on the chord and freestream velocity is used with the ESDIRK4 temporal discretization [18]. A grid and timestep dependence study was done for this problem and is discussed in the appendix of Krakos and Darmofal [19].

5.2. Results

The NACA 0012 problem at $\alpha = 9°$ was examined using several windows, with the goal of recreating the results of the van der Pol oscillator for the more complex system of equations, in particular Figure (7) and Figure (8). Shown in Figure (11a), the problem

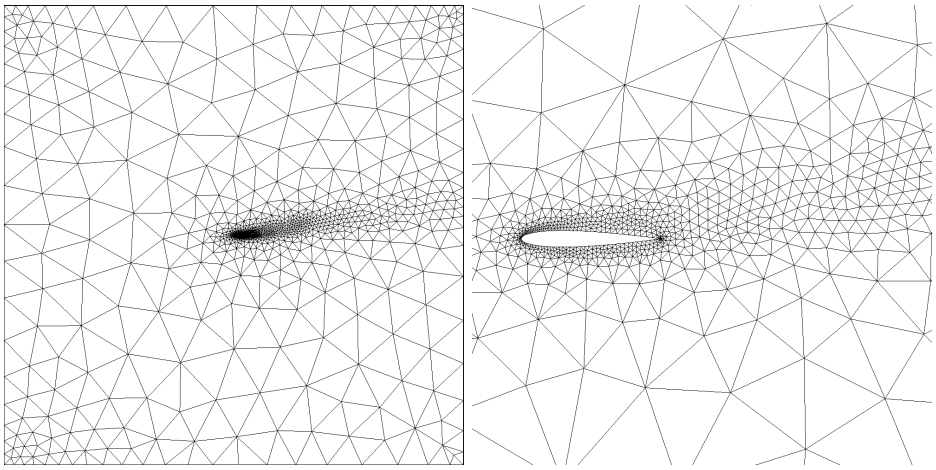

(a) Computational domain (b) Airfoil grid

Figure 10: NACA 0012 mesh

is initialized to a small perturbation of the stationary point solution and allowed to run to time $t = 200$, such that the initial transient has for the most part died out leaving an equilibrium oscillation. The windowed time-average lift, $\mathcal{J}_w$, was computed for a range of window spans from one period up to 1000 periods using a lead-in to equilibrate the tangent state of 175 periods. The period is approximated by interrogating the trough-to-trough time over approximately 90 periods. At $\alpha = 9°$, the period is approximately $T = 1.016123$. The NACA 0012 at these conditions has a non-negligible dependence of the oscillation period and the time-average lift coefficient, computed using a bump window with $M = 1000T$, on the angle of attack, see Figure (11b).

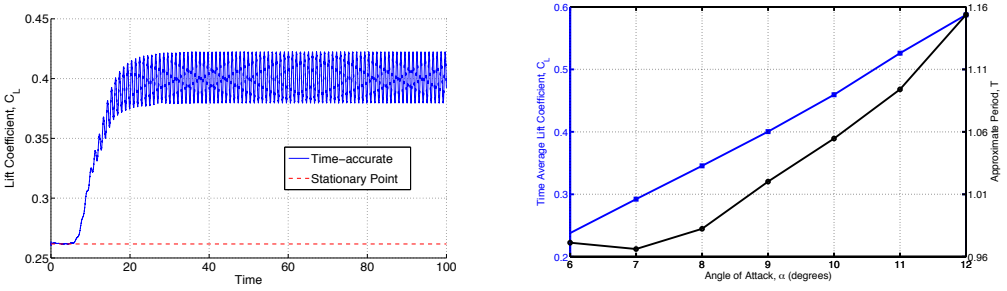

(a) Problem initialization: C_L vs. time (b) $\overline{C_L}$ & T vs. angle of attack

Figure 11: NACA 0012 - unsteady solution

Shown in Figure (12a), the output error of the square window is bounded by $\mathcal{O}(M^{-1})$. As with the van der Pol oscillator, the output weighted by the Hann window converges at $\mathcal{O}(M^{-3})$. The Hann-squared window converges at approximately $\mathcal{O}(M^{-5})$, but the convergence rate is not as uniform as the Hann window due to numerical precision. The

bump window error appears to follow an exponential trajectory through 30 periods, but levels off to converge at approximately the same rate as the Hann-squared windowed output. When the timestep is decreased to $\Delta t = 0.002$, the error for the Hann-squared and bump windows exhibit less noise and adhere more closely to the expected rates.

