https://ntrs.nasa.gov/search.jsp?R=20190000875 2020-05-09T09:53:11+00:00Z

Verifying Implementation of the Dryden Turbulence Model and MIL-F-8785 Gust Gradient

Michael M. Madden*
NASA Langley Research Center, Hampton, VA, 23681, USA

Turbulence modeling in human-in-the-loop simulation is important to assessing aircraft handling qualities and pilot performance and to provide additional realism for pilot training. In the simulation community, the Dryden turbulence spectra is a popular choice for modeling the linear turbulent gusts because its rational form is efficiently reproduced by passing white noise through linear filters. The MIL-F-8785 gust gradients similarly use additional linear filters to model the gradient of the turbulent gust over the wing, and it represents the gust gradients as perturbations to the air-relative rotational rates. The Cockpit Motion Facility at NASA Langley Research Center (LaRC) models continuous random turbulence using the Dryden one-dimensional spectra and MIL-F-8785 gust gradient. The facility recently reviewed and updated its verification of these models as part of an initiative to improve motion cueing under turbulence. This exercise introduced improved methods for verifying the turbulence models and led to rediscovery of model assumptions that informed improvements to implementation.

I. Nomenclature

Δt	=	the simulation time step (s)
Ω	=	spatial frequency (rad/ft)
Φ	=	power spectral density (PSD) of a turbulence quantity
σ	=	the turbulence intensity (ft/s)
ω	=	temporal frequency (rad/s)
C_1, C_3, C_3	=	coefficients of the digital form of a first order filter.
DFT	=	a discrete Fourier transform (DFT) sample
L	=	the characteristic length of a turbulence quantity at altitude (ft)
L_{FA}	=	the characteristic length of turbulence in the free atmosphere (ft)
N_{DFT}	=	the minimum number of points for the DFTs used to reconstruct Φ
N_{total}	=	the minimum total number of points to reconstruct Φ
V	=	the magnitude of the vehicle's velocity vector (ft/s)
V_{stall}	=	stall speed of a vehicle
W	=	white noise input for a gust component
a, b, c, d	=	Laplace coefficients of a general first order filter
b	=	wingspan (ft); also used as a Laplace coefficient of a first order filter.
$\bar{c}$	=	mean aerodynamic chord (ft)
f	=	execution frequency of the simulation (Hz)
h	=	altitude above ground level (AGL) (ft)
l_T	=	the wing-to-tail distance (ft)
p, q, r	=	roll, pitch, and yaw components of an angular velocity vector
p', q', r'	=	the angular rate perturbations of the turbulence (rad/s)
t	=	current time (s)
u, v, w	=	longitudinal, lateral, and vertical components of a velocity vector in body axes
u', v', w'	=	the linear perturbations of the turbulence (ft/s)
x, y	=	input and output, respectively, of a digitized filter

* Chief Engineer for Modeling and Simulation and AIAA Senior Member.

II. Introduction

About ten years ago, the Cockpit Motion Facility (CMF) at NASA Langley Research Center (LaRC) began conducting research on improving the motion cueing of aircraft response to atmospheric turbulence [1, 2]. Comments made by experienced pilots acting as test subjects raised the possibility that the simulation may be deficient in modeling the response of a large civil transport to turbulence. As a result, LaRC decided in 2013 to review its modeling of aircraft response to turbulence. This paper highlights one activity of that review, to revisit the software implementation and verification of the turbulence models. To model continuous turbulence experienced by an aircraft in flight, the CMF implements models derived from the Dryden one-dimensional turbulence spectra for the linear gusts and from the MIL-F-8785 gust gradient spectra to represent angular rate perturbations. The turbulence model code used in the CMF is part of the Langley Standard Real-Time Simulation in C++ (LaSRS++), an object-oriented simulation framework [3]. The LaRC review of the turbulence model consisted of both a code review and a functional verification. The code review compared the turbulence model code against algorithmic descriptions of the Dryden gust and MIL-F-8785 gust gradient found in the literature. Observations from the code review, which may be of general interest and necessary to replicate the verification described here, are discussed in section III. The functional verification assessed the time history output of the code using three criteria:

1. The mean of the output should be zero.
2. The variance of the linear gust outputs should equal the square of the turbulence intensity (σ) parameter set for each gust component, where σ^2 is the total turbuelene energy of the gust component. The total energy of the gust gradients are a fraction of the total energy of an associated linear gust, as a function of the aircraft wingspan. Thus, the variance of a gust gradient output should equal the total energy of that gust gradient.
3. The turbulent outputs should have a power spectral density (PSD) that reproduces the Dryden and MIL-F-8785 spectra.

Reproducing the spectra is the more difficult of the criteria to demonstrate. Section IV details the method used to reproduce the spectra. Section V presents the results of the functional verification against the three success criteria.

III. Implementation Algorithm for the Dryden and MIL-F-8785 Models

Equations (1) through (3) present the Dryden one-dimensional spectra Φ both in spatial form (Ω) and temporal form (ω) for the longitudinal (u), lateral (v), and vertical components (w) of the turbulent gust [4]:

$$\Phi_u(\Omega) = \sigma_u^2 \frac{2L_u}{\pi} \frac{1}{1+(L_u\Omega)^2} \quad \Phi_u(\omega) = \sigma_u^2 \frac{2L_u}{\pi V} \frac{1}{1+\left(L_u \frac{\omega}{V}\right)^2} \tag{1}$$

$$\Phi_v(\Omega) = \sigma_v^2 \frac{L_v}{\pi} \frac{1+3\,(L_v\Omega)^2}{[1+(L_v\Omega)^2]^2} \quad \Phi_v(\omega) = \sigma_v^2 \frac{L_v}{\pi V} \frac{1+3\,\left(L_v \frac{\omega}{V}\right)^2}{\left[1+\left(L_v \frac{\omega}{V}\right)^2\right]^2} \tag{2}$$

$$\Phi_w(\Omega) = \sigma_w^2 \frac{L_v}{\pi} \frac{1+3\,(L_w\Omega)^2}{[1+(L_w\Omega)^2]^2} \quad \Phi_w(\omega) = \sigma_w^2 \frac{L_v}{\pi V} \frac{1+3\,\left(L_w \frac{\omega}{V}\right)^2}{\left[1+\left(L_w \frac{\omega}{V}\right)^2\right]^2} \tag{3}$$

where V is the speed of the vehicle, σ is the turbulence intensity (square root of the total turbulent energy), and L is the characteristic length of the turbulence field. In the free atmosphere, the Dryden spectra is isotropic. Thus, for all three gust components, the turbulence intensity parameters are equal ($\sigma_u = \sigma_v = \sigma_w$), and the turbulence characteristic lengths are equal ($L_u = L_v = L_w$) [4].[†] However, in the planetary boundary layer, turbulence is approximately isotropic only in the horizontal plane [5]. Thus, the horizontal parameters remain equal ($\sigma_u = \sigma_v$; $L_u = L_v$) in a reference frame perpendicular to the local nadir, but the vertical parameters (σ_w and L_w) differ. MIL-STD-8785B [4] starts the free

[†] MIL-HDBK-1797 [7] defines $L_u = 2L_v = 2L_w$ in the free atmosphere using the definition of the distance to zero correlation between points; the spectra equations are adjusted accordingly so that the mathematics remain unchanged. However, Ref. [8] ascribes physical meaning only to the "integral scale of the turbulence" L while noting that the transverse scale is L/2. Retaining an equal L in all three directions more easily conveys a turbulence field that is frozen and isotropic. It also simplifies concepts described later in the paper.

atmosphere at an altitude of 1750 ft AGL ("above ground level"). MIL-STD-8785C [6] and MIL-HDBK-1797 [7] raise the altitude to 2000 ft AGL. However, MIL-STD-8785C and MIL-HDBK-1797 define the planetary boundary layer only for altitudes below 1000 ft AGL, leaving a gap in the model between 1000 ft and 2000 ft AGL.[‡] MIL-F-8785B defines the turbulence model in the planetary boundary layer with no gaps and no discontinuity at the transition to the free atmosphere. Therefore, the LaSRS++ turbulence code implements the MIL-F-8785B model for turbulence in the planetary boundary layer with a departure, described later, for modeling the turbulence intensity.

Equation (4) shows the turbulence characteristic length L as function of AGL altitude h within the planetary boundary layer [4]:

$$\left.\begin{array}{c} L_u = L_v = {L_{FA}}^{2/3}\sqrt[3]{h} \\ L_w = h \end{array}\right\} \quad h < L_{FA} \tag{4}$$

where L_{FA} is the turbulent characteristic length in the free atmosphere, i.e. 1750 ft. Equation (5) shows the relationship among the turbulence intensities in the planetary boundary layer. These relationships are derived from MIL-STD-8785B, but they use the longitudinal turbulence intensity σ_u as the independent variable.

$$\sigma_v = \sigma_u \quad \sigma_w = \sigma_u \sqrt[3]{\frac{h}{L_{FA}}} \tag{5}$$

The LaSRS++ turbulence code, however, does not implement the MIL-F-8785B plot of σ_w as a function of altitude nor does it apply the recommendation in MIL-F-8785C and MIL-HDBK-1797 that, in the planetary boundary layer, σ_w is set to 10% of the mean wind speed. Instead, σ_u (not σ_w) is simply set to a user-defined value. This permits the simulation to produce any level of turbulence intensity suitable to the purpose of an experiment.

