Theoretical Basis for the Resolution and Noise of SWOT Estimates of Sea-Surface Height

Dudley B. Chelton and Roger M. Samelson
Oregon State University

and

J. Thomas Farrar
Woods Hole Oceanographic Institution

May 14, 2015

1. Introduction

This note clarifies two important issues related to ocean observations of sea-surface height (SSH) by SWOT. The analysis presented here is intended to supplement the discussion in the SWOT science requirements document[1], the onboard processing document[2], and the mission performance and error budget document[3].

The first issue is the question of how the raw SWOT data will be filtered by the onboard processor and the related questions of what the resolution and decorrelation length scale are for the onboard pre-processed estimates of SSH. This information is provided in the onboard processing document[2], but without the mathematical basis for the conclusions. A rigorous development of the mathematical formalism for assessing the resolution and decorrelation length scale is presented in Section 2, along with a discussion of the error reduction that can be achieved by smoothing of the onboard pre-processed SWOT data in ground-based post-processing. Section 2 also includes a clarification of the relation between the feature resolution of smoothed SSH fields as defined in the onboard processing document[2] and the wavelength resolution corresponding to the half-power filter cutoff that is often used by oceanographers to characterize resolution.

The second issue is the question of the science requirement for the variance of the uncorrelated errors in the pre-processed estimates of SSH computed onboard the SWOT spacecraft. While this is important information for users since these data are fundamental to all oceanographic applications of SWOT data, the requirement for the noise variance of the onboard estimates of SSH is not explicitly specified in the present drafts of the SWOT documentation[1,2,3]. The analysis in Section 3 clarifies the relationship between the uncorrelated errors in the onboard pre-processed estimates of SSH and the error specification as characterized in the SWOT documentation. The latter is expressed as a white-noise spectrum only for wavelengths longer than 15 km over which the estimated signal-to-noise ratio is required to exceed unity for 68% of the world ocean (see the cyan line in Figure 1). This uncorrelated error specification is furthermore based on hypothetical smoothing of the onboard SWOT data with an idealized 2-dimensional filter that eliminates all variability with scales shorter than 15 km. The 1-dimensional (e.g., along-track) white-noise spectral representation of the residual uncorrelated measurement errors after this hypothetical filtering is shown by the dashed red line in Figure 1.

Since idealized filtering is not achievable in practice, the science requirement specification of SWOT measurement accuracy cannot be tested rigorously in post-launch verification. The instrument performance could be assessed from SWOT data smoothed with a realizable filter that has the same 15-km filter cutoff wavelength as the idealized filter. Such an assessment could only be qualitative, however, since the wavenumber content of the filtered data will be subject to the limitations of the gradual rolloff of the filter transfer function of the realizable filter. A quantitative assessment of SWOT measurement accuracy can only be made from the onboard pre-processed SWOT estimates of SSH without applying additional filtering. Moreover, the variance of the uncorrelated errors in the onboard SWOT data without filtering is essential information to enable users to make informed decisions about their choice of filtering in post-processing of the SWOT data to reduce the noise.

The white-noise spectrum of the uncorrelated errors in the onboard pre-processed SWOT data that is derived in Section 3d from the error specification with hypothetical filtering is shown by the dashed blue line in Figure 1. The RMS value of the uncorrelated errors in the onboard pre-processed SWOT data derived from this white-noise spectrum is 2.74 cm. Since this information is not explicitly provided in the SWOT documentation[1,2,3], this result and the discussion of its theoretical basis are perhaps the most important contributions of the analysis presented here.

In addition to uncorrelated errors, SWOT data are also contaminated by spatially correlated measurement errors from a variety of sources, including orbit errors, sea-state bias errors and environmental corrections of the two-way travel time of the altimetric range estimates. The spectrum of these long-wavelength measurement errors is expressed in the SWOT documentation[1,2,3] in terms of the hypothetical ground-based post-processed SWOT data after idealized smoothing with a filter cutoff wavelength of 15 km (see the dotted red line in Figure 1). The spectrum of total measurement errors in the hypothetically smoothed SWOT data is the sum of the spectra of uncorrelated errors and long-wavelength errors, which is shown by the solid red line in Figure 1. The associated spectrum of total measurement errors in the onboard pre-processed SWOT data without smoothing that is deduced in Section 3d is shown by the solid blue line. For reference, the mean and 68 percentile of the distribution of global SSH signal spectra that are presented in the SWOT science requirements document[1] are shown by the black and cyan lines, respectively.

The various sources of spatially correlated measurement errors have wavelengths much longer than the submesoscale and mesoscale variability with wavelengths shorter than ∼100 km that is the primary oceanographic emphasis of the SWOT mission. These long-wavelength measurement errors can thus be largely removed by spatially high-pass filtering the SWOT data in the along-track direction. The dominant source of measurement errors on scales shorter than ∼100 km is the spatially uncorrelated errors that are addressed in this note.

The uncorrelated SWOT measurement errors are treated here as if they are constant regionally and temporally. In reality, the measurement errors vary somewhat across the measurement swaths, with higher values near the edges of each swath. They also increase with increasing significant wave height, which varies regionally and temporally. As in the science requirements document[1], the uncorrelated measurement errors derived in this note for the onboard pre-processed 1 km × 1 km SWOT estimates of SSH are for the global, time-averaged estimates across the inner 50 km of the measurement swaths for conditions of 2-m significant wave height.

2. Filtering and Resolution

The raw SWOT estimates of SSH at the full resolution of ∼10 m of the Ka-band Radar Interferometer (KaRIn) instrument are far too noisy to be useful in oceanographic applications. To reduce the measurement noise, while at the same time reducing the unnecessarily large data volume over the oceans, the present plan[2] is to smooth the raw data onboard the spacecraft to achieve a resolution of 1 km. The meaning of "resolution" as defined in the SWOT onboard processing document[2] is discussed in Section 2c.

There are many possible filters that could be used for this smoothing. In one dimension, every filter can be characterized by its filter cutoff wavenumber k_c, which is defined below in Section 2a in terms of the filter transfer function of the smoother. The ideal filter has a filter transfer function that consists of values of 1 for wavenumbers $|k| \le k_c$ and 0 for higher wavenumbers. In practice, the filter transfer functions of real filters decrease gradually with increasing wavenumber across the low-wavenumber pass band and roll off steeply through the filter cutoff wavenumber k_c. Most real filters also have undesireable side lobes at wavenumbers higher than k_c. The simple uniform-weighted running average smoother considered in Section 2a has the largest side lobes of any low-pass filter.

As described in detail in the onboard processing document[2], the SWOT Project Office plans to smooth the raw SWOT data differently in the along-track and cross-track directions. Specifically, along-track smoothing will be applied using a Blackman-Harris smoother, and cross-track smoothing will be applied using the Parzen smoother that is defined in Section 2b. The parameters of each of these smoothers were carefully chosen by the SWOT Project Office to have essentially identical filtering properties. The filter transfer functions of both of these smoothers have side lobes that are much smaller than those of the uniform-weighted running average. The rationale for different smoothing in each dimension is not clear. The Blackman-Harris smoother has better side lobe suppression (see Figure 32 in the onboard processing document[2]; see also Harris, 1978[4]), but not enough so to make any significant, or even detectable, difference in the smoothed SSH values. The 2-dimensional smoothing of the raw SWOT data could thus be achieved equally well for all intents and purposes by smoothing with a Parzen window in both dimensions. It is therefore sufficient to consider only the Parzen smoother in this note. The filtering properties of the Parzen smoother are examined in detail in Sections 2b–e.

2a. The Running Average Smoother

The Parzen smoother to be used for cross-track smoothing in the SWOT onboard processor (see Sections 2b and 2c) is most easily understood by first considering the simple running average smoother. For simplicity, consider the case of a 1-dimensional spatial series that is continuous. The results can be extended straightforwardly to the case of 2-dimensional "block averaging." The results can also be extended to discrete sampling, but the equations become more cumbersome.

A running average with a span of L_1 applied to a spatial series $h(x)$ can be written as the convolution

$$h_1(x) = w_1(x) * h(x) \equiv \int_{-\infty}^{\infty} w_1(x-s)\, h(s)\, ds, \tag{2.1}$$

where the weighting function is

$$w_1(x) = \frac{1}{L_1}\,\Pi\left(\frac{x}{L_1}\right) \tag{2.2a}$$

and $\Pi(x/L_1)$ is shorthand notation for the rectangle function defined by

$$\Pi\left(\frac{x}{L_1}\right) \equiv \begin{cases} 1 & \text{if } -L_1/2 \le x \le L_1/2 \\ 0 & \text{otherwise.} \end{cases} \tag{2.2b}$$

The output (2.1) of the running average smoother with weighting function (2.2) is thus

$$h_1(x) = \frac{1}{L_1}\int_{-\infty}^{\infty} \Pi\left(\frac{x-s}{L_1}\right) h(s)\,ds = \frac{1}{L_1}\int_{x-L_1/2}^{x+L_1/2} h(s)\,ds. \tag{2.3}$$

The subscript 1 distinguishes this uniform-weighted running average from the Parzen weighted running average that is defined later in Section 2b.

The filter transfer function of any smoother is defined to be the Fourier transform of the weighting function associated with the smoother. For the case of the uniform-weighted running average (2.3), the Fourier transform of the weighting function (2.2) is

$$W_1(k) = \int_{-\infty}^{\infty} w_1(x)\,e^{-i2\pi kx}\,dx = \frac{1}{L_1}\int_{-L_1/2}^{L_1/2} e^{-i2\pi kx}\,dx = \frac{2}{L_1}\int_0^{L_1/2} \cos(2\pi kx)\,dx = \mathrm{sinc}(kL_1), \tag{2.4a}$$

where k is wavenumber and $\mathrm{sinc}(kL_1)$ is shorthand notation for

$$\mathrm{sinc}(kL_1) \equiv \frac{\sin(\pi kL_1)}{\pi kL_1}\,. \tag{2.4b}$$

Note that the multiplicative factor π in the numerator and denominator on the right side of (2.4b) is implicit in the definition of $\mathrm{sinc}(kL_1)$.

The wavenumber contents of the unfiltered and filtered data are defined by their Fourier transforms, which can be denoted as $H(k)$ and $H_1(k)$, respectively. By the Convolution Theorem, the convolution integral (2.1) in the space domain can be expressed as multiplication in the wavenumber domain. The wavenumber content $H_1(k)$ of the filtered output $h_1(x)$ is thus related to the wavenumber content $H(k)$ of the unfiltered data $h(x)$ by

$$H_1(k) = W_1(k)H(k)\,. \tag{2.5}$$

The filter transfer function $W_1(k)$ thus determines the wavenumber content of the filtered output.

The filtering properties of any smoother can be characterized by a filter cutoff wavenumber that is defined to be the wavenumber at which the squared value of the filter transfer function has a value of 0.5. This half-power filter cutoff wavenumber k_c for the filter transfer function (2.4) of the uniform-weighted running average smoother is thus the value of k_c for which

$$W_1^2(k_c) = \mathrm{sinc}^2(k_cL_1) = 0.5 \qquad \Rightarrow \qquad k_c = \frac{0.443}{L_1}\,. \tag{2.6}$$

The weighting function (2.2) and the squared value of the filter transfer function (2.4) of the uniform-weighted running average are shown for the case of $L_1 = 1$ km in Figure 2. A highly undesirable feature of the simple running average smoother with uniform weighting is that its filter transfer function has large side lobes outside of the main lobe that defines the range of wavenumbers $|k| \leq k_c$ that are of interest.

2b. The Parzen Smoother

The large side lobes of the filter transfer function of the uniform-weighted running average examined in Section 2a can be suppressed by replacing the uniform weighting with a tapered weighting. A simple approach to tapering is to apply a succession of uniform-weighted running averages. For example, the weighting function for two passes of a uniform-weighted running average with a span of L_1 can be shown to be equivalent to a single triangular weighted running average with a span of $L_2 = 2L_1$. More generally, p passes of a uniform-weighted running average with a span of L_1 are equivalent to a single weighted running average with a span of $L_p = pL_1$ and a weighting function that consists of a piecewise continuous polynomial of order $p - 1$.