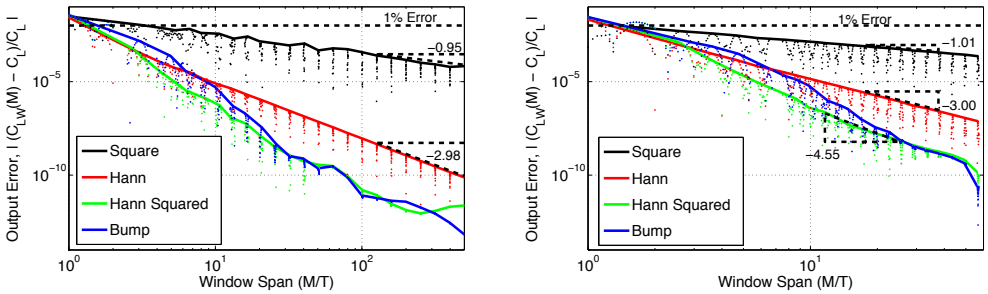

(a) $\Delta t = 0.1$ (b) $\Delta t = 0.002$

Figure 12: NACA 0012 - Output error vs. M/T (smaller max M for $\Delta t = 0.002$ due to cost)

The output sensitivity was computed using the tangent equations, and is investigated in Figure (13a). The output sensitivity error is computed relative to the sensitivity using the bump window with $M/T = 10^{2.5}$. For smaller window spans (from 1 to 10 periods), the square and Hann windows converge the output sensitivity at the same rate as for the output, but as the window spans continue to increase through 100, they both begin level off. This leveling off results in the nonconvergence of the square window and the error of the Hann window converges at $\mathcal{O}(M^{-2})$. The bump and Hann-squared windows exhibit lower error than the Hann window, but at approximately the same rate for this case. Again, when the timestep is decreased to $\Delta t = 0.002$, the error curves have smoother convergence rates with decreased noise and exhibit improved convergence rates for increased span windows compared to the larger timestep. For this problem, for both

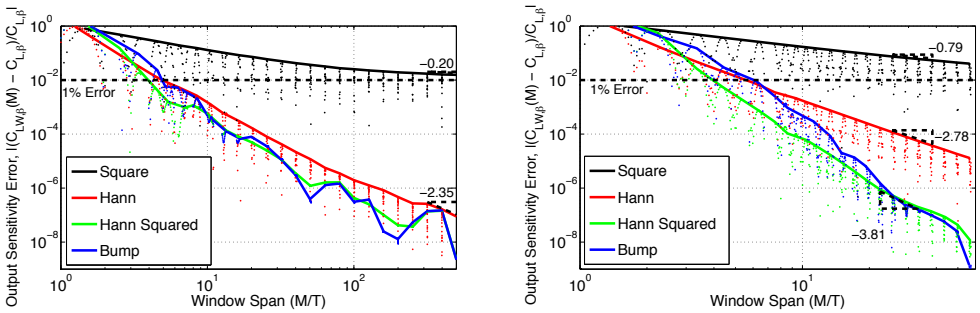

(a) $\Delta t = 0.1$ (b) $\Delta t = 0.002$

Figure 13: NACA 0012 - Output sensitivity error vs. M/T (smaller max M for $\Delta t = 0.002$ due to cost)

the output and the sensitivity to achieve a 1% relative error, any of the suggested non-

square windows will suffice with a small number of periods ($M < 10T$). The smoother windows are only necessary for tighter accuracy requirements, as the effect of increased convergence is not found until larger M.

For short-time windows, the output and output sensitivity of the NACA 0012 were calculated using the square, triangle, and piecewise quadratic, piecewise cubic, and piecewise quartic windows. For timesteps comparable to that used to compute the long-time windows ($\Delta t = 0.1$), error due to the misalignment of the timestep with the window endpoints is dominant and does not diminish with an improved period estimate. A smaller timestep is required to decrease the influence of this misalignment, and so $\Delta t = 0.002$ is used. The smaller timestep gives a more accurate solution than the large timestep used for the long-time windows, and so the $\Delta t = 0.1$ results are incompatible for determining convergence with regards to the period approximation. To correct this, the error is calculated relative to the bump window using $\Delta t = 0.002$ at ~ 60 periods, and the results are shown in Figure (14a) and Figure (14b). In the output, all five windows converge with $\mathcal{O}(\delta T^n)$ before leveling off, verifying the improved convergence rate of smoother windows. Testing with various timesteps suggests that for the square, triangle, and quadratic windows the leveling off is caused by the timestep misalignment. As with the outputs, the windowed sensitivities converge at $\mathcal{O}(\delta T^n)$ or better before leveling off, indicating that for larger period error and $\Delta t = 0.002$, the error is dominated by the $\mathcal{A}$ term. For smaller error, only the square and triangle window show the decreased convergence rate related to the $\mathcal{B}$ term. For both the output and sensitivity, the smoother windows level off at approximately the same error, but with a distinct decrease in error with increased window smoothness. Examination of the root mean square with the phase of the windows of both the output and output sensitivity show that the sensitivity to the phase levels off as well, revealing the limit of the accuracy of the simulation on the limit cycle: for the output, the flow solution; and for the sensitivity, the tangent state solution. Regardless of this limitation, all $n > 1$ short-time windows outperform the square window.