The MIL-F-8785 gust gradients are a Taylor series approximation of the gradient of the Dryden gust over a planar wing with a span of b [10]. Equations (6) through (8) show the mapping of the gust gradient to an angular rate perturbation and the resulting spatial spectra for that perturbation [4]:

$$p_g = -\frac{\partial w_g}{\partial y} \quad \Phi_{p_g}(\Omega) = \frac{\sigma_w^2}{L_w} \frac{\frac{4}{5}\sqrt[3]{\frac{\pi L_w}{4b}}}{1 + \left(\frac{4b}{\pi}\Omega\right)^2} \tag{6}$$

$$q_g = \frac{\partial w_g}{\partial x} \quad \Phi_{q_g}(\Omega) = \frac{\Omega^2}{1 + \left(\frac{4b}{\pi}\Omega\right)^2} \Phi_w(\Omega) \tag{7}$$

$$r_g = -\frac{\partial v_g}{\partial x} \quad \Phi_{r_g}(\Omega) = \frac{\Omega^2}{1 + \left(\frac{3b}{\pi}\Omega\right)^2} \Phi_v(\Omega) \tag{8}$$

To create a time series of turbulent gusts and their gradients, white noise W is fed into filters derived from the Dryden and MIL-F-8785 spectra. LaSRS++ generates white noise using the random number algorithm of Ref. [11] as modified by Ref. [12].[§] The white noise is then scaled by $\sqrt{\pi/\Delta t}$ to normalize the spectral power of its time series to one. Equations (9) and (10) show the Laplace form of the turbulence filters as they are represented in the post-review LaSRS++ code. This formulation is derived from Ref. [14] except that the 2nd order filters have been factored into a pair of first-order filters. The formulation also retains the term in the numerator of the first filter as part of the filter because V (vehicle speed) and L (the turbulence characteristic length) can vary with time. This was one of the changes made as a result of the code review. The pre-review code had pulled the numerator of the first filter out as if it were a constant scale factor. The discrete form of the filters in Eqs. (11) and (12) expose the problem with treating the numerator as a constant. In the equations, the numerator is the coefficient b of the Laplace form, and the coefficient a

[‡] Reference [9] uses linear interpolation to fill this gap.

[§] This is the same algorithm used in the predefined random number generator minstd_rand in C++-11 [13].

is zero. The coefficient b cannot directly factor out of the discrete equation as a scaling constant because it is not a multiplier to the past output $y_{t-\Delta t}$. The output $y_{t-\Delta t}$ only relates to b through its past values. It is that relationship through past values that permits the coefficient b to be factored out as a constant scale factor when it is invariant with time. However, this is not appropriate when it can change as a function of time. Retaining the numerator term in the filter creates differences in the time series produced when V or L vary with time. However, it doesn't affect the verification results presented here because the spectra reconstruction requires use of a constant V and L.

$$u' = W_u\sqrt{\frac{\pi}{\Delta t}}\,\frac{\sqrt{\frac{2V}{\pi L_u}}}{s+\frac{V}{L_u}}\sigma_u \quad v' = W_v\sqrt{\frac{\pi}{\Delta t}}\,\frac{\sqrt{\frac{3V}{\pi L_v}}}{\left[s+\frac{V}{L_v}\right]}\,\frac{s+\frac{V}{\sqrt{3}L_v}}{\left[s+\frac{V}{L_v}\right]}\sigma_v \quad w' = W_w\sqrt{\frac{\pi}{\Delta t}}\,\frac{\sqrt{\frac{3V}{\pi L_w}}}{\left[s+\frac{V}{L_w}\right]}\,\frac{s+\frac{V}{\sqrt{3}L_w}}{\left[s+\frac{V}{L_w}\right]}\sigma_w \tag{9}$$

$$p_g = W_p\sqrt{\frac{\pi}{\Delta t}}\,\frac{\left(\frac{\pi}{4b}\right)^{7/6}\sqrt{\frac{4V}{5}}}{\sqrt[3]{L_w}\,s+\sqrt[3]{L_w}\,\frac{\pi}{4b}V}\sigma_w \quad q_g = w'\,\frac{\frac{\pi}{4b}s}{s+\frac{\pi}{4b}V} \quad r_g = v'\,\frac{\frac{\pi}{3b}s}{s+\frac{\pi}{3b}V} \tag{10}$$

The LaSRS++ code digitizes the filters using a matched z-transform. In the pre-review LaSRS++ code, the digitized filters also modeled the white noise input using a first-order hold (FOH) because, at that time, the LaSRS++ filter library only contained digitized forms using a FOH.** However, a zero-order hold (ZOH) is the better representation for uncorrelated, random inputs. References [14] and [15] explicitly identify the use of the ZOH for discretizing the turbulence filters. Furthermore, Ref. [16] mathematically analyzes how the ZOH and FOH affect the variance of the discrete outputs. It concludes that the FOH will show a greater loss of variance correlated with the ratio of the vehicle travel per step (V * Δt) to the turbulence characteristic length (L). Section V.A examines variance loss under both the ZOH and the FOH. The results prompted a change to use a ZOH to discretize the white noise input of the initial first-order filter for each gust component shown in Eqs. (9) and (10), but the FOH is retained for discretizing the inputs to any subsequent filters. Equations (11) and (12) show the discrete algorithms that LaSRS++ employs for the general form of a first-order filter. Equation (11) shows the form using a FOH for the input, and Eq. (12) shows the form using a ZOH for the input.

$$\frac{as+b}{cs+d} \overset{FOH}{\Longrightarrow} y_t = C_1\,y_{t-\Delta t} + C_2\,x_t + C_3\,x_{t-\Delta t} \begin{cases} C_1 = e^{-d\Delta t/c} \\ C_2 = \dfrac{b}{d} + \dfrac{(1-C_1)\left(a - c\dfrac{b}{d}\right)}{\mathrm{d}\Delta t} \\ C_3 = \dfrac{b}{d}(1-C_1) - C_2 \end{cases} \tag{11}$$

$$\frac{as+b}{cs+d} \overset{ZOH}{\Longrightarrow} y_t = C_1\,y_{t-\Delta t} + C_2\,x_t + C_3\,x_{t-\Delta t} \begin{cases} C_1 = e^{-d\Delta t/c} \\ C_2 = a/c \\ C_3 = \dfrac{b}{d}(1-C_1) - C_2 \end{cases} \tag{12}$$

IV. Verification Method

A good estimate of the mean, variance, and power spectral density (PSD) can be obtained from a single time history if it is of sufficient length. Prior verifications performed by the simulator group at Langley appeared to arbitrarily pick tens of thousands to hundreds of thousands of points. Subsection A presents a quantitative method for computing the minimum number of points to obtain a good estimated PSD for a given test case. For completeness, the verification should also exercise a variety of scenarios using different values for vehicle speed (V), turbulence characteristic length (L), and simulation time step (Δt). For this review, the author designed cases representative of subsonic flight for a variety of aircraft (general aviation aircraft, fighters, and large transports) in both the free atmosphere and the planetary boundary layer; see subsection C for more detail.

** The filter library was primarily designed to support control system models where inputs are expected to be continuous processes or functions.

A. Determining the number of data points for each scenario

The criteria to reconstruct the spectra function for each turbulence component drives the number of data points required for each simulation scenario. Because the turbulence time history is stochastic, reconstructing the spectra equations requires an average of multiple independent discrete Fourier transform (DFT) samples. Furthermore, each DFT sample must be of sufficient resolution to capture the characteristic features of the spectra function.

First, the number of points needed to construct a single DFT sample of sufficient resolution is determined. Table 1 shows the selected characteristic points and the value of the temporal frequency ω where these characteristic points are located in the spectra functions Φ_u (longitudinal) and Φ_w (vertical) as defined by Eqs. (1) and (3). The selected characteristic points are the values of the function that are fractions of the power at zero frequency (using fractions that are multiples of ¼), points of maxima and minima ($d\phi/dt = 0$) and points of inflection ($d^2\phi/dt^2 = 0$).

		Longitudinal	Vertical
Fraction of zero frequency power	1	0	0 V/L
	¾	0.57735 (V/L)	1.46789 (V/L)
	½	(V/L)	2.05817 (V/L)
	¼	1.73205 (V/L)	3.20804 (V/L)
Maxima and Minima		0	0
			0.57735 (V/L)
Inflection Points		0.57735 (V/L)	0.27395 (V/L) 1.21676 (V/L)

Table 1 Frequency Location of Characteristic Points for the Dryden Spectra

The smallest unique point other than zero is approximately ¼ V/L for the vertical component. The other points are approximately integer multiples of this point. Therefore, to reconstruct the shape of the vertical component's spectra, the first point of a DFT must represent this frequency (in rad/s) or an approximate whole fraction of it. Because the frequency range in the DFT is one half the frequency at which the simulation is run, the number of points required to produce this minimum frequency increases with increasing frame rate (or decreasing time step) and is doubled. Equation (13) shows the resulting formula for the approximate number of data points per DFT in the expressions following the ~ operator:

$$N_{DFT} \sim \frac{2\,\pi\,f}{\Delta\omega} \sim \frac{8\,\pi\,f\,L}{V} = 2^{\left\lceil \log_2 \frac{8\,\pi\,f\,L}{V} \right\rceil} \tag{13}$$

where f is the simulation frequency, i.e. 1/Δt. This number must then be rounded up to the next power of two. The final expression containing this rounding is shown following the equal sign. It uses the ceiling operator $\lceil x \rceil$ on the base-2 logarithm of the approximate expression and raises 2 to the resulting power. For example, if the length scale is 1750, the velocity is 300, and the frequency is 20 Hz, then the approximate formula returns $N_{DFT} = 2932$. This must then be rounded up to the next power of two using the final expression, resulting in $N_{DFT} = 4096$. For the verification presented here, the number of points per DFT was quadrupled from the minimum specified by Eq. (13) to better capture the curvature about the inflection and maxima points of the Φ_v (lateral) and Φ_w (vertical) spectra.

If one used only enough points for one DFT, one would get a jagged looking DFT because the turbulence is a random stochastic process. Each DFT represents a random power value. To reproduce the power spectra requires enough sample points to produce multiple DFTs and average the results. However, to obtain a good average, each DFT sample must also be independent. One characteristic of the turbulence filters is that they produce random data that is correlated over a time scale equal to L/V. Thus, only collections of DFTs numbering (L f)/V (rounded up to the next whole number) can be treated as independent. Furthermore, dependent on how closely one wants to the narrow the confidence interval of the reconstructed PSD about the spectra function, one must have multiple independent sample sets of DFTs. If one wants the estimated power level to be within one computed standard deviation of the true power level with 99% confidence, then nine independent DFT sample sets are required. For the PSD reconstruction, 36 independent sets of DFT samples were used to achieve 95% confidence of being within 1/3 of a standard deviation. Therefore, the equation for the minimum number of points to reproduce a PSD with high confidence bounds is:

$$N_{total} \sim 36\left(\frac{L\,f}{V}\right)\left(\frac{8\,\pi\,f\,L}{V}\right) \sim 288\pi\left(\frac{L\,f}{V}\right)^2 = 36\left\lceil\frac{L\,f}{V}\right\rceil\left(2^{\left\lceil\log_2\frac{8\,\pi\,f\,L}{V}\right\rceil}\right) \tag{14}$$

For the prior example (L=1750, V=300, f=20), this equation indicates that 17,252,352 points are needed in the time series. This is 4096 points for each DFT multiplied by 117 DFTs per independent set and multiplied again by 36 independent sets for the target confidence interval. For a frequency of 50 Hz, the number of points increases to 85,105,152. As stated in the prior paragraph, the number of points used in the verification results presented here for these examples was quadrupled to improve the ω resolution of the reconstructed PSD.