Mathematically, the weighting function $w_p(x)$ of p passes of a uniform-weighted running average with a span of L_1 consists of p convolutions of the rectangle weighting function (2.2) of the uniform-weighted running average. The Fourier transform of this smoother (i.e., its filter transfer function) is therefore easily determined from the Convolution Theorem to be the multiplicative product of p of the filter transfer functions (2.4) of the uniform-weighted running average,

$$W_p(k) = \mathrm{sinc}^{\,p}(kL_1)\,. \tag{2.7}$$

The extremum of the j^{th} side lobe of this filter transfer function has a value of $\mathrm{sinc}^{\,p}(k_j L_1)$, where $k_j = (j + 1/2)/L_1,\ j = 1, 2, 3, \ldots$. Since $\mathrm{sinc}(k_j L_1) = (-1)^j / [(j + 1/2)\pi]$ has a magnitude less than 1 for all j, the side lobe extrema decrease rapidly in magnitude both with increasing j and increasing p. For the case of the dominant $j = 1$ side lobe, the sinc function has a value of $\mathrm{sinc}(1.5) = \sin(1.5\pi)/(1.5\pi) = -0.212$. The dominant side lobe of the filter transfer function for p passes of the uniform-weighted running average thus has a magnitude of $0.212^{\,p}$ and a squared magnitude of $0.212^{\,2p}$.

There is no practical advantage to smoothing with more than $p = 4$ passes of a uniform-weighted running average since the extremum of the dominant side lobe of its squared filter transfer function is 4.08×10^{-6}, i.e., an attenuation factor of more than 50 db. The quadruple running average is called the Parzen smoother. Mathematically, the weighting function of the Parzen smoother can be written as a quadruple convolution of the weighting function (2.2) of the uniform-weighted running average,

$$w_4(x) = \frac{1}{L_1^4}\left[\Pi\left(\frac{x}{L_1}\right) * \Pi\left(\frac{x}{L_1}\right) * \Pi\left(\frac{x}{L_1}\right) * \Pi\left(\frac{x}{L_1}\right)\right]. \tag{2.8}$$

The output of the Parzen smoother is the convolution

$$h_4(x) = w_4(x) * h(x) = \int_{-\infty}^{\infty} w_4(x - s)\, h(s)\, ds\,. \tag{2.9}$$

The subscript 4 signifies that the Parzen smoother is equivalent to 4 passes of a uniform-weighted running average. The filter transfer function (2.7) with $p = 4$ for the Parzen weighting function is

$$W_4(k) = \mathrm{sinc}^4(kL_1) = \mathrm{sinc}^4\left(\frac{kL_4}{4}\right), \tag{2.10}$$

where $L_4 = 4L_1$ and L_1 is the span of each of the four passes of the uniform-weighted running average. The significance of L_4 becomes apparent below. The filter cutoff wavenumber of the Parzen smoother, defined as in (2.6) to be the wavenumber k_c at which the squared value of the filter transfer function (2.10) has a value of 0.5, is

$$W_4^2(k_c) = \mathrm{sinc}^8(k_c L_1) = \mathrm{sinc}^8(k_c L_4/4) = 0.5 \qquad \Rightarrow \qquad k_c = \frac{0.228}{L_1} = \frac{0.910}{L_4}\,. \tag{2.11}$$

While it would be very tedious to calculate the weights of the Parzen smoother from the quadruple convolution (2.8), it is straightforward with the help of integral tables to determine these weights from the inverse Fourier transform of the filter transfer function (2.10). The result is the piecewise cubic polynomial defined by

$$w_4(x) = \begin{cases} \dfrac{8}{3L_4}\left(1 - \dfrac{24x^2}{L_4^2} + \dfrac{48|x|^3}{L_4^3}\right) & \text{if } 0 \le |x| \le L_4/4 \\ \dfrac{8}{3L_4}\left(2 - \dfrac{12|x|}{L_4} + \dfrac{24x^2}{L_4^2} - \dfrac{16|x|^3}{L_4^3}\right) & \text{if } L_4/4 \le |x| \le L_4/2 \\ 0 & \text{if } |x| > L_4/2. \end{cases} \tag{2.12}$$

It is thus apparent that the Parzen smoother has a full span of L_4 that is defined above to be four times larger than the span L_1 of each of the four passes of the uniform-weighted smoother upon which it is based.

In principle, the Parzen smoother (2.9) can be implemented either as a succession of four passes of the uniform-weighted running average (2.3) with span L_1 or as a single pass of (2.9) with the piecewise cubic weighting (2.12) and a span of $L_4 = 4L_1$. In practice, however, application to discretely sampled data results in small differences in the filtered output by the two procedures. When applied to data with a uniformly spaced sample interval of Δx, for example, an arbitrarily specified value of L_4 can result in a value of $L_1 = L_4/4$ that is a non-integer multiple of Δx. Moreover, it is desirable for the smoothing span to be an odd multiple of Δx so that the weighting function is symmetric. This assures that there is no phase shift at any of the wavenumbers in the filtered output. It is clearly not possible for both L_1 and L_4 to be odd integer multiples of Δx. The preferred implementation of the Parzen smoother with uniformly spaced data is therefore as a single pass of (2.9) with the piecewise cubic weighting (2.12) and L_4 equal to an odd integer multiple of Δx.

The weighting function (2.12) and the square of the filter transfer function (2.10) of the Parzen smoother are shown for the case of $L_4 = 1$ km in Figure 3. This is equivalent to a span of $L_1 = L_4/4 = 0.25$ km for each of the four passes of the uniform-weighted running average (2.3).

For a given value of L_1, it can be noted from (2.6) and (2.11) that the filter cutoff wavenumber k_c is a factor of $0.228/0.443 \approx 1/2$ smaller for the Parzen smoother than for the single-pass uniform-weighted running average. The Parzen smoother is thus a stronger low-pass filter, i.e., it restricts the filtered output to lower wavenumbers and hence longer wavelengths.

To achieve the same half-power filter cutoff wavenumber k_c with both the uniform-weighted running average and the Parzen smoother, the span L_4 can be adjusted accordingly. A Parzen smoother with a span of $L_4 = 2$ km, for example, has a half-power filter cutoff wavenumber of $k_c = 0.455$ cpkm (cycles per km). As shown in the middle panel of Figure 4, this is very close to the half-power filter cutoff wavenumber of $k_c = 0.443$ cpkm for the uniform-weighted running average with a span of $L_1 = 1$ km that was shown previously in Figure 2. Note again the much smaller side lobes of the filter transfer function of the Parzen smoother compared with those of the uniform-weighted running average. This improved side lobe suppression comes at the price of somewhat more gradual roll off through the filter cutoff wavenumber k_c of the filter transfer function of the Parzen smoother (see the middle panel of Figure 4).

2c. Onboard Pre-Processing of SWOT Data

The present plan[2] is to smooth the raw SWOT data in the onboard processor to achieve the science requirement of 1 km × 1 km resolution for ocean observations. (The accuracy requirement for these onboard pre-processed estimates of SSH is derived in Section 3d.) As noted previously, this resolution will be achieved using a Blackman-Harris smoother in the along-track direction and a Parzen smoother in the cross-track direction. The parameters of these two smoothers have been chosen by the SWOT Project Office so that their filter transfer functions are nearly identical. The filtering properties of the onboard pre-processed estimates of SSH are therefore discussed here in the context of the Parzen smoother that was examined in detail in Section 2b.

The resolution of smoothed estimates of SSH could be defined in a variety of ways. The definition in the SWOT onboard processing document[2] is based on the lagged autocorrelation function associated with the Parzen smoother. This autocorrelation function can be determined by noting that the square of the filter transfer function (2.10) of the Parzen smoother is proportional to the power spectral density of the low-pass filtered output of the Parzen smoother applied to data consisting of uncorrelated ("white") noise [see (2.21b) in Section 2e below with the white noise spectrum of unfiltered onboard SWOT data on the right side of the equation defined by (2.23)]. The autocorrelation of this filtered output can be obtained as the inverse Fourier transform of its power spectral density, i.e., as the inverse Fourier transform of the squared value of the filter transfer function (2.10) of the Parzen smoother.

In the x dimension, the lagged autocorrelation function for a Parzen smoother with a span of L_4 is thus given by

$$R(x) = \int_{-\infty}^{\infty} \mathrm{sinc}^8\left(\frac{kL_4}{4}\right) e^{i2\pi kx}\, dk = 2\int_{0}^{\infty} \mathrm{sinc}^8\left(\frac{kL_4}{4}\right) \cos(2\pi kx)\, dk\,. \tag{2.13}$$

This can be evaluated with the help of integral tables. The result is a piecewise continuous 7^{th}-order polynomial that is symmetric about lag $x = 0$. This is shown for positive lags x in the bottom panel of Figure 4 for a span of $L_4 = 2$ km.

The scale of features that can be resolved in the smoothed data is based somewhat subjectively on the lag at which the autocorrelation decays to a value of 0.5. For a Parzen smoother with a span of $L_4 = 1$ km, it can be shown that an autocorrelation of 0.5 occurs at a lag of $x = 0.248$ km. Since the autocorrelation is symmetric about zero lag, this can be interpreted as the radius of features that can be resolved. The feature diameter resolution after filtering with a Parzen smoother with a span of $L_4 = 1$ km is therefore approximately 0.5 km.

Doubling the span L_4 of the Parzen smoother doubles the lag at which the autocorrelation decays to a value of 0.5 (see the bottom panel of Figure 4), and hence doubles the feature resolution scale of the smoothed SSH estimates. The Parzen smoother with a span of $L_4 = 2$ km shown in Figure 4 thus has a feature diameter resolution capability of ∼1 km. It can be seen from the bottom panel of Figure 4 that the autocorrelation function for this Parzen smoother decays to a very small value of 0.050 at a lag of 1 km. Smoothed estimates of SSH with a Parzen smoothing span of $L_4 = 2$ km are therefore essentially uncorrelated at a lag of 1 km.

The requirement for the onboard pre-processing of SWOT data is for SSH estimates on a 1 km × 1 km grid to be statistically uncorrelated. From the discussion above, this can be achieved using a 2-dimensional Parzen smoother with a span of $L_4 = 2$ km in each dimension[5]. In the x dimension, the half-power filter cutoff wavenumber (2.11) for this span is $k_c \approx 0.5$ cpkm (see the middle panel of Figure 4), which is also the Nyquist wavenumber for a sample interval of $\Delta x = 1$ km. The resolution and sampling requirements for the onboard pre-processed SWOT estimates of SSH can thus be achieved using a Parzen smoother with a span of $L_4 = 2$ km in each dimension, and posting of the smoothed SSH estimates on a 1 km × 1 km grid[6].

2d. The Feature Diameter Resolution of SSH in Ground-Based Post-Processing of SWOT Data

Most oceanographic applications of SWOT data will require additional smoothing of the onboard pre-processed estimates of SSH to reduce the uncorrelated measurement noise. The formalism developed in Section 2c is extended in this section to characterize the feature resolution in SSH fields smoothed in ground-based post-processing with any desired half-power filter cutoff wavenumbers k_c and l_c in the x and y dimensions.

For the case of the Parzen smoother considered above, it was shown in Figure 4 that smoothing SSH in the x dimension with a span of $L_4 = 2$ km results in an autocorrelation of 0.5 at a lag of about 0.5 km and therefore a feature diameter resolution scale of about 1 km. More generally, the feature diameter resolution scale for a Parzen smoother in the x dimension with an arbitrary span of L_4 is

$$\text{Feature Diameter Resolution} \approx \frac{L_4}{2}\,. \tag{2.14a}$$

Since the half-power filter cutoff wavenumber (2.11) for a Parzen smoother with a span of L_4 is $k_c \approx L_4^{-1}$, the feature diameter resolution scale (2.14a) can be expressed equivalently in terms of the half-power filter cutoff wavenumber as

$$\text{Feature Diameter Resolution} \approx \frac{1}{2k_c}\,. \tag{2.14b}$$

In terms of half-power filter cutoff wavelength $\lambda_c = k_c^{-1}$, this is

$$\text{Feature Diameter Resolution} \approx \frac{\lambda_c}{2}\,. \tag{2.14c}$$

While the details of the filter transfer function differ for different smoothers, every smoother can be characterized by the half-power filter cutoff wavenumber k_c for a particular choice of the parameters of the filter (e.g., the span L_4 for the case of the Parzen smoother). The primary distinctions between different smoothers are in their abilities to reduce the large-amplitude side-lobes of the filter transfer function of the uniform-weighted running average considered in Section 2a. Regardless of the specific formulation of the smoothing procedure, the relationships (2.14b) and (2.14c) between feature diameter resolution and filter cutoff wavenumber k_c or wavelength λ_c developed for the Parzen smoother are approximately applicable to any smoother.