At 1% error in the period approximation, the triangle is 2 orders of magnitude more accurate than the square window, the quadratic another 2, and the cubic and quartic windows more than 6 orders of magnitude more accurate. To achieve 1% error relative error in the output, any of the $n > 1$ short-time windows will suffice with even a 10% error in period approximation. To achieve the same relative error in the output sensitivity, any of the $n > 1$ windows will suffice with 1% error in the period approximation. The quartic window gives less than 1% sensitivity relative error for 10% error in the period approximation, demonstrating the insensitivity to inaccuracy in the period approximation for the smoother, larger n windows. For this calculation, the lead-in to equilibrate the derivative variable (tangent or adjoint) is much larger than the window span ($t_r = 175 \gg 5$), thus the marginal increase in total calculation is small between a triangle or quadratic window and the longer cubic or quartic windows.

To achieve a 1% error in the time-average lift, the choice of long-time windows make little difference: all, including the square window, require less than 2 periods. For $dC_L/d\alpha$, however, the square window asymptotes to error greater than 1%. The non-square windows require between 4 and 6 periods, with the Hann-squared requiring the shortest window. Short-time windows, as with the van der Pol oscillator, require only a crude approximation of the period to reach the 1% error tolerance. Again, as with the van der Pol oscillator, the lead-in time combined with the need to compute the period of the oscillation, however approximate, may make long-time windows the desirable option.

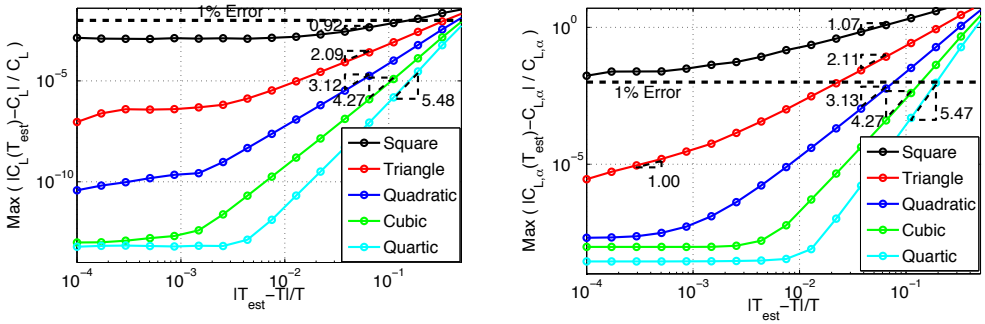

(a) Output error (b) Sensitivity error

Figure 14: Effect of period estimate error on periodic time-average error

6. Conclusion

In this work, we investigated the convergence of the outputs and output parameter sensitivities of quasi-periodic problems using finite time simulation. We showed that outputs and sensitivities computed in this manner can take excessive time to converge, and that parameter sensitivities can fail altogether to converge to the periodic value. We developed a theoretical basis for this lack of convergence, showing that it results from the dependency of the period on the parameter of interest, and demonstrated it using ODE and PDE examples. To mitigate this lack of convergence, we investigated long-time and short-time windows applied to the instantaneous output: testing examples from both classes, demonstrating their respective convergence rates, and identifying their advantages and disadvantages. These windows enabled the accurate computation of periodic output sensitivities and decreased simulation time to compute time-averaged outputs.

Simulations using both long- and short-time windows require a relaxation time to equilibrate the solution and adjoint or tangent variables. For engineering problems, this lead-in time may be much longer than the required window span. Short-time windows demonstrate that the output of interest can be windowed over a very small number of periods, but additionally require knowledge of the period of oscillation. Additionally, with short-time windows the temporal discretization must accurately capture the shape of the window at any discontinuities, and failure to do so can introduce error into the output or sensitivity. Long-time windows still require a relaxation time to equilibrate adjoint or tangent states, but can significantly decrease the required simulation time without the need to compute the period of oscillation. In our experiments, long-time windows have been less sensitive to interactions between the temporal discretization and the window shape with only a small increase in the span of the window to achieve a given error tolerance.

The benefit of output windows is limited to deterministic problems. Preliminary testing with a chaotic Lorenz oscillator confirms this, with non-square windows performing no better than the square window for either output or parameter sensitivity calculations. This is due to the exponential growth of perturbations exhibited by chaotic systems.

Future work will investigate alternate methods to compute parameter sensitivities of time-average outputs for chaotic problems.