B. Spectra Reproduction Method

To reproduce the spectra, the output time series is divided into a number of DFT samples using the number of points in Eq. (13) multiplied by four to improve the resolution of the curvature about the characteristic points. The Fast Fourier Transform (FFT) method then transforms the sample points into a DFT. The DFTs are combined into an estimated PSD using Bartlett's method, which is a simple average of the DFTs. The PSD is then scaled by ($N_{DFT}\,\Delta t/\pi$), the length of time represented by the half cycle in radians of the one-sided periodogram; this converts the PSD from a function of Hz to a function of rad/s. Lastly, the periodogram was multiplied by the vehicle speed to be a function of rad/ft (Ω). This scaling enables co-plotting of the reconstructed PSDs for a given turbulence component from multiple test cases with the same turbulence characteristic length L (i.e. at the same altitude). Then, the reconstructed PSDs for that turbulence component can be directly compared against its ideal spectra function as presented in Eqs. (1), (2), (3), (6), (7), or (8). Equation (15) shows the total scaling and averaging.

$$\Phi(\Omega) = V\frac{N_{DFT}\,\Delta t}{\pi}\frac{1}{n}\sum_{i=1}^{n} DFT_i \tag{15}$$

C. Verification Scenarios

The verification scenarios were designed to represent flight from a wide variety of aircraft (general aviation, large transport, and fighters) at subsonic speeds in both the free atmosphere and the planetary boundary layer. In the free atmosphere, aircraft were assumed to fly no slower than initial approach speed and no faster than cruise speed. This range of speeds is approximately 150 to 900 ft/s. For flight within the planetary boundary layer but above 200 feet, aircraft were assumed to fly no slower than 120 ft/s and no faster than 422 ft/s (250 kts). By the time aircraft reach decision height (200 ft AGL), they are assumed to be landing at a speed that is $1.3*V_{stall}$. Table 2 shows the 18 pairs of altitude and speed used in the verification presented here.

Altitude (ft)	**Speed (fps)**					
3000	150	300	450	600	750	900
1000	120	180	240	300	360	420
200	110	150	200	240	280	350

Table 2 Altitude-Speed Pairs Used for Verification

To verify that the turbulence models exhibit no dependencies on execution rate, the turbulence model was run, for each altitude-speed pair, at rates of 20 Hz, 32 Hz, and 50 Hz considered representative of flight simulation execution rates. This results in a total of 54 test cases. Because Eqs. (9) and (10) apply the turbulence intensity σ as a scale factor to the filter output, a variety of values for σ_u, σ_v, and σ_w is not necessary, but they should have a value other than 1.0 ft/s to avoid the potential of false negative results. All test cases used a value of 2.0 ft/s. Similarly, all test cases used the same random number seeds: 123456789 for the longitudinal white noise (W_u), 2010924726 for the lateral white noise (W_v), 417893401 for the vertical white noise (W_w), and 1255 for the roll white noise (W_p). Lastly, for the MIL-F-8785B gust gradient components, the test cases used the wingspan b of a Boeing 757-200 (124.8 ft).

D. An Important Caution on Execution Frequency of Turbulence Models

The verification documented in this paper was completed before a LaRC review of the vehicle response modeling to turbulence. The results of that review were documented in Ref. [17] which makes the following observation:

> …current practice is to apply turbulent gusts to the vehicle model by adding the gusts to the freestream inputs of the aerodynamics model… The freestream, however, is defined at a length scale that is a multiple of the vehicle characteristic lengths, and continuous turbulence models represent turbulence as the accumulation of energies from a continuum of wavelengths including wavelengths smaller than the vehicle characteristic lengths. Thus, there is a limit to the turbulence wavelengths applicable using [this] approximation.

The smallest turbulence wavelength λ contributing to each output sample of the turbulence model is defined by the relation $2V/f$ where V is the vehicle speed and f is the simulation execution frequency in Hz.[††] Thus, as the execution frequency of the turbulence model increases for a given vehicle speed, the turbulence wavelengths contributing to each gust output also become smaller. When the output samples include contributions from turbulence wavelengths smaller than the characteristic lengths of the vehicle, then vehicle response at the higher frequencies devolves into noise. For most aircraft, the freestream assumption is most limiting for the longitudinal wavelength λ_u. If one applies only the linear gusts, which approximates the aircraft as a point, then the valid longitudinal wavelengths λ_u are limited to $\geq 60\ \bar{c}$ (the mean aerodynamic chord) if the vehicle is compact[‡‡] [18] or, for wing-tail configurations, $\geq 120\ l_T$ (the wing-tail length) [19]. Adding the gust gradient, which approximates the aircraft as a rigid planar wing, reduces these numbers by an order of magnitude, i.e. $\lambda_u \geq 6\ \bar{c}$ for compact aircraft or $\geq 12\ l_T$ for wing-tail configurations [18, 19]. The wavelength limit can be reduced again by half for the wing-tail configuration ($\lambda_u \geq 6\ l_T$) if turbulence is applied separately to the tail by tracking when turbulence computed at the wing reaches the tail [19]. Ref. [20] suggests a method for accomplishing this that does not require separate aerodynamic build-up equations for the tail. Ref. [17] applies these limits to an F16-XL and B737-600. Using only linear gusts, the valid execution frequency f for the turbulence model is akin to the short-period frequency for the F16-XL (0.3 – 1.2 Hz) and is akin to the phugoid frequency for the B737-600 (0.05 to 0.2 Hz). Adding the gust gradient increases the valid range of f by an order of magnitude. The valid execution rate of 3 to 12 Hz for F16-XL may be able to produce vehicle response frequencies of 1.5 to 6 Hz. Those frequencies approach the range of 4 to 8 Hz which generates human discomfort. However, the valid execution rates of 0.5 to 2 Hz for the B737-600 may produce vehicle response not indistinguishable from short-period motion. These validity limits are based on an estimated error in vehicle response of 10% [18]. Since turbulence is often used in human-in-the-loop simulation to generate motion cues of shaking in the cockpit, the subjective perception of test subjects may permit greater errors in vehicle response. However, pushing execution rates into frequencies that include energy from turbulence wavelengths λ_u approaching $\bar{c}$ for compact aircraft or l_T for wing-tail configurations is likely to result in a perceived loss of fidelity. Therefore, the turbulence model will likely need to execute at an integer divisor of the simulation's execution frequency to avoid generating high frequency noise in the vehicle response.

V. Functional Verification Results

A. Mean and Variance

For all six turbulence components across the 54 cases, the mean of the time series output was 0 ± 0.003 ft/s or better regardless of the type of hold used to digitize the turbulence filters. On the other hand, use of a zero-order hold (ZOH) or a first order hold (FOH) does impact the variance of the time series output. Figure 1 plots the variance ratio, i.e. the ratio of the variance of the time series output (σ^2_{output}) to the square of the turbulence intensity setting ($\sigma^2_{setting}$), for the following four digital implementations of the linear turbulence filters (u', v' and w'):

1. FOH All Filters. The turbulence filters are factored into a series of 1st order filters. All filters are digitized using a first order hold (FOH). This is the implementation used in the pre-review LaSRS++ code; however, to enable a more direct comparison of FOH and ZOH on the turbulent outputs, the results presented in this section include all other changes made to the post-review LaSRS++ code.
2. ZOH All Filters. The turbulence filters are factored into a series of 1st order filters. All filters are digitized using a zero order hold (ZOH).
3. ZOH No Factoring. The turbulence filters for the linear gusts are not factored into a series of first order filters, i.e., the turbulence filters for the lateral (v') and vertical (w') gusts are digitized as 2nd order filters rather than a pair of first order filters. A ZOH is used to model the inputs of each filter.
4. ZOH First Filter. The turbulence filters are factored into a series of first order filters. A ZOH is used for the first filter in a series. A FOH is used for any subsequent filters. This is the implementation retained for the post-review LaSRS++ code.

[††] This relationship between λ and f uses the Nyquist frequency ($f/2$) that results from the execution frequency f.
[‡‡] An aircraft is 'compact' if it does not have a large tail or other aerodynamic surface at a large distance (e.g., $> 2\bar{c}$) from the aerodynamic center of its wing.

The points are plotted against the normalized length scale (V * Δt/L), the ratio of the vehicle distance per time step to the turbulence length scale. This normalized length scale tends to be small in the free atmosphere where L is 1750 feet and become large near the surface where L is proportional to the AGL altitude (w') or its cube root (u' and v').

In the FOH All Filters case, the time series output shows a variance loss of up to 7% as the normalized length scale increases. Switching all the first order filters to ZOH (ZOH All Filters) causes the opposite trend, variance of the time series output exceeds the set variance by up to 3% as the normalized length scale increases. However, that exceedance trend appears only for the lateral and vertical gusts; the longitudinal gust, which is a single first order filter, remains flat around 1.0. That result prompted another run using a digitized second order filter for the lateral and vertical gusts (ZOH No Factoring). That implementation shows a near and consistent preservation of the variance in the time series output for all three gust components across all normalized length scales. With this knowledge, an attempt was made to improve the results of the factored series of first order filters by using a ZOH only on the first filter (ZOH First Filter). This exercise was necessary, in part, to determine the method for digitizing the turbulent filters for the q' and r' gust gradients, which are first order filters applied to the w' and v' linear gusts respectively. The ZOH First Filter implementation also produced a near-level variance ratio around one for all three gust components, and the graph is nearly indistinguishable from the ZOH No Factoring implementation. The ZOH First Filter result indicates that the ZOH is only needed to properly model the uncorrelated white noise input into the first filter. However, the output of the first filter (and subsequent filters) is a partially correlated time series. Therefore, the FOH better models inputs received from a predecessor filter. The ZOH First Filter approach was further verified using the gust gradients. Figure 2 shows the variance ratios for the time series output from the gust gradients for both the FOH All Filters (pre-review code) and ZOH First Filter (post-review code) implementations. When all filters use a FOH, the time series of the gust gradients show losses of up to 19% as the normalized length scale increases. When the first filter is modified to use a ZOH, then the variance losses largely disappear though a trend remains visible, showing losses of up to 1.4% with increasing normalized length scale. Clearly, a ZOH for the filter is necessary to preserve the selected variance of the turbulence for the linear gusts and, especially, for the gust gradients.