The preceding analysis clarifies the relation between the autocorrelation-based definition (2.14a) of feature diameter resolution and the half-power filter cutoff wavelength $\lambda_c = k_c^{-1}$ that is often used by oceanographers to characterize the resolution of filtered SSH fields. From (2.14c), the feature diameter resolution of SSH fields smoothed with a half-power filter cutoff wavenumber k_c is half as large as the wavelength resolution $\lambda_c = k_c^{-1}$. This definition of resolution is coarser by 25% than the feature diameter resolution of $0.4\,\lambda_c$ as defined by Chelton et al. (2011)[7].

2e. Reduction of Uncorrelated Error Variance in Ground-Based Post-Processing of SWOT Data

The variance of the uncorrelated errors in the onboard pre-processed SWOT estimates of SSH can be reduced to any desired level with sufficient smoothing in ground-based post-processing. The objective of this section is to quantify the reduction that can be expected for given choices of the half-power filter cutoff wavenumbers k_c and l_c of the smoothing in the x and y dimensions. The multiplicative reduction factor derived in this section is used later in Section 3e to quantify the residual uncorrelated error variance in SSH fields after smoothing the onboard pre-processed SWOT data in ground-based post-processing with arbitrary user-specified filter cutoff wavenumbers.

To illustrate the filtering of SWOT data in ground-based post-processing, the weighting function, squared filter transfer function and lagged autocorrelation function are shown by the green lines in Figure 5 for the case of a Parzen smoother in the x dimension with a span of $L_4 = 14$ km. The blue lines in Figure 5 correspond to a uniform-weighted running average with a span of $L_1 = L_4/2 = 7$ km. It can be seen that the half-power filter cutoff wavenumber and the lag at which the autocorrelation decays to a value of 0.5 for this uniform-weighted running average are essentially the same as those of the Parzen smoother with a span of $L_4 = 14$ km. Note again the superior performance of the Parzen smoother in terms of side lobe suppression.

A rigorous estimate of the amount by which filtering with the Parzen smoother reduces the variance of uncorrelated errors compared with block averaging (i.e., compared with the uniform-weighted running average) can be derived from the filter transfer functions (2.4) and (2.10) of the two smoothers. As discussed in Section 2c, present plans[2] are for the raw SWOT data to be filtered onboard the satellite and posted on a grid with dimensions $\Delta x \times \Delta y = 1$ km $\times$ 1 km on which the pre-processed estimates of SSH are statistically uncorrelated. Consider a 1-dimensional discretely sampled spatial series $\epsilon(x_m)$ of errors in the onboard estimates of SSH at M locations

$x_m = m\Delta x,\ m = 0, 1, \ldots, M-1$. The discrete Fourier transform of this spatial series of errors in the x dimension is

$$E(k_{m'}) = \frac{1}{M}\sum_{m=0}^{M-1} \epsilon(x_m)\, e^{-i2\pi k_{m'} m\Delta x} = \frac{1}{M}\sum_{m=0}^{M-1} \epsilon(x_m)\, e^{-i2\pi m' m/M}\,, \tag{2.16}$$

where $k_{m'} = m'\Delta k,\ m' = (-M/2+1), (-M/2+2), \ldots, 0, \ldots, M/2$ are the Fourier wavenumbers that are separated by the constant interval $\Delta k = (M\Delta x)^{-1}$ that is defined by the record length $M\Delta x$. The 1-dimensional wavenumber power spectral density is defined in terms of the discrete Fourier transform (2.16) to be

$$\Sigma_{1d}(k_{m'}) = \frac{1}{\Delta k}\, E^*(k_{m'})E(k_{m'}), \quad m' = (-M/2+1), (-M/2+2), \ldots, 0, \ldots, M/2\,. \tag{2.17}$$

The asterisk in this expression denotes the complex conjugate.

Since $\epsilon(x_m)$ is a real variable, it is easy to show that the spectrum (2.17) is symmetric, i.e., $\Sigma_{1d}(-k_{m'}) = \Sigma_{1d}(k_{m'})$. It is therefore conventional to double the 1-dimensional sample power spectral density values (2.17) at all but the wavenumbers corresponding to indices $m' = 0$ and $M/2$ and display the spectrum over only the positive wavenumbers. The index $m' = M/2$ corresponds to the Nyquist wavenumber $k_{\mathcal{N}} \equiv (2\Delta x)^{-1}$ that is the highest resolvable wavenumber for the discrete sample interval Δx. The 1-sided representation of the 1-dimensional wavenumber sample power spectral density, which will be denoted with a prime, is thus related to the 2-sided, 1-dimensional wavenumber sample power spectral density (2.17) by

$$\Sigma'_{1d}(k_{m'}) = \begin{cases} \Sigma_{1d}(k_{m'}) & \text{if } m' = 0 \text{ or } M/2 \\ 2\,\Sigma_{1d}(k_{m'}) & \text{otherwise.} \end{cases} \tag{2.18}$$

Analogous to the expressions (2.1) and (2.9) for the output after filtering with the uniform-weighted running average and the Parzen smoother, the output $\overline{\epsilon}(x_m)$ of an arbitrary linear filter applied to the errors $\epsilon(x_m)$ in the onboard estimates of SSH can be written as a convolution

$$\overline{\epsilon}(x_m) = w(x_m) * \epsilon(x_m)\,, \tag{2.19}$$

where $w(x_m)$ is the filter weighting function for the specific linear filter. By the Convolution Theorem, this can be written as multiplication in the wavenumber domain,

$$\overline{E}(k_{m'}) = W(k_{m'})E(k_{m'})\,, \tag{2.20}$$

where $\overline{E}(k_{m'})$ and $W(k_{m'})$ are the discrete Fourier transforms at wavenumber $k_{m'}$ of the filtered output $\overline{\epsilon}(x_m)$ and the weighting function $w(x_m)$ of the linear filter. The 1-sided power spectral density of the filtered output can thus be expressed in terms of the 1-sided power spectral density (2.17) of the unfiltered errors by[8]

$$\overline{\Sigma}'_{1d}(k_{m'}) \;=\; \frac{2}{\Delta k}\,\overline{E}^{\,*}(k_{m'})\overline{E}(k_{m'}) \tag{2.21a}$$

$$\;=\; W^*(k_{m'})\,W(k_{m'})\,\Sigma'_{1d}(k_{m'})\,. \tag{2.21b}$$

Parseval's Theorem for the sample variance σ_ϵ^2 of the uncorrelated errors in the 1 km $\times$ 1 km SWOT data can be expressed in terms of either the 2-sided, 1-dimensional power spectral density (2.17) or the 1-sided, 1-dimensional power spectral density (2.18) in the x dimension as

$$\sigma_\epsilon^2 = \sum_{\substack{m'=-M/2+1 \\ m'\neq 0}}^{M/2} \Sigma_{1d}(k_{m'})\,\Delta k \;=\; \sum_{m'=1}^{M/2} \Sigma'_{1d}(k_{m'})\,\Delta k\,. \tag{2.22}$$

Note that the zero wavenumber corresponding to the index $m' = 0$ is excluded from the sum (2.22). This is because the power spectral density at zero wavenumber is equivalent to the square of the mean value, which does not contribute to the sample variance. Since the errors $\epsilon(x_m)$ are uncorrelated on the $\Delta x = 1$ km grid (see Section 2c), the wavenumber sample power spectral density $\Sigma'_{1d}(k_{m'})$ is constant at all wavenumbers $k_{m'}$. Uncorrelated errors are therefore referred to as "white noise" in analogy with the constant electromagnetic spectrum of the color white. The constant white-noise spectral value obtained from (2.22) with $\Delta k = (M\Delta x)^{-1}$ is

$$\Sigma'_{1d}(k_{m'}) \;=\; \frac{\sigma_\epsilon^2}{(M/2)\Delta k} \;=\; \frac{\sigma_\epsilon^2}{k_{\mathcal{N}}}\,, \qquad m' = 0, 1, \ldots, M/2\,, \tag{2.23}$$

where $k_{\mathcal{N}} = (2\Delta x)^{-1}$ is the Nyquist wavenumber defined previously.

Parseval's Theorem for the sample variance $\overline{\sigma}_\epsilon^2$ of the residual uncorrelated errors $\overline{\epsilon}(x_m)$ after smoothing of the onboard SWOT data can similarly be expressed in terms of its 1-sided, 1-dimensional power spectral density (2.21) as

$$\overline{\sigma}_\epsilon^2 = \sum_{m'=1}^{M/2} \overline{\Sigma}'_{1d}(k_{m'})\,\Delta k\,.$$

Substituting (2.21b) and (2.23) into the right side of this equation gives

$$\overline{\sigma}_\epsilon^2 = \sum_{m'=1}^{M/2} W^*(k_{m'})\,W(k_{m'})\,\Sigma'_{1d}(k_{m'})\,\Delta k = \frac{\sigma_\epsilon^2}{k_{\mathcal{N}}} \sum_{m'=1}^{M/2} W^*(k_{m'})W(k_{m'})\,\Delta k\,. \tag{2.24a}$$

It is readily seen from (2.4) and (2.10) that the filter transfer functions $W_1(k)$ and $W_4(k)$ for the cases of the uniform-weighted running average and the Parzen smoother are both real. This is a general property of symmetric smoothers. For any symmetric linear filter, (2.24a) thus becomes

$$\overline{\sigma}_\epsilon^2 = \frac{\sigma_\epsilon^2}{k_{\mathcal{N}}} \sum_{m'=1}^{M/2} W^2(k_{m'})\,\Delta k\,. \tag{2.24b}$$

Equation (2.24b) can be interpreted as a discretized statement that the sample variance of smoothed white-noise errors is equal to the integrated area under the squared filter transfer function. This is easily seen by considering the limit as the record length $M\Delta x$ approaches infinity. The wavenumber interval $\Delta k = (M\Delta x)^{-1}$ then becomes an infinitesimally small value dk and the

discrete wavenumbers $k_{m'}$ become continuous. The subscript m' can then be dropped and the discrete summation becomes an integral so that (2.24b) can be written as

$$\overline{\sigma}_\epsilon^2 = \frac{\sigma_\epsilon^2}{k_\mathcal{N}} \int_0^{k_\mathcal{N}} W^2(k)\, dk\,. \tag{2.25}$$

The variance of the discretely sampled smoothed white-noise errors $\overline{\epsilon}(x_m)$ with infinitely long record length $M\Delta x$ is thus proportional to the area under the squared value of the filter transfer function integrated from the zero wavenumber to the Nyquist wavenumber $k_\mathcal{N} = (2\Delta x)^{-1}$. The Nyquist wavenumber that defines the upper bound of the integral is imposed by the discrete sample interval Δx, regardless of the record length $M\Delta x$. [As noted above, the zero wavenumber at the lower bound of the integral (2.25) must be excluded from the discretized form (2.24b) of this integral since it does not contribute to the sample variance.]

It is well known that block averaging of L uncorrelated errors $\epsilon(x_m)$ in uniform-weighted running averages with a span of $L_1 = L$ reduces the uncorrelated error variance σ_ϵ^2 by a factor of L. It can be anticipated from the close agreement of the half-power filter cutoff wavenumbers of the two filter transfer functions in Figure 5 that approximately the same variance reduction factor of L can be achieved using a Parzen smoother with a span of $L_4 = 2L$. An approximate analytical expression for the ratio of the variances $\overline{\sigma}_1^2$ and $\overline{\sigma}_4^2$ of, respectively, uniform-weighted and Parzen-weighted running averages of uncorrelated white noise can be derived from (2.25) by considering the ratio

$$R \;=\; \frac{\overline{\sigma}_4^2}{\overline{\sigma}_1^2} \;=\; \frac{\displaystyle\int_0^{k_\mathcal{N}} W_4^2(k)\, dk}{\displaystyle\int_0^{k_\mathcal{N}} W_1^2(k)\, dk} \;=\; \frac{\displaystyle\int_0^{k_\mathcal{N}} \mathrm{sinc}^8(kL)\, dk}{\displaystyle\int_0^{k_\mathcal{N}} \mathrm{sinc}^2(kL)\, dk}\,, \tag{2.26}$$

where $W_1(k)$ and $W_4(k)$ are the filter transfer functions (2.4) and (2.10) of the uniform-weighted running average with a span of $L_1 = L$ and the Parzen smoother with a span of $L_4 = 4L$. The latter yields $\mathrm{sinc}^4(kL)$ on the right side of the expression (2.10) for the filter transfer function $W_4(k)$ of the Parzen smoother.