References

[1] He, J.-W., Glowinski, R., Metcalfe, R., and Periaux, J., "Active Control and Drag Optimization for Flow Past a Circular Cylinder," *J. Comput. Phys.*, Vol. 163, 2000, pp. 83–117.

[2] Nadarajah, S. K., McMullen, M. S., and Jameson, A., "Optimum Shape Design for Unsteady Flows Using Time Accurate and Non-Linear Frequency Methods," AIAA 2003-3875, 2003.

[3] Rumpfkeil, M. P. and Zingg, D. W., "Unsteady Optimization Using a Discrete Adjoint Approach Applied to Aeroacoustic Shape Design," AIAA 2008-18, 2008.

[4] Mani, K. and Mavriplis, D. J., "An Unsteady Discrete Adjoint Formulation for Two-Dimensional Flow Problems with Deforming Meshes," AIAA 2007-60, 2007.

[5] Nadarajah, S. K. and Jameson, A., "Optimum Shape Design for Unsteady Flows With Time-Accurate Continuous and Discrete Adjoint Methods," *AIAA Journal*, Vol. 45, No. 7, 2007, pp. 1478–1491.

[6] Duta, M. C., Giles, M., and Campobasso, M. S., "The Harmonic Adjoint Approach to Unsteady Turbomachinery Design," *Internat. J. Numer. Methods Fluids*, Vol. 40, No. 3–4, 2002, pp. 323–332.

[7] Thomas, J. P., Hall, K. C., and Dowell, E. H., "Discrete Adjoint Approach for Modeling Unsteady Aerodynamic Design Sensitivities," *AIAA Journal*, Vol. 43, No. 9, 2005, pp. 1931–1936.

[8] Nadarajah, S. K. and Jameson, A., "Optimum Shape Design for Unsteady Three-Dimensional Viscous Flows Using a Nonlinear Frequency-Domain Method," *AIAA Journal of Aircraft*, Vol. 44, No. 5, 2007, pp. 1513–1527.

[9] Srinath, D. N. and Mittal, S., "An Adjoint Method for Shape Optimization in Unsteady Viscous Flows," *J. Comput. Phys.*, Vol. 229, 2010, pp. 1994–2008.

[10] Wilkins, A. K., Tidor, B., White, J., and Barton, P. I., "Sensitivity Analysis for Oscillating Dynamical Systems," *SIAM J. Sci. Comput.*, Vol. 31, No. 4, 2009, pp. 2706–2732.

[11] Tomovic, R. and Vulkobratovic, M., *General Sensitivity Theory*, American Elsevier, New York, NY, 1972.

[12] Barth, T., "Space-Time Error Representation and Estimation in Navier-Stokes Calculations," *Complex Effects in Large Eddy Simulations*, edited by T. J. Barth, M. Griebel, D. E. Keyes, R. M. Nieminen, D. Roose, T. Schlick, S. C. Kassinos, C. A. Langer, G. Iaccarino, and P. Moin, Vol. 56 of *Lecture Notes in Computational Science and Engineering*, Springer Berlin Heidelberg, 2007, pp. 29–48, 10.1007/978-3-540-34234-2_3.

[13] Boyd, J. P., *Chebyshev and Fourier Spectral Methods*, Dover, Mineola, NY, 2001, http://www-personal.engin.umich.edu/jpboyd/.

[14] Bassi, F. and Rebay, S., "High-order accurate discontinuous finite element solution of the 2D Euler equations," *J. Comput. Phys.*, Vol. 138, No. 2, 1997, pp. 251–285.

[15] Bassi, F. and Rebay, S., "A High-order discontinuous finite element method for the numerical solution of the compressible Navier-Stokes equations," *J. Comput. Phys.*, Vol. 131, 1997, pp. 267–279.

[16] Cockburn, B. and Shu, C.-W., "Runge-Kutta discontinuous Galerkin methods for convection-dominated problems," *J. Sci. Comput.*, 2001, pp. 173–261.

[17] Fidkowski, K. J., Oliver, T. A., Lu, J., and Darmofal, D. L., "p-Multigrid solution of high-order discontiunous Galerkin discretizations of the compressible Navier-Stokes equations," *J. Comput. Phys.*, Vol. 207, No. 1, 2005, pp. 92–113.

[18] Wang, L. and Mavriplis, D. J., "Implicit Solution of the Unsteady Euler Equations for High-Order Accurate Discontinuous Galerkin Discretizations," AIAA 2006-0109, 2006.

[19] Krakos, J. and Darmofal, D., "Effect of Small-Scale Output Unsteadiness on Adjoint-Based Sensitivity," *AIAA Journal*, Vol. 48, No. 11, Nov. 2010, pp. 2611–2623.