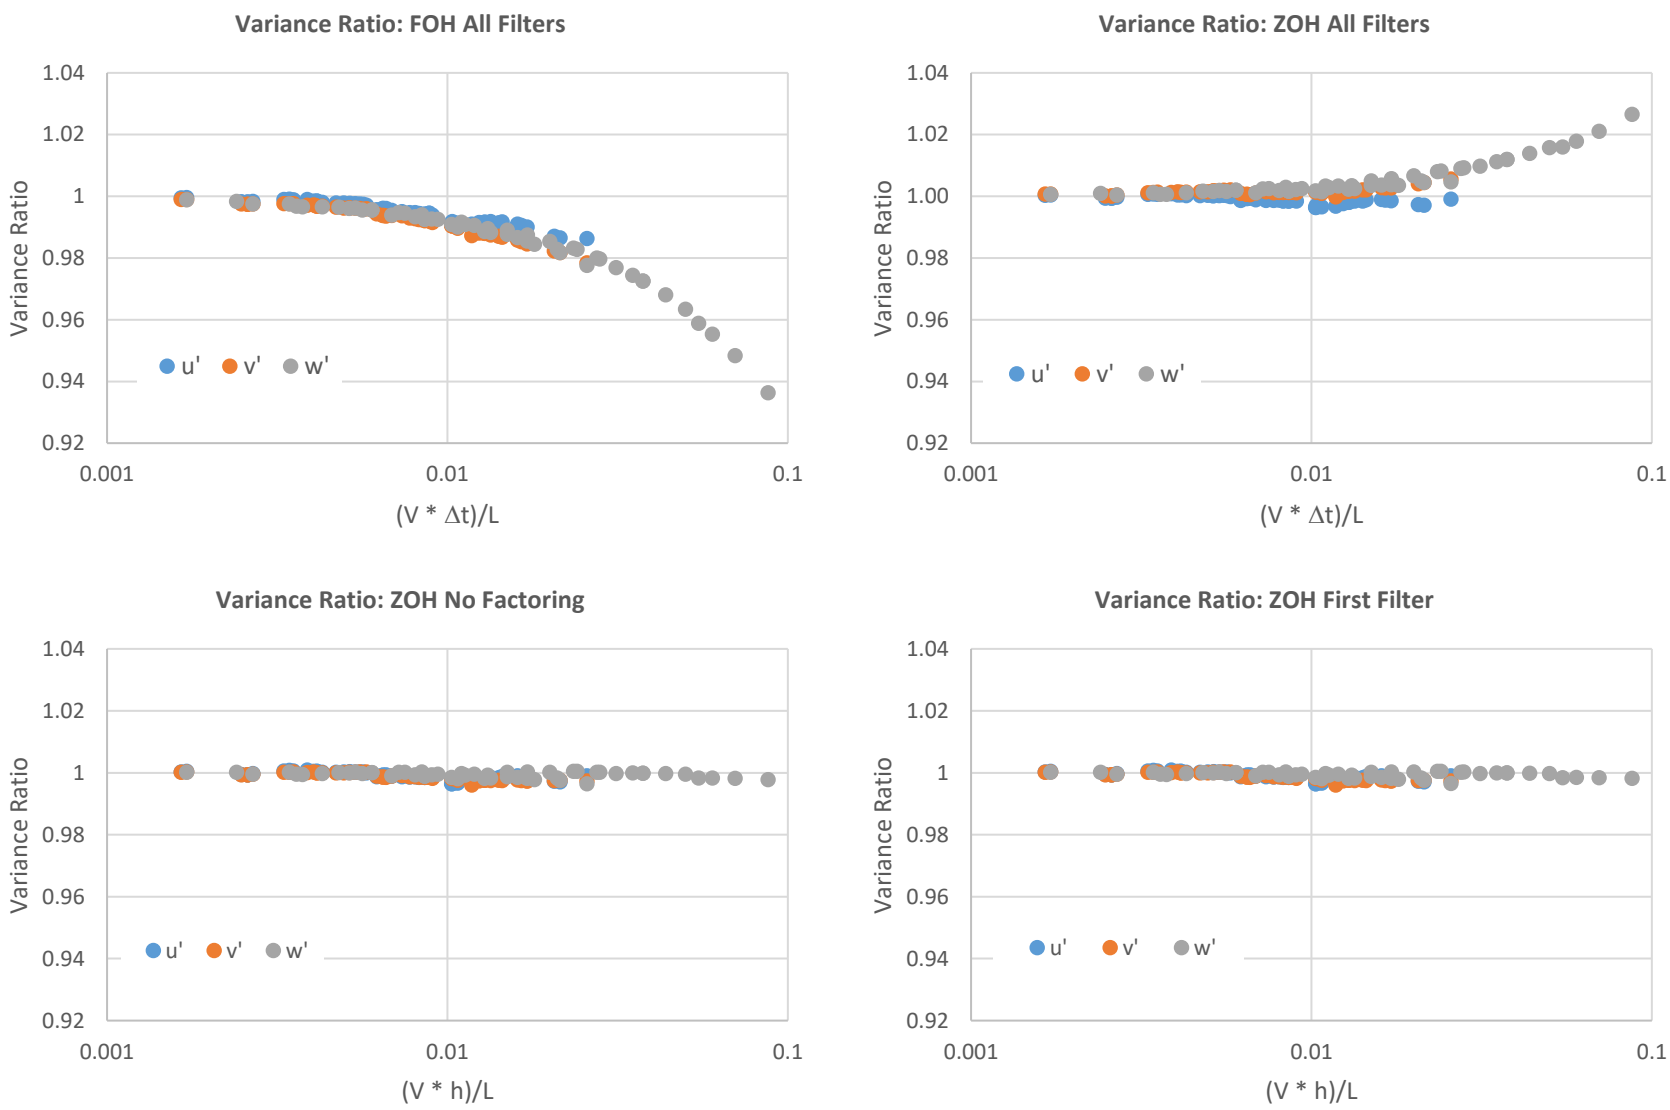

Figure 1 Variance Ratio of Linear Gusts for various ZOH and FOH Configurations

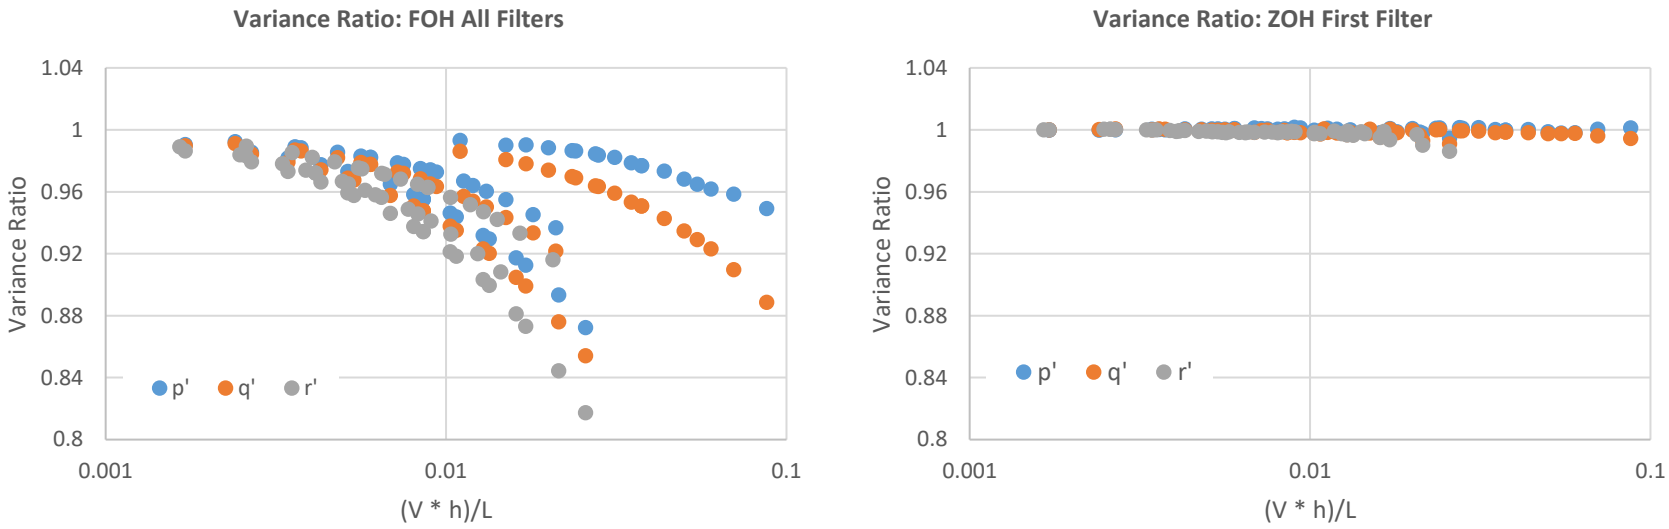

Figure 2 Variance Ratio of Gust Gradient Outputs Before and After ZOH Change to First Filter

B. Spectra Reproduction

Reconstructed spectra are shown in Figure 3 through Figure 6. These figures use time series data from the ZOH First Filter implementation of the post-review LaSRS++ code. Each figure shows two pairs of charts (four total). Each pair shows the reconstructed spectra for a turbulence component at a given altitude. The top pair of charts show the reconstructed spectra at the free atmosphere altitude of 3000 ft AGL. The bottom pair of charts show the reconstructed spectra at an altitude of 200 ft AGL in the planetary boundary layer. These two altitude cases represent extremes in reconstructing the spectra; the 1000 ft case lies somewhere in between and is not shown. The right chart of a pair shows the reconstructed spectra using a log scale, and the spatial frequency (Ω) range is set to the smallest range among the 18 test cases at the given altitude. The frequency range for a test case is [0, $2\pi f$/V]; therefore the test case with the lowest frequency and the highest vehicle speed has the shortest range. For example, at 200 ft AGL, the lowest frequency is 20 Hz and the highest speed is 350 ft/s; the resulting range as displayed in Figure 3 and Figure 4 is [0, 0.18]. The log scale of the right chart permits a better view of the high frequencies where the power contributions are very small. However, to better view reconstructed spectra at the low frequencies, the left chart of the pair uses a linear scale, and Figure 3 and Figure 4 further truncate the spatial frequency range displayed. Figure 3 shows the reconstructed spectra for the longitudinal gust u', and Figure 4 shows the reconstructed spectra for the vertical gust w'. v' is not shown due to similarity in form to w', and w' has the smaller turbulence scale length L at 200 ft, representing a more extreme case. Figure 5 shows the reconstructed spectra for the roll gust gradient p', and Figure 6 shows the reconstructed spectra for the yaw gust gradient r'. Again, the pitch gust gradient q' is not shown due to similarity of form to r', and the time series for r' show a greater loss of variance in Figure 2.