Exact analytical solutions cannot be obtained analytically for the two integrals on the right side of (2.26). However, solutions can be found in integral tables if the upper bound of the integrals is ∞ rather than the Nyquist wavenumber $k_\mathcal{N}$. The ratio (2.26) can thus be rewritten as

$$R \;=\; \frac{I_8 - \displaystyle\int_{k_\mathcal{N}}^{\infty} \mathrm{sinc}^8(kL)\, dk}{I_2 - \displaystyle\int_{k_\mathcal{N}}^{\infty} \mathrm{sinc}^2(kL)\, dk}\,, \tag{2.27a}$$

where

$$I_2 \;=\; \int_0^{\infty} \mathrm{sinc}^2(kL)\, dk \;=\; \frac{1}{2L} \tag{2.27b}$$

$$I_8 \;=\; \int_0^{\infty} \mathrm{sinc}^8(kL)\, dk \;=\; \frac{151}{630L} \tag{2.27c}$$

The solutions (2.27b) and (2.27c) for I_2 and I_8 are derived from the general solution obtained from integral tables for the integral from 0 to ∞ of $\mathrm{sinc}^p(kL)$ with $p = 2$ and 8.

Because the side lobes of $\mathrm{sinc}^8(kL)$ decay so much faster than the side lobes of $\mathrm{sinc}^2(kL)$ (see Figures 2–5), the correction term in the numerator of (2.27a) is much smaller than the correction term in the denominator. The ratio of variances of uncorrelated errors smoothed with the Parzen smoother and the uniform-weighted running average can therefore be approximated as

$$R \;\approx\; \frac{I_8}{I_2 - \displaystyle\int_{k_{\mathcal{N}}}^{\infty} \mathrm{sinc}^2(kL)\, dk}\,, \tag{2.28}$$

By approximating each of the side lobes of $\mathrm{sinc}^2(kL)$ as rectangles, it can be shown that the correction term in the denominator of this equation is approximately $\pi^{-2}L^{-2}$ for sufficiently large L. The ratio (2.28) of variances of Parzen-weighted and uniform-weighted running averages then becomes

$$R \;\approx\; \frac{I_8}{I_2 - \pi^{-2}L^{-2}}\,. \tag{2.29}$$

The analytical approximation (2.29) for R is shown by the dots in Figure 6 for values of L ranging from 2 km to 25 km. For comparison, the solid line corresponds to the exact solution obtained by numerical integration of the numerator and denominator of (2.26) with $k_{\mathcal{N}} = (2\Delta x)^{-1}$ and a sample interval of $\Delta x = 1$ km. The analytical approximation is slightly too large for $L = 2$ km but is indistinguishable from the numerical solution for $L > 2$ km. It is thus seen that the large L approximation (2.29) is applicable to all practical choices of smoothing of the 1 km $\times$ 1 km onboard pre-processed SWOT estimates of SSH. The ratio R is 0.5 for $L = 5$ km. It is slightly larger than 0.5 for $L < 5$ km and slightly smaller than 0.5 for $L > 5$ km. For all intents and purposes, however, R can be considered to have a value of 0.5 for any choice of $L \geq 2$ km. In other words, the residual variance σ_4^2 of white-noise errors smoothed using a Parzen smoother with a span of $L_4 = 4L$ is about half as large as the residual variance σ_1^2 of white noise errors smoothed using a uniform-weighted running average with a span of $L_1 = L$. Halving the span of the Parzen smoother to $L_4 = 2L$ doubles the variance σ_4^2, thus resulting in essentially the same variance σ_1^2 as the uniform-weighted running average with a span of $L_1 = L$. The Parzen smoother with span $L_4 = 2L$ has slightly higher variance for $L < 5$ and slightly lower variance for $L > 5$.

The above analysis in one dimension can be extended straightforwardly to the case of filtering in two dimensions. A Parzen smoother with a 2-dimensional span of $L_4(x) = L_x\Delta x$ in the x dimension and $L_4(y) = L_y\Delta y$ in the y dimension achieves approximately the same error variance reduction as block averaging of $L_x/2 \times L_y/2$ of the 1 km $\times$ 1 km onboard pre-processed estimates of SSH. The number of independent measurements in a Parzen smoothed estimate of SSH with spans of $L_x\Delta x \times L_y\Delta y$ is thus approximately $L_xL_y/4$. The multiplicative error variance reduction factor for the Parzen smoother is therefore

$$\text{Multiplicative Error Variance Reduction Factor} \;\approx\; \frac{4}{L_xL_y}\,. \tag{2.30a}$$

For the case shown in Figure 5, the Parzen smoother with a span of 14 km × 14 km achieves essentially the same error variance reduction as block averaging the 1 km × 1 km onboard pre-processed estimates of SSH over a 7 km × 7 km area, but with better filtering properties (see Figure 5). Since the onboard pre-processed estimates of SSH at 1 km × 1 km spacing are statistically uncorrelated (see Section 2c), the variance of their uncorrelated errors is reduced by a multiplicative factor of approximately $1/7^2 = 1/49$ in the SSH fields smoothed in ground-based post-processing using a Parzen smoother with a 2-dimensional span of 14 km × 14 km.

Because the half-power filter cutoff wavenumber (2.11) of the Parzen smoother is $k_c \approx (L_x \Delta x)^{-1}$ in the x dimension and $l_c \approx (L_y \Delta y)^{-1}$ in the y dimension, a 2-dimensional span of $L_x \Delta x \times L_y \Delta y$ can be expressed alternatively as approximately $k_c^{-1} \times l_c^{-1}$. This smoothing is analogous to block averaging of $(2\Delta x\, k_c)^{-1} \times (2\Delta y\, l_c)^{-1}$ of the 1 km × 1 km onboard pre-processed estimates of SSH. The reduction of uncorrelated error variance after smoothing the onboard estimates of SSH with a Parzen smoother with these half-power filter cutoff wavenumbers is thus

$$\text{Multiplicative Error Variance Reduction Factor} \approx 4\,\Delta x\,\Delta y\,k_c\,l_c. \tag{2.30b}$$

For smoothing with a half-power filter wavelength of $\lambda_c = k_c^{-1} = l_c^{-1}$ in both dimensions, this becomes

$$\text{Multiplicative Error Variance Reduction Factor} \approx \frac{4\,\Delta x \Delta y}{\lambda_c^2}. \tag{2.31}$$

The multiplicative error reduction factor (2.30b) derived above for the Parzen smoother provides an approximate characterization of the reduction of uncorrelated error variance for other smoothers with parameters chosen to give the same filter cutoff wavenumbers k_c and l_c. The residual uncorrelated error variance after filtering with arbitrary user-specified filter cutoffs is quantified in Section 3e after the white noise error variance σ_ϵ^2 of the onboard pre-processed SWOT data without smoothing is derived in Section 3d based on the science requirement specification in terms of hypothetical filtering with an idealized smoother.

3. Noise of SWOT Estimates of SSH

The uncorrelated errors of the onboard pre-processed SWOT estimates of SSH are most easily characterized in terms of an error variance. However, as noted in the Introduction and shown by the dotted red line in Figure 1 (see also Figure 7 below), the SWOT data are also subject to long-wavelength measurement errors with variance that increases with decreasing wavenumber (increasing wavelength). To be compatible with a spectral specification of long-wavelength measurement errors, the SWOT documentation characterizes the uncorrelated measurement errors as a white-noise spectrum rather than an error variance.

The manner in which the white-noise spectrum of measurement errors is specified in the SWOT documentation[1,2,3] is not useful to oceanographic users for two reasons. Firstly, the white noise errors are specified only for filtered SWOT data, and, moreover, for the case of idealized but unrealizable filtering. And secondly, this spectral representation of residual uncorrelated errors is not easily transformed into a variance that can be used to assess the limitations of 2-dimensional maps of SWOT data for investigation of submesoscale variability. To be able to quantify the

residual noise after user-specific smoothing of the SWOT data in ground-based post-processing, the uncorrelated error variance of the onboard pre-processed SWOT estimates of SSH must be specified. Converting the white noise spectrum as formulated in the the SWOT science requirement to a form that is useful to users requires a detailed and tedious derivation of the mathematical formalism for the relationships between the 2-dimensional wavenumber spectrum, the 1-dimensional (e.g., along-track) wavenumber spectrum and the variance. These relationships are presented in Sections 3a–c.

The analysis proceeds as follows. The spectral representation of uncorrelated error variance is derived in Section 3a in terms of the 2-dimensional wavenumber power spectral density, and in Section 3b in terms of the along-track 1-dimensional wavenumber power spectral density. The latter is the basis for the science requirement specification of uncorrelated measurement errors. This science requirement explicitly includes hypothetical smoothing of the onboard pre-processed SWOT estimates of SSH in ground-based post-processing with an idealized low-pass filter that has a specific half-power filter cutoff wavelength of $\lambda_c = k_c^{-1} = l_c^{-1} = 15$ km in both dimensions.

A general discussion of the spectral characteristics of smoothed SWOT data, and the resulting reduction of measurement noise, are presented in Section 3c. While an idealized filter cannot be implemented in practice, the hypothetical result allows a derivation in Section 3d of the uncorrelated error variance in the onboard pre-processed SWOT estimates of SSH. In particular, the systematic development in Sections 3a–c ultimately yields equation (3.14c) below that relates the uncorrelated error variance σ_ϵ^2 of the onboard pre-processed SWOT estimates of SSH on the right side of the equation to the along-track 1-dimensional wavenumber spectrum on the left side of the equation that is the science requirement specification after hypothetical idealized smoothing. This equation is used in Section 3d to determine the accuracy requirement for the unsmoothed onboard pre-processed SWOT estimates of SSH.

For practical use, a general characterization of the residual uncorrelated error variance in ground-based post-processing of SWOT data after smoothing with a realizable 2-dimensional low-pass filter is given in Section 3e in a general form that is applicable to any user-specified half-power filter cutoffs. This provides the information needed by oceanographic users to quantify the residual noise in SWOT estimates of SSH after applying their specific choice of smoothing of the onboard pre-processed SWOT data.

3a. The 2-Dimensional Spectrum of White Noise

Consider a 2-dimensional spatial field $h(x, y)$ of SSH sampled at discrete locations (x_m, y_n) for $x_m = m\Delta x,\ m = 0, 1, \ldots M-1$ and $y_n = n\Delta y,\ n = 0, 1, \ldots, N-1$, where Δx and Δy are the sample intervals in the x and y dimensions. The sample variance σ_h^2 of this spatial field is related to its 2-dimensional wavenumber sample power spectral density $S_{2d}(k_{m'}, l_{n'})$ at wavenumbers $k_{m'} = m'/(M\Delta x)$ and $l_{n'} = n'/(N\Delta y)$ by Parseval's Theorem,

$$\sigma_h^2 \;=\; \sum_{\substack{m'=-M/2+1 \\ m'\neq 0}}^{M/2} \;\sum_{\substack{n'=-N/2+1 \\ n'\neq 0}}^{N/2} S_{2d}(k_{m'}, l_{n'})\, \Delta l\, \Delta k, \tag{3.1}$$

where

$$\Delta k = \frac{1}{M\Delta x} \tag{3.2a}$$

$$\Delta l = \frac{1}{N\Delta y} \tag{3.2b}$$

are the discrete Fourier wavenumber intervals that are defined by the record lengths $M\Delta x$ and $N\Delta y$ in the x and y dimensions. Equation (3.1) is the 2-dimensional extension of Parseval's Theorem (2.22) in one dimension that was considered in Section 2e. As in the 1-dimensional case, the zero wavenumbers corresponding to indices $m' = n' = 0$ are excluded from the sums in (3.1) because the spectral values at these wavenumbers correspond to the squared value of the sample mean, which does not contribute to the sample variance σ_h^2. The indices $m' = M/2$ and $n' = N/2$ that define the upper limits of the summations in (3.1) correspond to the Nyquist wavenumbers $k_{\mathcal{N}}$ and $l_{\mathcal{N}}$ in the x and y dimensions, respectively, that are defined for the discrete sample intervals Δx and Δy to be

$$k_{\mathcal{N}} = \frac{M}{2}\,\Delta k \;=\; \frac{M}{2}\,\frac{1}{M\Delta x} \;=\; \frac{1}{2\Delta x} \tag{3.3a}$$

$$l_{\mathcal{N}} \;=\; \frac{N}{2}\,\Delta l \;=\; \frac{N}{2}\,\frac{1}{N\Delta y} = \frac{1}{2\Delta y}\,. \tag{3.3b}$$

The 2-dimensional form (3.1) of Parseval's Theorem can be interpreted as a discretized statement that the sample variance is equal to the volume under the 2-dimensional wavenumber sample power spectral density. Analogous to the 1-dimensional case considered in Section 2e, this is easily seen by considering the limit as the record lengths $M\Delta x$ and $N\Delta y$ in the x and y dimensions approach infinity. The wavenumber intervals Δk and Δl then become infinitesimally small values dk and dl and the discrete wavenumbers $k_{m'}$ and $l_{n'}$ become continuous. The subscripts m' and n' can then be dropped and the discrete summations become integrals so that (3.1) can be written in simpler form as[9]

$$\sigma_h^2 = \int_{-k_{\mathcal{N}}}^{k_{\mathcal{N}}} \int_{-l_{\mathcal{N}}}^{l_{\mathcal{N}}} S_{2d}(k,l)\, dl\, dk\,. \tag{3.4}$$

The variance of the discretely sampled spatial field $h(x_m, y_n)$ with infinitely long record lengths $M\Delta x$ and $N\Delta y$ in the two dimensions is thus the integrated volume under the 2-dimensional sample power spectral density. A point that is important to the analysis in Sections 3c and 3d is that the ranges of integration are finite. The Nyquist wavenumbers (3.3) that define the lower and upper bounds of the integrals are imposed by the discrete sample intervals Δx and Δy, regardless of the record lengths $M\Delta x$ and $N\Delta y$.