Overall, the reconstructed spectra are in good agreement with the associated spectra function. Nevertheless, there is some visible scattering of power values about the ideal spectra function at low frequencies. However, those errors seem to have no pattern relating them to the vehicle speed or execution frequency of the test cases. At high frequencies, the reconstructed spectra show a fanning out that over-predicts the power contribution from these frequencies. This fanning does appear to be a function of both the vehicle speed and the execution frequency of the model. Though difficult to see from the charts, the errors increase with the vehicle distance per time step (V Δt). In other words, for a given frequency, the over prediction increases with increasing velocity, and, for a given velocity, the over-prediction increases with increasing time step (i.e., decreasing execution frequency). Therefore, the largest over prediction occurs for the pair with the largest vehicle speed and lowest frequency among the 18 test cases. The fanning has limited impact on outcomes because it occurs in a region where the power contributions are an order of magnitude less than the power contribution from the lower frequencies. The chart titles in Figure 3 and Figure 4 present the percentage of the total turbulent energy represented by the frequency range of each chart. The difference in the percentages between the linear (for low frequencies) and log plots (for high frequencies) show that the region of the fanning covers only around 3% to 5% of the total turbulent energy. Even so, this over-prediction appears to benefit the modeled turbulence by partly counteracting the energy lost due to the truncated frequency range associated with the digital execution rate.[§§] In fact, the energy loss due to frequency truncation follows the same correlation; it grows with increasing (V Δt) given that the truncation frequency is $2\pi f$/V = 2π/(V Δt). The cancelation of these two sources of error is evident in Figure 1 and Figure 2 where the time series output from the ZOH First Filter digitization presents very little loss in variance. (Recall that the total energy of the linear gust spectra is the variance.)

Figure 7 illustrates how differences in the reconstructed spectra reflect differences in the variance preservation seen in Figure 1. This figure shows a pair of reconstructed spectra under the most extreme case for w', a vehicle speed of 350 ft/s at an altitude of 200 ft AGL and an execution rate of 20 Hz. The bottom line is the reconstructed spectra from the FOH First Filter implementation, and the top line is the reconstructed spectra from the ZOH First Filter implementation. The ZOH First Filter case shows the over-prediction also seen in Figure 4 and discussed in the prior paragraph. The FOH First Filter case shows an under-prediction at the high frequencies. Moreover, the under prediction drops by over three orders of magnitude near the end of the frequency range. Figure 8 shows all 18 test cases for w' at 200 ft using the FOH First Filter implementation. This figure shows fanning of the under-prediction that increases with V*Δt. This source of lost energy combines with the loss of energy due to frequency truncation (which also correlates with increasing V*Δt) to produce the variance loss shown Figure 1 that increases with (V Δt)/L.

[§§] The energy loss due to frequency truncation can be gleaned from the title of the log scale plots in Figure 3 and Figure 4. The total energy expressed in the plot titles implies a less than 0.5% energy loss due to frequency truncation among the plotted linear gust cases except for w' at 200 ft AGL where the loss is less than 1.9%.

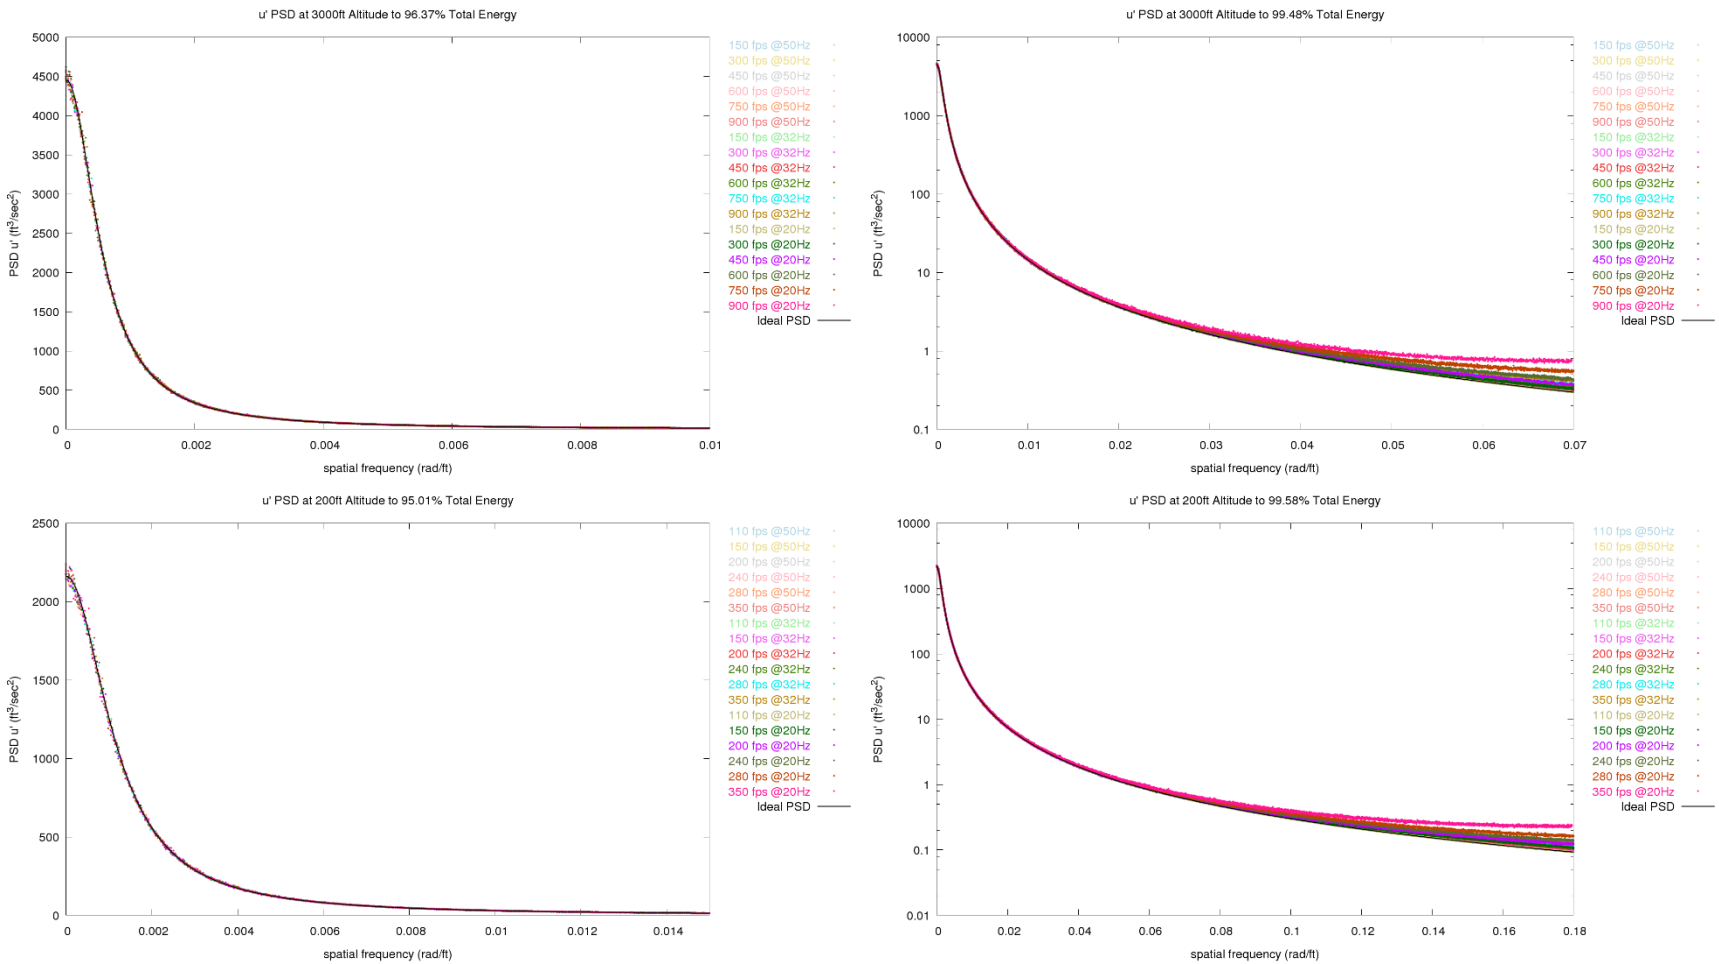

Figure 3 Reconstructed Spectra for u' at 3000 ft and 200 ft AGL Using Linear and Log Scales

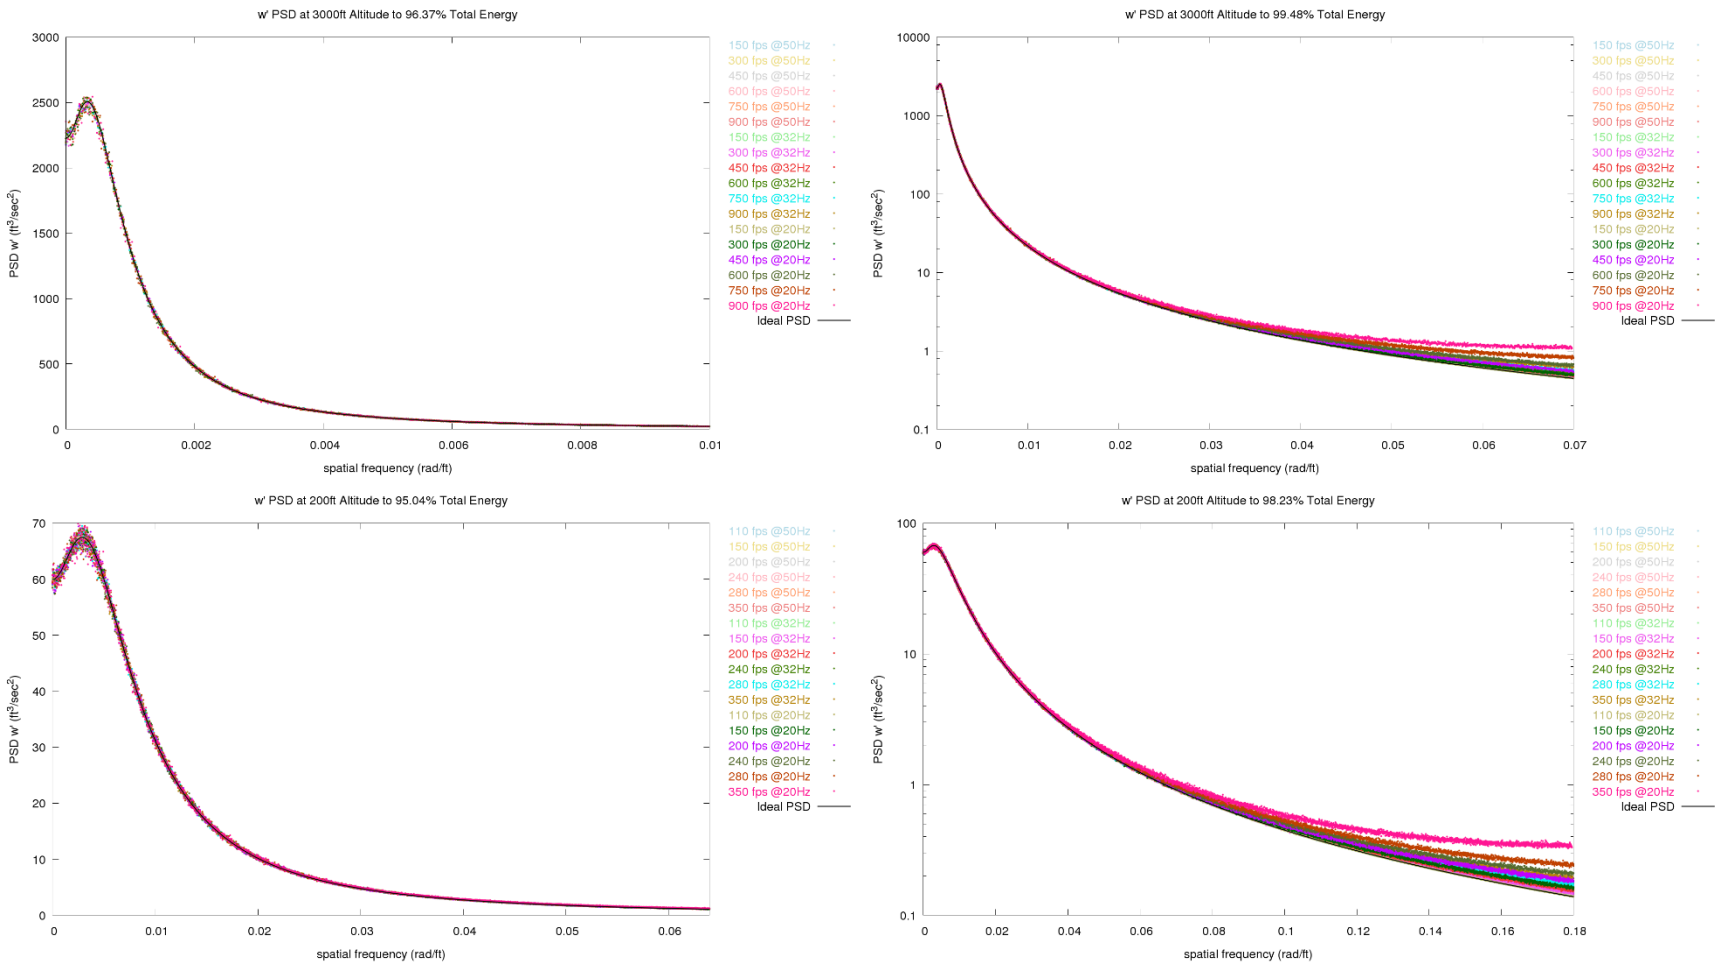

Figure 4 Reconstructed Spectra for w' for 3000 ft and 200 ft AGL Using Linear and Log Scales

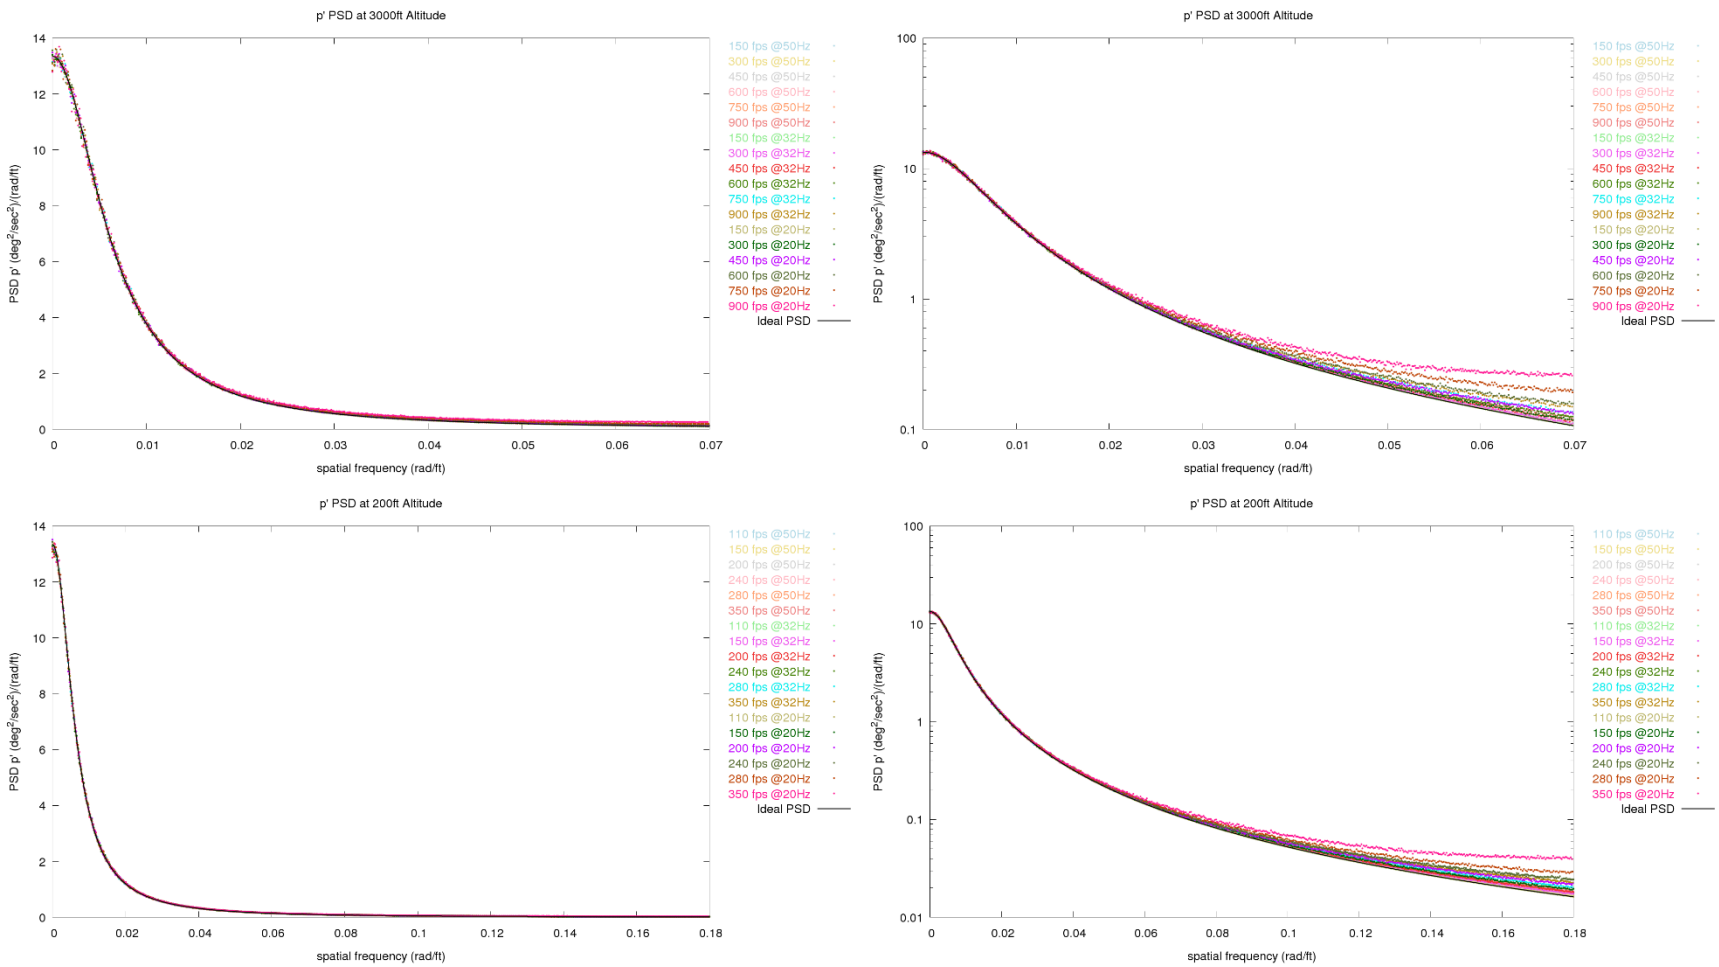

Figure 5 Reproduced Spectra for p' at 3000 ft and 200 ft AGL Using Linear and Log Scales