The integral representation (3.4) of Parseval's Theorem greatly simplifies the notation in the analysis that follows. It should be kept in mind, however, that the resolutions (3.2a,b) of the wavenumbers k and l in the x and y dimensions are finite and are imposed by the finite record lengths $M\Delta x$ and $N\Delta y$.

For the contribution of uncorrelated errors $\epsilon(x, y)$ with variance σ_ϵ^2 to the total SSH variance σ_h^2, the 2-dimensional wavenumber sample power spectral density is constant at all wavenumbers

k and l (i.e., "white"). This constant value can be defined to be Σ_{2d}. Then Parseval's Theorem (3.4) for the spatially uncorrelated measurement error contribution to the variance of the SSH measurements reduces to

$$\sigma_\epsilon^2 = \int_{-k_\mathcal{N}}^{k_\mathcal{N}} \int_{-l_\mathcal{N}}^{l_\mathcal{N}} \Sigma_{2d}\, dl\, dk = 4k_\mathcal{N} l_\mathcal{N} \Sigma_{2d} = \frac{\Sigma_{2d}}{\Delta x \Delta y}\,.$$

The constant 2-dimensional power spectral density for a specified white noise error variance σ_ϵ^2 is therefore

$$\Sigma_{2d}(k,l) = \Delta x\, \Delta y\, \sigma_\epsilon^2, \quad \text{for } -k_\mathcal{N} \le k \le k_\mathcal{N} \text{ and } -l_\mathcal{N} \le l \le l_\mathcal{N}. \tag{3.5}$$

3b. The 1-Dimensional Spectrum of White Noise

The 1-dimensional wavenumber sample power spectral density of $h(x,t)$ in the x dimension (which could be defined to be parallel to the satellite ground track) can be obtained from the 2-dimensional wavenumber sample power spectral density by integrating over all wavenumbers in the y dimension,

$$S_{1d}(k) = \int_{-l_\mathcal{N}}^{l_\mathcal{N}} S_{2d}(k,l)\, dl\,, \quad \text{for } -k_\mathcal{N} \le k \le k_\mathcal{N}\,. \tag{3.6}$$

Since $h(x,y)$ is a real variable, the 1-dimensional wavenumber sample power spectral density is symmetric, i.e., $S_{1d}(-k) = S_{1d}(k)$. As discussed in Section 2e, it is therefore conventional to double[10] the 1-dimensional sample power spectral density values (3.6) and display them over only the positive wavenumbers k. This 1-sided representation of the 1-dimensional wavenumber sample power spectral density, which will be denoted as in Section 2e with a prime, is thus related to the 2-sided, 1-dimensional wavenumber sample power spectral density $S_{1d}(k)$ and the 2-dimensional wavenumber sample power spectral density $S_{2d}(k,l)$ by

$$S'_{1d}(k) = 2S_{1d}(k) = 2\int_{-l_\mathcal{N}}^{l_\mathcal{N}} S_{2d}(k,l)\, dl\,, \quad \text{for } 0 \le k \le k_\mathcal{N}\,. \tag{3.7}$$

The 1-sided, 1-dimensional sample power spectral density for the specified white noise variance σ_ϵ^2 can thus be related to the constant 2-dimensional white noise spectrum Σ_{2d} by

$$\Sigma'_{1d}(k) = 4l_\mathcal{N}\, \Sigma_{2d} = \frac{2\Sigma_{2d}}{\Delta y}\,, \quad \text{for } 0 < k \le k_\mathcal{N}\,. \tag{3.8a}$$

From (3.5), this constant 1-sided, 1-dimensional white noise spectrum can be expressed in terms of the uncorrelated error variance σ_ϵ^2 as

$$\Sigma'_{1d}(k) = 2\Delta x\, \sigma_\epsilon^2\,, \quad \text{for } 0 < k \le k_\mathcal{N}\,. \tag{3.8b}$$

3c. Spectral Characteristics of Smoothed SWOT Data

As described in Section 2c, the present plan[2] for onboard pre-processed estimates of SSH can be understood by considering smoothing of the raw SWOT data in the form of a 2-dimensional

Parzen smoother with a span of $L_4(x) \times L_4(y) = 2$ km $\times$ 2 km. The present plan further calls for posting of the onboard pre-processed estimates of SSH on a 2-dimensional grid with dimensions $\Delta x \times \Delta y = 1$ km $\times$ 1 km. The variance σ_ϵ^2 of the uncorrelated errors in these onboard pre-processed SWOT estimates of SSH could be reduced in ground-based post-processing by applying 2-dimensional block averaging over $J \times K$ pre-processed estimates of SSH. Since the 1 km $\times$ 1 km onboard estimates of SSH are statistically uncorrelated (see Section 2c), the $J \times K$ block averaging reduces the uncorrelated error variance σ_ϵ^2 by a multiplicative factor of $(JK)^{-1}$. The residual noise variance in the post-processed estimates of SSH after block averaging would thus be

$$\overline{\sigma}_\epsilon^2 = \frac{\sigma_\epsilon^2}{JK} \,. \tag{3.9}$$

As in Section 2d, the overbar distinguishes the variance $\overline{\sigma}_\epsilon^2$ of the residual uncorrelated errors in the ground-based post-processed estimates of SSH from the variance σ_ϵ^2 of the uncorrelated errors of the onboard pre-processed estimates of SSH.

In practice, block averaging of the onboard pre-processed SWOT estimates of SSH is undesirable. When such block averages are constructed as overlapping averages on the 1 km $\times$ 1 km grid, this is equivalent to 2-dimensional running averages of the SWOT data. As discussed in Section 2b, smoothing with the Parzen smoother has much better filter side-lobe suppression (see Figure 5). From (2.30a) with $L_x = 2J$ and $L_y = 2K$, smoothing the onboard pre-processed SWOT estimates of SSH using a 2-dimensional Parzen smoother with a span of $L_4(x) \times L_4(y) = 2J\Delta x \times 2K\Delta y$ reduces the variance of the uncorrelated errors by a multiplicative factor of $(JK)^{-1}$ (see Figure 5 for the case of $J = K = 7$ and $\Delta x = \Delta y = 1$ km). From (2.11), the half-power filter cutoff wavenumbers of post-processed estimates of SSH obtained by 2-dimensional Parzen smoothing with spans of $L_4(x) = 2J\Delta x$ and $L_4(y) = 2K\Delta y$ are

$$k_c \approx \frac{1}{2J\Delta x} \tag{3.10a}$$

$$l_c \approx \frac{1}{2K\Delta y} \tag{3.10b}$$

Parseval's Theorem in the integral form (3.4) that relates the variance $\overline{\sigma}_h^2$ of the post-processed estimates of SSH to the 2-dimensional power spectral density $\overline{S}_{2d}(k,l)$ of the post-processed estimates of SSH is

$$\overline{\sigma}_h^2 = \int_{-k_\mathcal{N}}^{k_\mathcal{N}} \int_{-l_\mathcal{N}}^{l_\mathcal{N}} \overline{S}_{2d}(k,l)\, dl\, dk \,. \tag{3.11a}$$

For the discretized samples and spectra, it is assumed in (3.11a) that the post-processed estimates of SSH obtained after smoothing with filter cutoff wavenumbers of k_c and l_c are constructed at the same $M \times N$ grid locations (x_m, y_n) as the onboard pre-processed estimates of SSH. The Nyquist wavenumbers $k_\mathcal{N}$ and $l_\mathcal{N}$ in (3.11a) are thus unchanged from (3.3a,b).

Neglecting the imperfections of the filter transfer function of the smoothing applied in the ground-based post-processing of the onboard pre-processed SWOT estimates of SSH (i.e., the gradual rolloff of the filter transfer function through the filter cutoff wavenumbers k_c and l_c and the side lobes at wavenumbers higher than k_c and l_c), the 2-dimensional wavenumber power spectral

density $\overline{S}_{2d}(k,l)$ of the post-processed SWOT data after smoothing is zero for all wavenumbers $|k| > k_c$ and $|l| > l_c$. Parseval's Theorem (3.11a) for the post-processed estimates of SSH can thus be written with truncated ranges of integration as

$$\overline{\sigma}_h^2 = \int_{-k_c}^{k_c} \int_{-l_c}^{l_c} \overline{S}_{2d}(k,l)\, dl\, dk\,. \tag{3.11b}$$

Analogous to (3.7), the 1-sided, 1-dimensional (e.g., along-track) wavenumber power spectral density of the post-processed SSH estimates after smoothing can be obtained from the 2-dimensional wavenumber power spectral density $\overline{S}_{2d}(k,l)$ of the post-processed SSH estimates by

$$\overline{S}'_{1d}(k) = 2 \int_{-l_{\mathcal{N}}}^{l_{\mathcal{N}}} \overline{S}_{2d}(k,l)\, dl\,.$$

Again neglecting imperfections of the filter transfer function of the smoothing applied to the pre-processed SWOT estimates of SSH, this can be written equivalently with a truncated range of integration as

$$\overline{S}'_{1d}(k) = \begin{cases} 2 \displaystyle\int_{-l_c}^{l_c} \overline{S}_{2d}(k,l)\, dl & \text{if } 0 \le k \le k_c \\ 0 & \text{otherwise.} \end{cases} \tag{3.12}$$

The 2-dimensional white-noise wavenumber sample power spectral density of the residual uncorrelated errors in the post-processed SSH estimates after smoothing can be denoted as $\overline{\Sigma}_{2d}$. From Parseval's Theorem (3.11b), the noise contribution (3.9) to the variance of the post-processed SSH estimates reduces to

$$\overline{\sigma}_\epsilon^2 = \int_{-k_c}^{k_c} \int_{-l_c}^{l_c} \overline{\Sigma}_{2d}\, dl\, dk = 4 k_c l_c\, \overline{\Sigma}_{2d}\,.$$

The 2-dimensional power spectral density of the residual white noise in the post-processed estimates of SSH for the idealized low-pass filter considered here is thus related to the residual noise variance $\overline{\sigma}_\epsilon^2$ by

$$\overline{\Sigma}_{2d}(k,l) = \begin{cases} \dfrac{\overline{\sigma}_\epsilon^2}{4 k_c l_c} & \text{if } -k_c \le k \le k_c \text{ and } -l_c \le l \le l_c \\ 0 & \text{otherwise.} \end{cases} \tag{3.13a}$$

Substituting (3.9) for $\overline{\sigma}_\epsilon^2$ expresses this equivalently in terms of the variance σ_ϵ^2 of the onboard pre-processed estimates of SSH without filtering as

$$\overline{\Sigma}_{2d}(k,l) = \begin{cases} \dfrac{\sigma_\epsilon^2}{4 k_c l_c J K} & \text{if } -k_c \le k \le k_c \text{ and } -l_c \le l \le l_c \\ 0 & \text{otherwise.} \end{cases} \tag{3.13b}$$

With the expressions (3.10a,b) for the filter cutoff wavenumbers k_c and l_c, this gives

$$\overline{\Sigma}_{2d}(k,l) = \begin{cases} \Delta x \Delta y\, \sigma_\epsilon^2 & \text{if } -k_c \le k \le k_c \text{ and } -l_c \le l \le l_c \\ 0 & \text{otherwise.} \end{cases} \tag{3.13c}$$

For wavenumbers $|k| \le k_c$ and $|l| \le l_c$, the 2-dimensional power spectral density values for the filtered white noise are thus unchanged from the constant spectral values (3.5) of the unfiltered white noise. Filtering with the idealized filter transfer function considered here thus eliminates the 2-dimensional white noise power spectral density only at wavenumbers $|k| > k_c$ and $|l| > l_c$.