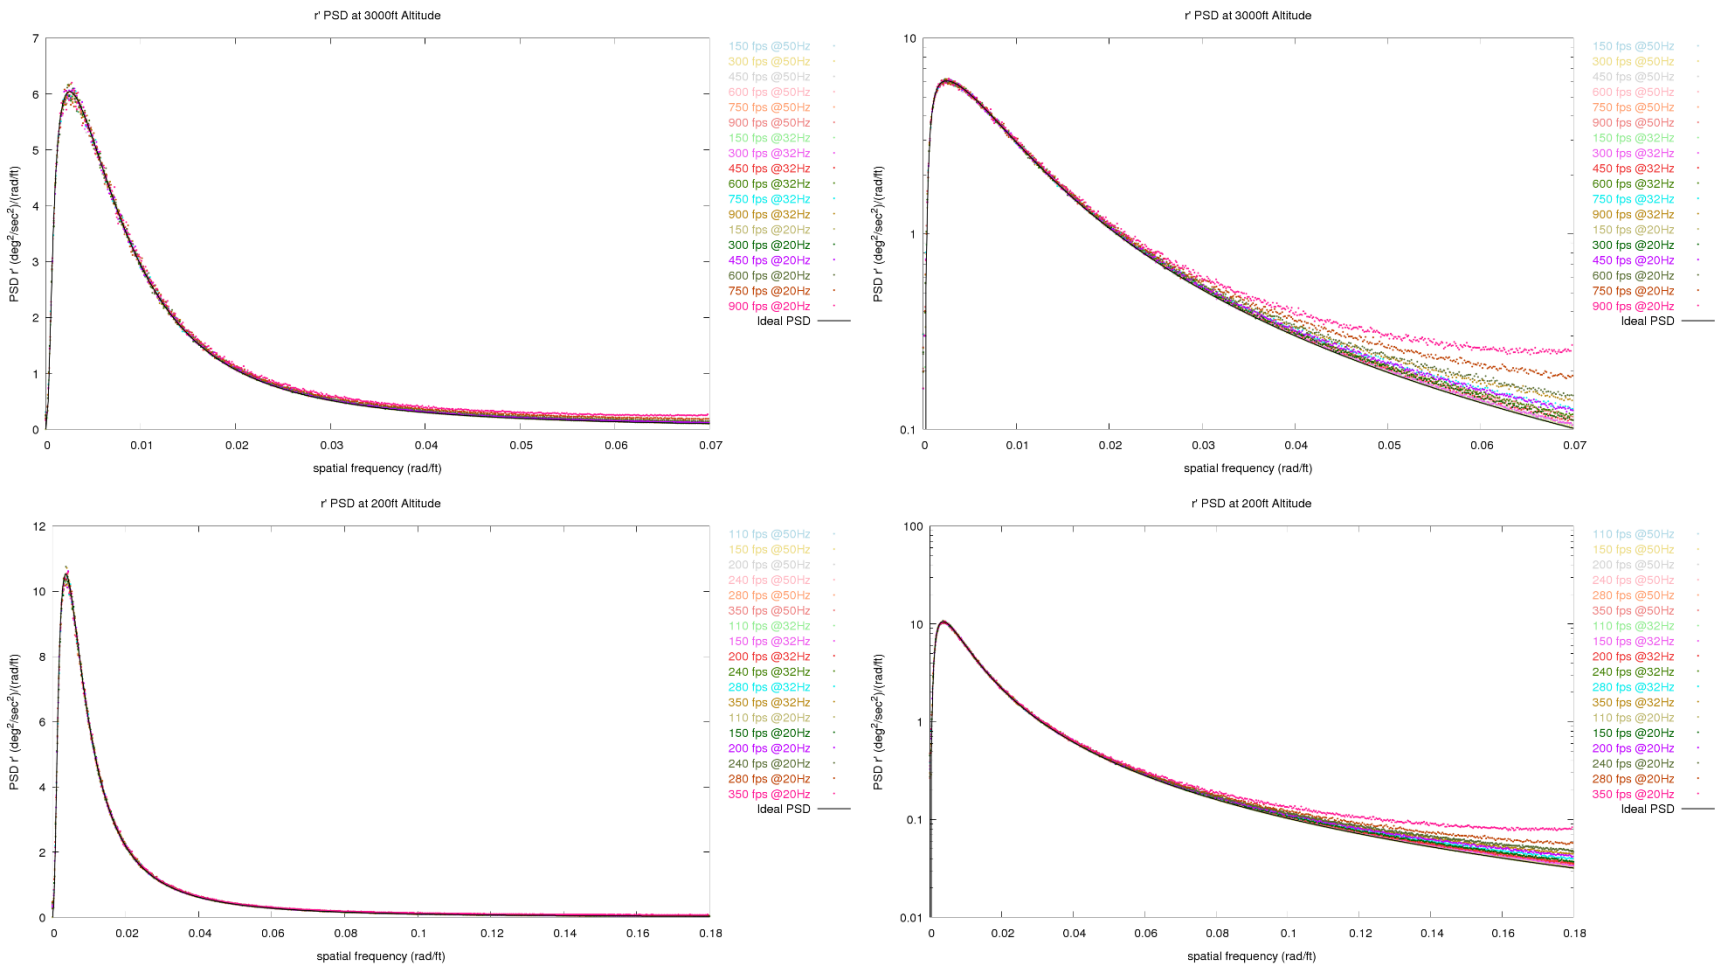

Figure 6 Reconstructed Spectra for r' at 3000 ft and 200 ft AGL Using Linear and Log Scales

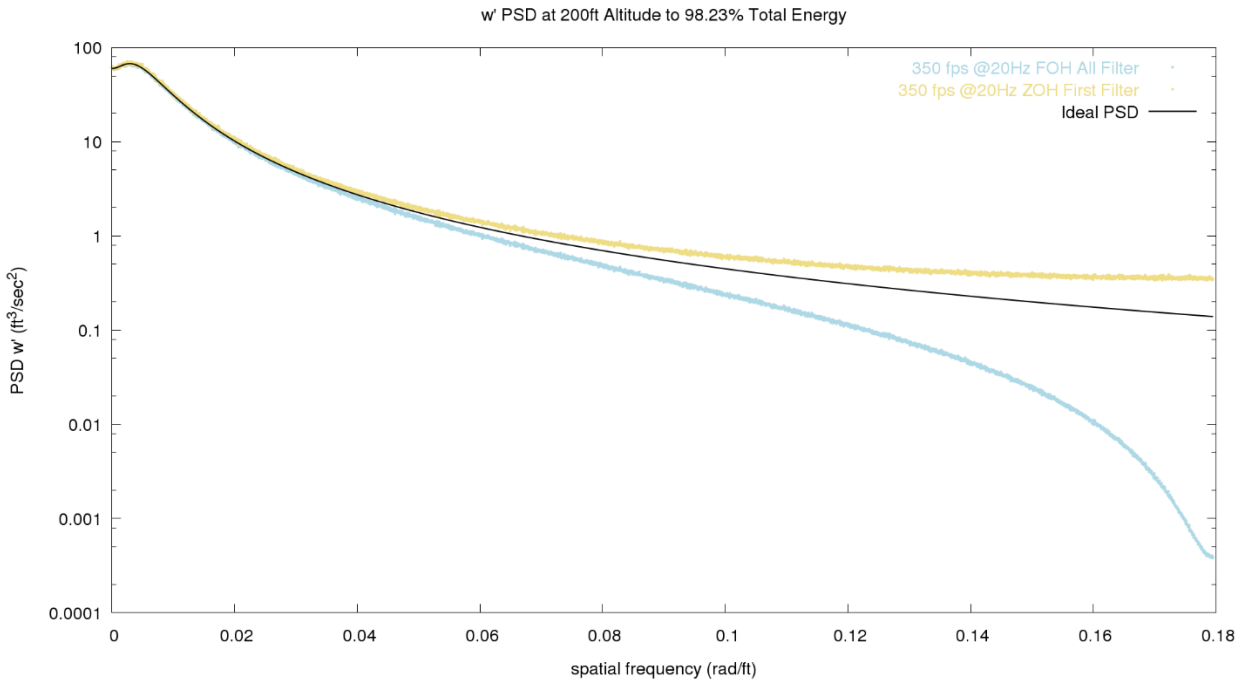

Figure 7 Reconstructed Spectra Comparison between FOH and ZOH on First Filter

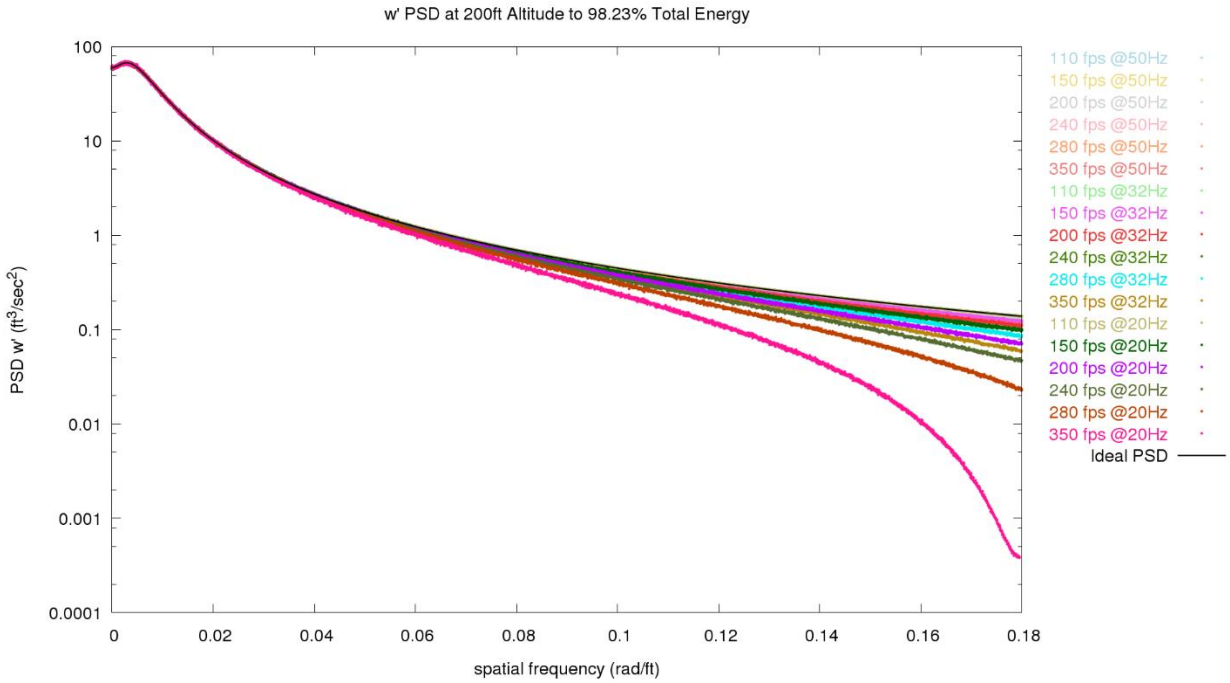

Figure 8 Reconstructed Spectra for w' at 200 ft Using FOH on First Filter

VI. Conclusion

In 2013, the Cockpit Motion Facility (CMF) at NASA Langley Research Center performed a review of its modeling and motion cueing of cockpit motion under turbulence. This paper covers one step of that review, revisiting the implementation and verification of the code that models the continuous turbulence. The code, part of the Langley Standard Real-time Simulation in C++ (LaSRS++), uses the Dryden spectra for linear gusts and the MIL-F-8785 spectra for gust gradients. The LaRC review of the turbulence model consisted of both a code review and a functional verification. The code review compared the turbulence model code against algorithmic descriptions of the Dryden gust and the MIL-F-8785 gust gradient found in the literature. The functional verification assessed the time series output of the code using three criteria: 1) the mean of the time series should be zero, 2) the variance of the time series should match the user-defined total turbulence energy, and 3) the time series should reproduce the turbulence spectra.