From (3.12), the residual uncorrelated noise contribution (3.9) to the 1-sided, 1-dimensional power spectral density in the post-processed SSH estimates after smoothing with the idealized filter considered here is

$$\overline{\Sigma}'_{1d}(k) = \begin{cases} 4l_c\overline{\Sigma}_{2d} & \text{if } 0 < k \le k_c \\ 0 & \text{otherwise.} \end{cases} \tag{3.14a}$$

Substituting (3.13a) for $\overline{\Sigma}_{2d}$, the 1-sided, 1-dimensional power spectral density of the residual noise in terms of the residual noise variance $\overline{\sigma}_\epsilon^2$ after smoothing with an idealized filter can be written equivalently as

$$\overline{\Sigma}'_{1d}(k) = \begin{cases} \dfrac{\overline{\sigma}_\epsilon^2}{k_c} & \text{if } 0 < k \le k_c \\ 0 & \text{otherwise.} \end{cases} \tag{3.14b}$$

This 1-sided, 1-dimensional power spectral density can be expressed in terms of the variance σ_ϵ^2 of the onboard pre-processed estimates of SSH without filtering by substituting (3.9) for $\overline{\sigma}_\epsilon^2$ and (3.10a) for the filter cutoff wavenumber k_c to get

$$\overline{\Sigma}'_{1d}(k) = \begin{cases} \dfrac{2\Delta x\, \sigma_\epsilon^2}{K} & \text{if } 0 < k \le k_c \\ 0 & \text{otherwise.} \end{cases} \tag{3.14c}$$

In contrast to the 2-dimensional power spectral density (3.13c) that is unchanged in the low-wavenumber pass band $|k| \le k_c$ and $|l| \le l_c$, the smoothing in the y dimension reduces the 1-dimensional power spectral density values (3.14c) within the pass band $|k| \le k_c$ in the x dimension by a multiplicative factor of K^{-1} compared with the 1-dimensional white noise spectrum (3.8b) of the unfiltered white noise. This perhaps non-intuitive point is critically important to the derivation in Section 3d of the noise variance of the onboard pre-processed SWOT estimates of SSH from the science requirement specification in terms of the 1-sided, 1-dimensional power spectral density of SWOT data smoothed with a hypothetical idealized filter.

3d. White Noise Error Specification for SWOT Estimates of SSH

The procedure for specifying the variance σ_ϵ^2 of the uncorrelated errors in the onboard pre-processed SWOT estimates of SSH on the 1 km × 1 km grid is to "work backwards" from a specification of a baseline requirement for the residual uncorrelated errors in smoothed estimates of SSH after hypothetical ground-based post-processing with a specific idealized 2-dimensional[11] low-pass filter. As expressed by (2.30a) with $L_x = L_y = 2J$, 2-dimensional smoothing of the 1 km × 1 km onboard estimates of SSH with half-power filter cutoff wavenumbers of $k_c = l_c = (2J\Delta x)^{-1}$ reduces the uncorrelated error variance σ_ϵ^2 by a multiplicative factor of approximately J^{-2}. The science requirement for measurement accuracy could therefore be specified in terms of

the variance of the residual uncorrelated errors in post-processed SSH estimates after idealized 2-dimensional smoothing. Instead, however, the baseline requirement for instrument noise is specified in terms of the 1-sided, 1-dimensional power spectral density (3.14) of the residual white noise after smoothing with an idealized (but unrealizable) filter. In particular, the 1-sided, 1-dimensional (along-track) wavenumber power spectral density of the residual white-noise component of the total measurement errors is specified as

$$\overline{\Sigma}'_{1d}(k) = 2 \text{ cm}^2/\text{cpkm}, \quad \text{if } 1/1000 \le k \le 1/15 \text{ cpkm}. \tag{3.15a}$$

The instrumental noise error is thus specified only for wavelengths $\lambda = k^{-1}$ longer than a filter cutoff wavelength of $\lambda_c = k_c^{-1} = 15$ km.

While the science requirement (3.15a) for uncorrelated errors is specified in the form of a 1-sided, 1-dimensional along-track power spectral density, it is important to keep in mind that one of the most significant contributions of the SWOT mission is that the KaRIn instrument will measure SSH 2-dimensionally across a pair of 70-km swaths that straddle the satellite ground track separated by a 20-km gap. Specification of the measurement errors as a 1-sided, 1-dimensional spectrum allows a comparison of the SWOT performance with the accuracy and resolution of SSH achieved from traditional nadir altimetry that measures SSH only along the satellite ground track.

Post-launch verification of the spectral specification (3.15a) of the residual uncorrelated measurement errors would require filtering of the onboard pre-processed 1 km × 1 km SWOT estimates of SSH. The error specification for idealized filtering would then have to be extended to wavenumbers higher than the filter cutoff of 1/15 cpkm, i.e.,

$$\overline{\Sigma}'_{1d}(k) = 0\,, \quad \text{if } 1/15 < k \le 1/2 \text{ cpkm}. \tag{3.15b}$$

The upper wavenumber of this specification is the Nyquist wavenumber (3.3a) for the planned sample interval of $\Delta x = 1$ km. The residual white noise error spectrum after smoothing with an idealized filter is shown by the dashed red line in Figure 7.

The variance of the residual white noise in ground-based post-processed estimates of SSH after smoothing with an idealized filter to eliminate variability with scales shorter than $\lambda_c = 15$ km can be determined by inverting (3.14b) with $\Sigma'_{1d}(k)$ given by (3.15) to obtain[12]

$$\overline{\sigma}_\epsilon^2 = k_c \overline{\Sigma}'_{1d} = \frac{\overline{\Sigma}'_{1d}}{15} = 0.133 \text{ cm}^2. \tag{3.16}$$

The requirement for ground-based post-processed estimates of SSH after filtering specifically with the hypothetical idealized filter is thus that the residual uncorrelated errors after filtering have an RMS value of 0.365 cm.

From (2.31), a filter cutoff wavelength of $\lambda_c = k_c^{-1} = l_c^{-1} = 15$ km in both dimensions is equivalent to averaging over approximately $\lambda_c^2/(4\Delta x \Delta y)$ uncorrelated onboard pre-processed estimates of SSH. The variance of the uncorrelated errors in the onboard pre-processed SWOT data with $\Delta y = \Delta x = 1$ km can thus be determined by substituting (3.16) into (3.9) with $J = K = \lambda_c/2 \approx 7.5$, which gives[13]

$$\sigma_\epsilon^2 = J^2 \overline{\sigma}_\epsilon^2 = 7.48 \text{ cm}^2. \tag{3.17}$$

The requirement for the 1 km × 1 km onboard pre-processed estimates of SSH is therefore that the uncorrelated errors have an RMS value of 2.74 cm.

The uncorrelated measurement errors in the SWOT estimates of SSH are superimposed on the previously noted spatially correlated (long-wavelength) measurement errors from orbit errors and various environmental effects. The science requirement for these red-noise errors is that they have a 1-sided, 1-dimensional (along-track) power spectral density no larger than

$$S_{1d}^{red}(k) = 0.0125\, k^{-2} \text{ cm}^2/\text{cpkm}, \quad \text{for } 1/1000 < k < 1/15 \text{ cpkm}, \tag{3.18}$$

which is shown extended to higher wavenumbers by the dotted red line in Figure 7. It can be seen from Figure 7 that smoothing with a half-power filter cutoff wavelength of $\lambda_c = k_c^{-1} = 15$ km would have very little attenuation effect on the red noise. The science requirement specification of the along-track power spectral density of the total measurement errors in the SWOT data is thus given approximately by the sum of the red-noise spectrum (3.18) and the spectrum (3.15) of the residual white noise after idealized low-pass filtering with a half-power filter cutoff wavelength of 15 km. This total measurement error spectrum is shown by the solid red line in Figure 7.

The preceding analysis neglects the effects of imperfections of the filter transfer function of any realizable low-pass filter applied in ground-based post-processing of the 1 km × 1 km onboard pre-processed SWOT estimates of SSH. Post-launch verification of the SWOT instrument performance as specified by (3.15) for 2-dimensionally filtered SWOT data would therefore be subject to imperfections of the filter transfer function. For example, the dashed green line in Figure 7 shows the 1-sided, 1-dimensional power spectral density of residual white noise errors after smoothing 2-dimensionally using a Parzen smoother with a half-power filter cutoff wavelength of $\lambda_c = k_c^{-1} = l_c^{-1} = 15$ km. The solid green line is the sum of this realizable low-pass filtered white noise spectrum and the red-noise spectrum (3.18). It is evident that the flattening of the residual noise spectrum to the value specified by (3.15a) would not be detectable when the onboard pre-processed SWOT data are smoothed with a realizable filter. The uncorrelated measurement errors in the SWOT estimates of SSH are thus untestable in the form (3.15) specified in the SWOT documentation based on a hypothetical idealized low-pass filter.

In practice, the complications of the imperfections of the filter transfer function need not be addressed in post-launch verification. From (3.8b), the along-track white noise spectrum for the uncorrelated errors in the unfiltered onboard pre-processed SWOT estimates of SSH with variance (3.17) derived above from the untestable specification (3.15) is

$$\Sigma_{1d}'(k) = 2\Delta x\, \sigma_\epsilon^2 = 14.96 \text{ cm}^2/\text{cpkm}, \quad \text{for } 0 < k \le k_{\mathcal{N}}, \tag{3.19}$$

where $k_{\mathcal{N}} = 1/2$ cpkm is the Nyquist wavenumber (3.3a) for the planned sample interval $\Delta x = 1$ km of the onboard SWOT data. The uncorrelated error spectrum (3.19) is shown by the dashed blue line in Figure 7.

While the uncorrelated errors are attenuated at all wavelengths by 2-dimensional smoothing in accord with (3.14c), this is not the case for the long-wavelength measurement errors. The latter are highly correlated in the cross-track direction over the 70-km widths of the two measurement swaths. Except at the short wavelengths over which the total measurement errors are dominated

by the contribution from uncorrelated errors, the long-wavelength errors thus cannot be be reduced appreciably by cross-track smoothing. The spectrum of long-wavelength measurement errors shown by the dotted red line in Figure 7 is therefore approximately the same with and without smoothing. The spectrum of total errors in the onboard pre-processed SWOT data without smoothing can thus be approximated by the sum of the dashed blue line and the dotted red line, which is shown as the solid blue line in Figure 7.

In contrast to the solid red line that is based on a hypothetical idealized low-pass filter, the along-track spectral specification of total errors in the onboard pre-processed SWOT estimates of SSH without ground-based filtering (the solid blue line) can be tested in post-launch verification.

3e. Noise Reduction in User-Defined Smoothing of SWOT Data

The analysis in Section 3d establishes the requirement (3.17) for the variance $\sigma_\epsilon^2 \approx 7.5$ cm^2 of uncorrelated errors in the onboard pre-processed estimates of SSH. Most users interested in oceanographic applications of SWOT data are likely to want to smooth the onboard pre-processed estimates of SSH to reduce the RMS noise of 2.74 cm. But they may prefer to smooth with a filter cutoff wavelength that differs from the value of $\lambda_c = k_c^{-1} = l_c^{-1} = 15$ km that forms the basis for the measurement noise requirement as summarized in Section 3d. Moreover, users may choose to use some filter other than the Parzen smoother that has been used in this note to illustrate the effects of filtering. Furthermore, they may wish to filter the SWOT data differently in each dimension, e.g., with less smoothing in the cross-shore direction than in the alongshore direction in coastal regions. As long as the half-power filter cutoff wavenumbers k_c and l_c of a user-defined smoother are understood, the formalism in Sections 2d, 2e and 3d quantifies the feature diameter resolution and residual error variance after smoothing of the onboard pre-processed estimates of SSH.

For example, the analysis in Section 2e shows that 2-dimensional smoothing with a half-power filter cutoff wavelength of $\lambda_c = k_c^{-1} = l_c^{-1}$ in both dimensions reduces the noise by the multiplicative factor (2.31). For the planned SWOT sample spacing of $\Delta x = \Delta y = 1$ km, smoothing with $\lambda_c = 10$ km (which could be achieved with a 2-dimensional Parzen window with spans of $L_4 = 9$ km in each dimension) is thus equivalent to averaging over approximately $5^2 = 25$ of the uncorrelated 1 km $\times$ 1 km onboard pre-processed estimates of SSH. The uncorrelated error variance (3.17) that is derived in Section 3d is then reduced by a multiplicative factor of approximately 1/25, resulting in a residual uncorrelated error variance of $\overline{\sigma}_\epsilon^2 = 0.299$ cm^2 and hence a residual RMS noise of 0.547 cm in the post-processed estimates of SSH with 10 km wavelength resolution. From (2.14c), this smoothing is equivalent to a feature diameter resolution scale of approximately 5 km.

4. Summary

This note presents a detailed mathematical analysis of the resolution and noise variance in the onboard pre-processed SWOT estimates of SSH. The Parzen smoother that is defined and examined in detail in Section 2b has been used throughout the analysis as a basis for assessing the resolution and residual noise of filtered SWOT data. The motivation for using the Parzen filter is that it is simple to implement and has good filtering properties in terms of filter side lobe suppression. Moreover, it is the filter that will be applied in the cross-track direction in the onboard processor

of the SWOT instrument. The conclusions in Sections 2 and 3 would be essentially the same for any other filtering procedure with parameters chosen to give the same half-power filter cutoff wavenumbers k_c and l_c in the x and y dimensions.

The analysis in Section 2c provides a metric for characterizing the feature diameter resolution of the onboard pre-processed SWOT estimates of SSH based on the spatial autocorrelation function of white noise after onboard smoothing of the raw SWOT data. For the planned smoothing in the onboard processor with a filter cutoff wavelength of $\lambda_c \approx 2$ km in each dimension, the feature diameter resolution (2.14c) is approximately 1 km. The resulting onboard pre-processed SWOT estimates of SSH on a 1 km $\times$ 1 km grid with this smoothing are shown in Section 2c to be uncorrelated.

Most oceanographic applications of SWOT data will require additional smoothing in ground-based post-processing to reduce the noise in the onboard pre-processed SWOT estimates of SSH. The analysis in Section 2c is extended in Section 2d to characterize the feature diameter resolution for arbitrary smoothing in post-processing of the 1 km $\times$ 1 km SWOT data with arbitrary user-specified filter cutoff wavenumbers k_c and l_c in the x and y dimensions.

In order to assess the reduction of uncorrelated measurement errors that is achieved from user-defined smoothing of the 1 km $\times$ 1 km SWOT data, it is necessary to determine the error variance reduction factor that is associated with a given half-power filter cutoff wavenumber of the smoother. This is addressed in Section 2e by comparing the error reduction after filtering using the Parzen smoother with the error reduction in block averages (equivalent to a uniform-weighted running average) with the same filter cutoff wavenumber.

In practice, the uniform-weighted running average is inferior to all other smoothers because of its serious side-lobe contamination (see Section 2a). The Parzen smoother is a good choice for ground-based post-processing. Many other filters offer equally good, and some even better, filtering properties. Regardless of what filter is used, smoothing with arbitrary half-power filter cutoff wavenumbers of k_c and l_c in the two dimensions reduces the uncorrelated error variance of the onboard pre-processed SWOT estimates of SSH by approximately the multiplicative factor (2.30b).

Application of the error variance reduction factor (2.30b) to estimate the residual error variance $\overline{\sigma}_\epsilon^2$ after smoothing of the 1 km $\times$ 1 km SWOT data requires knowledge of the error variance σ_ϵ^2 of the unsmoothed onboard pre-processed SWOT estimates of SSH. This information is not explicit in the present drafts of the SWOT documentation[1,2,3]. Instead, the science requirement for uncorrelated measurement errors is specified in the form (3.15) of an along-track 1-dimensional power spectral density after hypothetical 2-dimensional smoothing of the onboard pre-processed SWOT data with an idealized filter. The lengthy mathematical analysis in Sections 3a–c derives the relation (3.14c) between this idealized spectral specification and the uncorrelated error variance σ_ϵ^2. This relation allows the "reverse engineering" in Section 3d that establishes the requirement (3.17) for the uncorrelated error variance of the onboard pre-processed SWOT estimates of SSH in order to achieve the spectral form (3.15) of the science requirement specification for SWOT. The theoretical basis for (3.17) and the RMS uncorrelated error of 2.74 cm that this equation implies are perhaps the most important contributions of the analysis presented here since this information is not provided in any of the SWOT documentation[1,2,3].

When used in concert with the multiplicative error reduction factor (2.30b) in terms of the half-power filter cutoff wavenumbers k_c and l_c in the two dimensions, the uncorrelated error variance requirement (3.17) for the onboard pre-processed SWOT estimates of SSH allows users to determine the reduction of uncorrelated errors that can be achieved from smoothing in ground-based post-processing. Regardless of the specific smoother that is used, the residual uncorrelated error variance after smoothing is given approximately by

$$\overline{\sigma}_\epsilon^2 \;\approx\; 4\,\Delta x\,\Delta y\,k_c\,l_c\,\sigma_\epsilon^2 \;\approx\; 30\,\Delta x\,\Delta y\,k_c\,l_c\ \mathrm{cm}^2\,, \tag{4.1}$$

where Δx and Δy are the sample intervals of the onboard pre-processed SWOT estimates of SSH in the two dimensions. Present plans are for a uniform grid spacing of $\Delta x = \Delta y = 1$ km.

The expression (4.1) for the residual uncorrelated error variance assumes that the accuracy requirement (3.17) for the onboard pre-processed SWOT data is achieved. If the uncorrelated error variance σ_ϵ^2 of the unfiltered onboard pre-processed SWOT data is higher or lower than the 7.5 cm^2 requirement (3.17), the numerator on the right side of (4.1) will have to be adjusted accordingly.

Acknowledgments. We thank Lee-Lueng Fu, Rosemary Morrow and Michael Schlax for helpful discussions and constructive comments that significantly improved the text. This analysis was supported by NASA grant NNX13AE32G.

Footnotes and References

1. Rodriguez, E., 2015: *Surface Water and Ocean Topography Mission (SWOT) Project Science Requirements Document.* Jet Propulsion Laboratory document JPL D-61923, Initial Release, February 12, 2015.
2. Peral, E., 2015: *KaRIn: Ka-band Radar Interferometer On-Board Processor (OBP) Algorithm Theoretical Basis Document (ATBD).* Jet Propulsion Laboratory document JPL D-79130, Draft Release, November 21, 2014.
3. Esteban Fernandez, D., B. Pollard, P. Vaze and R. Abelson, 2013: *SWOT Project Mission Performance and Error Budget.* Jet Propulsion Laboratory document JPL D-79084, October 7, 2013.
4. Harris, F. J., 1978: On the use of windows for harmonic analysis with the discrete Fourier transform. *Proc. IEEE*, **66**, 51–83.
5. As noted previously, the filtering that will actually be applied in the onboard processor consists of a Parzen smoother in the cross-track direction but a Blackman-Harris smoother in the along-track direction. Since the differences between these two smoothers are negligible, the sampling characteristics and filtering in the onboard processor can be equally well understood in terms of application of a Parzen smoother in both dimensions.
6. For reasons that are not explained in the JPL onboard processing document[2], the Parzen window to be used in the onboard processor will have a slightly shorter span of 1.96 km. The filtering properties and feature resolution capability for a Parzen smoother with a span of $L_4 = 1.96$ km are indistinguishably different from those for a span of $L_4 = 2$ km.

7. Chelton, D. B., M. G. Schlax and R. M Samelson, 2011: Global observations of nonlinear mesoscale eddies. *Progr. Oceanogr.*, **91**, 167–216.

8. In accord with the definition (2.18) of the 1-sided spectrum, the multiplicative factor of 2 in (2.21a) should be 1 for wavenumber indices $m' = 0$ and $M/2$, but this technical detail will be ignored for mathematical convenience to simplify the derivation that follows.

9. The zero wavenumbers $k = l = 0$ that are included in the integral (3.4) must be excluded from the discretized form (3.1) of this integral since they do not contribute to the sample variance. Likewise, the Nyquist wavenumbers at the lower range of each of the integrals must also be excluded from (3.4). To simplify the discussion, these distinctions between the integral and discrete forms of Parseval's Theorem will not be mentioned again in the analysis that follows.

10. As discussed in the derivation of (2.18) in Section 2e, the power spectral density values at the zero and Nyquist wavenumbers are not doubled in the discretized version of the 1-sided power spectral density, but this technical detail will be ignored for mathematical convenience to simplify the discussion in the analysis that follows.

11. Although not explicit in the science requirements document[1], the filtering required to achieve the SWOT specification for uncorrelated errors would have to be applied in both dimensions.

12. The variance on the right side of (3.16) for the uncorrelated errors in the ground-based post-processed estimates of SSH after filtering actually extends the lower bound of the wavenumber range of the spectral error specification (3.15) down to zero wavenumber. Truncating the range of integration at a lower wavenumber of 1/1000 cpkm would change the residual noise variance (3.16) to 0.131 cm^2.

13. For the specific case of the Parzen smoother considered in Section 2c, the half-power filter cutoff wavelength for a span of L_4 is $\lambda_c = k_c^{-1} = L_4/0.910$. For $\lambda_c = 15$ km, this gives $L_4 \approx 14$ km. The corresponding multiplicative error reduction factor (2.30a) for filtered SWOT data with $L_x = L_y = 14$ km is 1/49, thus resulting in an uncorrelated error variance (3.17) of $\sigma_\epsilon^2 = 49\,\overline{\sigma}_\epsilon^2 = 6.52$ cm^2. The RMS value of the uncorrelated measurement error would then be 2.55 cm. The difference from the value of 2.74 cm deduced from (3.17) is small in consideration of the various approximations used in the analysis presented here, the most significant of which is the use of the idealized filter considered in the SWOT documentation[1,2,3].