The code review did find that the code largely adhered to the representation of the Dryden and MIL-F-8785 models found MIL-F-8785B [4] and Ref. [14]. (See Section III for a description of the models and for details of the digital algorithm in the LaSRS++ code.) However, the review produced two major observations. First, the pre-review code pulls out the numerator of the first filter as if it were a constant but this term contains two parameters that can vary with time, namely the vehicle speed (V) and the turbulence characteristic length (L). As a result, the code was modified to retain the term in the filter execution. Second, the pre-review code used a first-order hold (FOH) to model the white noise input when digitizing the turbulence filters. However, a zero-order hold (ZOH) is the better representation for uncorrelated, random inputs. References [14] and [15] explicitly identify the use of the zero-order hold for discretizing the turbulence filters. Moreover, Ref. [16] mathematically determined that the first-order hold should produce a greater loss in the turbulence energy present in the resulting time series. Therefore, the functional verification would assess the outcome of selecting between FOH and ZOH in digitizing the turbulence filters.

Before proceeding with the functional verification, the method of verification was updated and improved. Past attempts to reconstruct the spectra from a time series output of the turbulence model code used an arbitrarily selected number of points ranging from several thousand to several hundred thousand points. This paper presents a new method that determines the number of points derived from the need to reproduce defined characteristic points of the spectra and to target a high confidence bound. This method first determines the size of the discrete Fourier transform (DFT) required to capture the characteristic points of the linear gust spectra for a given test case. That analysis revealed that the minimum resolution of the DFT in rad/s needs to be ¼ V/L resulting in a minimum DFT length of $8\pi f L/V$ rounded up to the nearest power of two. The functional verification was performed using quadruple the minimum DFT size in order to capture some of the curvature about the characteristic points. Next, the method determines the number of DFTs to produce an averaged periodogram within a specified confidence bounds about the spectra function. The functional verification used a 95% confidence level for a bounds of 1/3 standard deviation, which requires 36 independent samples. However, each DFT cannot be treated as independent because the turbulence filters correlate data over a period of L/V; thus, only sets of DFTs of size fL/V, rounded up to the next integer, can be treated as independent. This results in an expression for the minimum number of time series points required to reproduce the spectra as a function of f, L, V, and the desired confidence interval. The full expression is shown in Eq. (14). In addition to better quantifying the number of points required in the time series, the test cases were expanded to cover a representative range of vehicle speeds and simulation execution frequencies at altitudes in both the free atmosphere and in the planetary boundary layer.[***] A total of 54 test cases were developed to represent the subsonic flight of a variety of aircraft under cruise, approach, and landing conditions. The mean, variance, and a periodogram were calculated from the time series for each of the 54 cases.

The functional verification was performed on four digital implementations of the turbulence filters. Three of the implementations factor the turbulence filters into a series of first order filters. The FOH All Filters implementation, taken from the pre-review code, applies a FOH on each filters in a series. The ZOH All Filters implementation applies a ZOH to each filter in a series. The ZOH First Filter implementation applies the ZOH to the first filter in a series and applies a FOH to subsequent filters. Finally, the ZOH No Factoring implementation doesn't factor the linear turbulence filters and applies a ZOH to each filter. All implementations succeed in producing a near zero mean for the turbulence outputs. However, differences are seen in the variance of the resulting time series from each implementation, and these differences grow with increasing $(V\,\Delta t)/L$ among the test cases. The FOH All Filters implementation exhibits a loss of variance that reaches 7%. The ZOH All Filters implementation exhibits an over prediction of the set variance that reaches 3%. However, both the ZOH No Factoring and the ZOH First Filter implementations produce almost no loss of variance across all test cases. As a result, the LaSRS++ code was modified to use the ZOH First Filter

[***] Section III.D provides a cautionary note that running the turbulence models at typical simulation execution rates may produce high frequency noise in the vehicle response, especially for large wing-tail transports.

implementation. The periodograms produced from the time series of the ZOH First Filter implementation also showed good agreement with the turbulence spectra functions. The periodograms do exhibit an over-prediction of the power contribution from high frequencies. However, this over-prediction appears to increase the turbulence energy of the frequency range just enough to counteract the energy lost due to the frequency truncation that results from running the filters at a fixed execution rate. Examining periodograms from the time series of the FOH First Filter implementation reveals the nature of its variance loss. The periodograms show a profound under prediction of power contributions from higher frequencies. That under prediction increases with increasing ($V\ \Delta t$) for the test cases at a given altitude (and, therefore, a given L). These losses combine with the energy loss due to frequency truncation, which also correlates with increasing ($V\ \Delta t$), to produce the variance loss that increases with ($V\ \Delta t$)/L. Overall, the functional verification provided compelling evidence that the improved turbulence code using ZOH First Filter implementation is an accurate simulation of the Dryden turbulence model and the MIL-F-8785 gust gradient model.

References

[1] Ercole, A. V., Cardullo, F. M., Zaychik, K., Kelly, L., and Houck, J., "Motion Cueing Algorithm Modification for Improved Turbulence Simulation," AIAA-2009-6247, AIAA Modeling and Simulation Technologies Conference, Chicago, IL, August 10-13, 2009.

[2] Ercole, A. V., Cardullo, F. M., Kelly, L., and Houck, J., "Algorithm for Simulating Atmospheric Turbulence and Aeroelastic Effects on Simulator Motion Systems," NASA/TM–2012-217601, NASA Langley Research Center, Hampton, VA, August 2012.

[3] R. Leslie, D. Geyer, K. Cunningham, M. Madden, P. Kenney, and P. Glaab, "LaSRS++: An Object-Oriented Framework for Real-Time Simulation of Aircraft," AIAA Modeling and Simulation Technologies Conference and Exhibit, August 1998, Boston, Massachusetts, AIAA-98-4529.

[4] "Flying Qualities of Piloted Airplanes," US Military Specification MIL-F-8785B, August 7, 1969.

[5] Chalk, C. R., Neal, T. P., Harris, T. M., Pritchard, F. E., and Woodcock, R. J., "Guide for MIL-F-8785B(ASG), 'Military Specification-Flying Qualities of Piloted Airplanes'," Technical Report AFFDL-TR-69-72, Air Force Dynamics Laboratory, August 1969.

[6] "Flying Qualities of Piloted Airplanes," US Military Specification MIL-F-8785C, November 5, 1980.

[7] "Flying Qualities of Piloted Aircraft," Department of Defense Handbook MIL-HDBK-1797, December 19, 1997.

[8] Etkin, B., "Turbulent Wind and Its Affect on Flight," AIAA-80-1836R, Journal of Aircraft, Vol. 18, No. 5, May 1981, pp. 327-345.

[9] "Dryden Wind Turbulence Model (Continuous)," The Mathworks, Inc., URL: https://www.mathworks.com/help/aeroblks/drydenwindturbulencemodelcontinuous.html [retrieved 25 April 2018].

[10] Etkin, B. "Theory of the Flight of Airplanes in Isotropic Turbulence - Review and Extension," AGARD Report 372, Advisory Group for Aeronautics Research and Development, Paris, France, April 1961.

[11] Park, S. K. and Miller, K. W., "Random number generators: good ones are hard to find," Communications of the ACM Vol. 31, Issue 10, October 1988, pp. 1192-1201. doi:10.1145/63039.63042

[12] Park, S. K., Miller, K. W., and Stockmeyer, P.K., "Technical Correspondence: Response" Communications of the ACM. Vol. 36 Issue 7, July 1993, pp. 108–110. doi:10.1145/159544.376068.

[13] "Psuedo random number generation," cppreference.com, URL: http://en.cppreference.com/w/cpp/numeric/random [retrieved 26 April 2018]

[14] McFarland, R., "A Standard Kinematic Model for Flight Simulation at NASA-Ames," NASA CR-2497, Computer Sciences Corporation, January 1975.

[15] Neuman, F. and Foster, J. D., "Investigation of a Digital Automatic Aircraft Landing System in Turbulence", NASA-TN-D-6066, NASA, Washington, DC, October 1970.

[16] Martin, D. J., "Real-Time simulation of Atmospheric Turbulence," Thesis, College of William and Mary, Williamsburg, VA, May 1979.

[17] Madden, M. M., "Improving Modeling and Motion Cueing of Flight in Turbulence," Proceedings of the Flight Simulation Conference: Striving for Effectiveness in Flight Simulation, Royal Aeronautical Society, London, 13-14 June 2017, URL: http://aerosociety.sds06.websds.net/2017/sefs/pages/downloads.php?s=7

[18] Etkin, B., Dynamics of Atmospheric Flight, John Wiley & Sons, Inc. New York, 1972, ISBN 0-471-24620-4.

[19] Barr, N. M., Gangsaas, D., and Schaeffer, D. R., "Wind Models for Flight Simulator Certification of Landing and Approach Guidance and Control Systems," FAA-RD-74-206, Federal Aviation Administration, Washington, D. C., 1974.

[20] Beal, T. R., "Digital Simulation of Atmospheric Turbulence for Dryden and von Karman Models," Journal of Guidance, Control, and Dynamics, Vol. 16, No. 1, January – February, 1993, pp. 132-138.