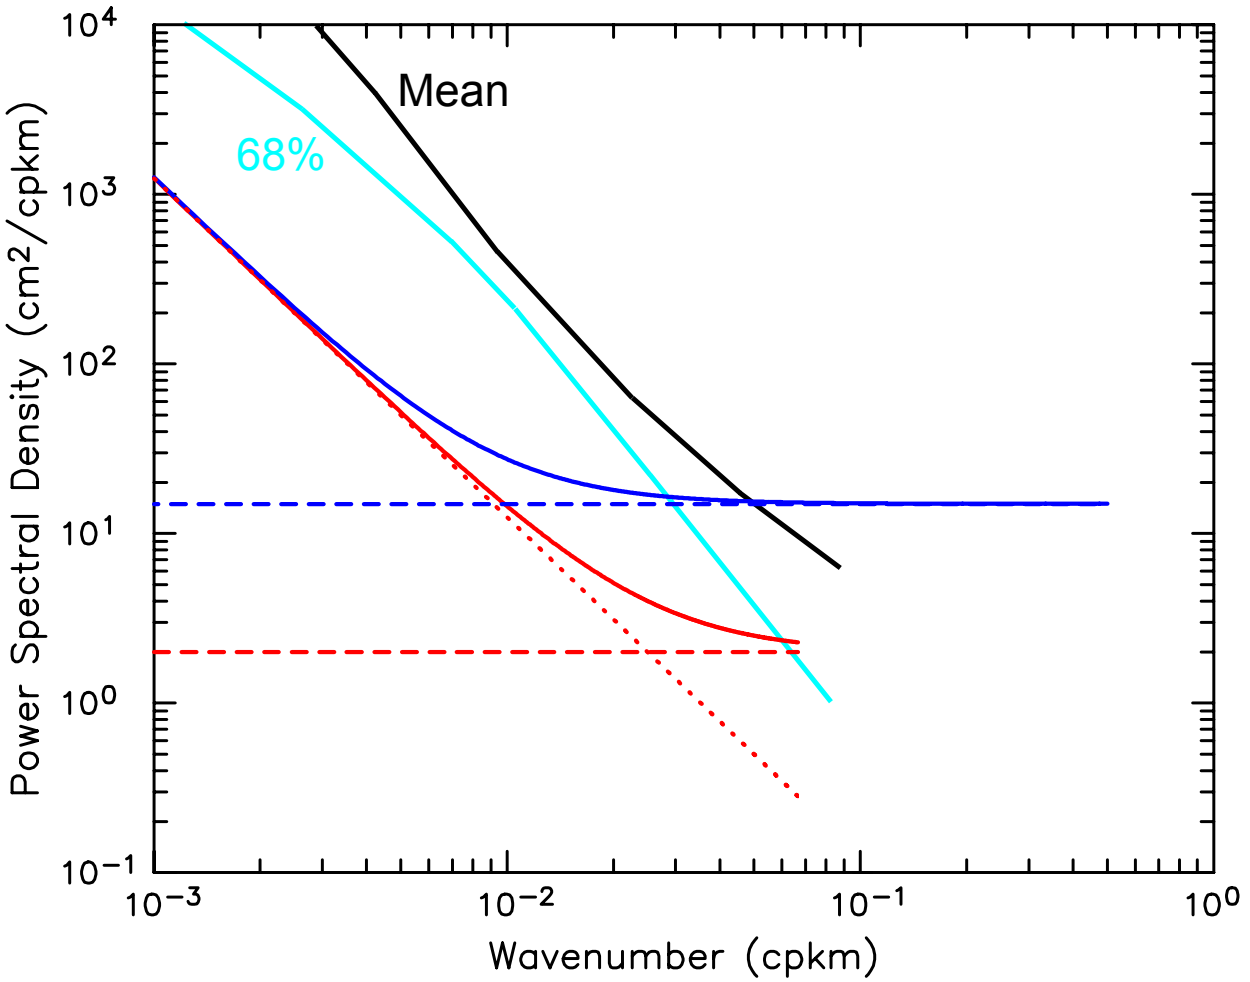

Figure 1. The 1-sided, 1-dimensional power spectral density of the science requirements for measurement errors of SWOT data before and after (blue and red lines, respectively) smoothing 2-dimensionally to eliminate variability with wavelengths shorter than a filter cutoff wavelength of λ_c=15 km. The dotted red line corresponds to the requirement for red noise from orbit errors and long-wavelength measurement errors. The dashed red line corresponds to the requirement for residual uncorrelated errors after smoothing with an idealized filter that has a magnitude of 1 for wavelengths longer than λ_c and 0 for shorter wavelengths. The solid red line is the sum of the power spectral densities of the red noise and the idealized low-pass filtered uncorrelated errors. The dashed blue line corresponds to the uncorrelated errors in onboard pre-processed SWOT data with no smoothing that is derived in Section 3d. The solid blue line is the sum of the power spectral densities of the red noise and the white noise in the onboard pre-processed SWOT data. For comparison, the mean and 68 percentile SSH signal power spectral densities from the SWOT Science Requirements Document are shown as black and cyan lines, respectively.

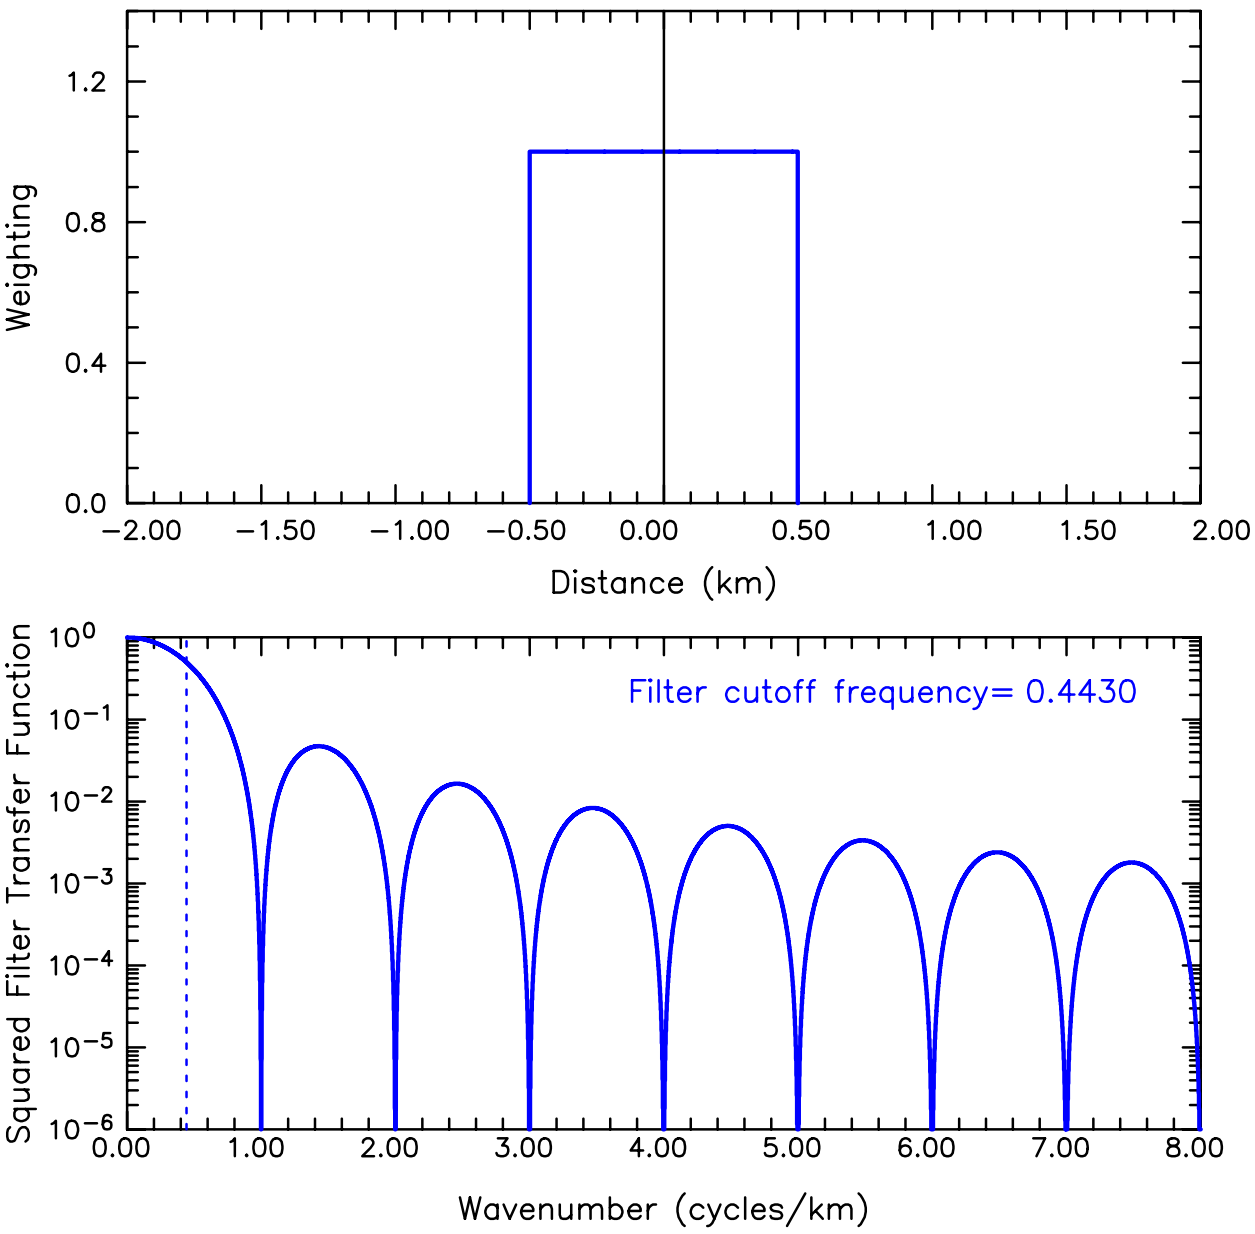

Figure 2. The weighting function (top) and squared filter transfer function (bottom) for a uniform-weighted running average with a span of L_1=1 km. The corresponding half-power filter cutoff wavenumber is $k_c = \lambda_c^{-1} = 0.443$ cpkm, shown by the vertical dashed line in the bottom panel.

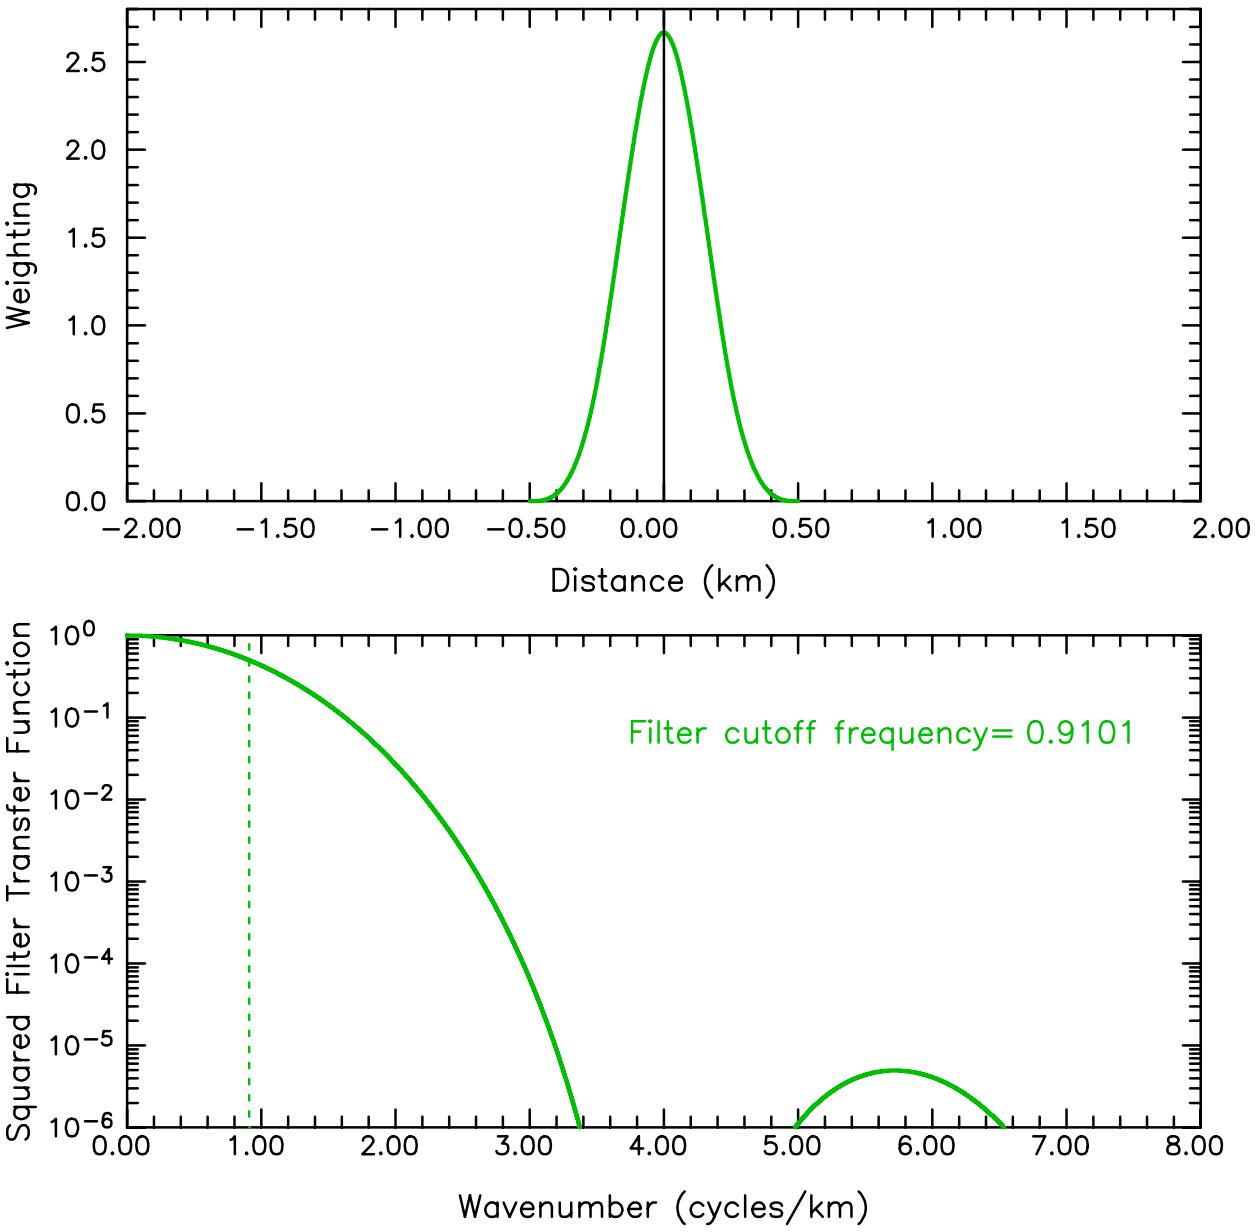

Figure 3. The weighting function (top) and squared filter transfer function (bottom) for a Parzen smoother with a span of L_4=1 km. The corresponding half-power filter cutoff wavenumber is $k_c = \lambda_c^{-1} = 0.910$ cpkm, shown by the vertical dashed line in the bottom panel.

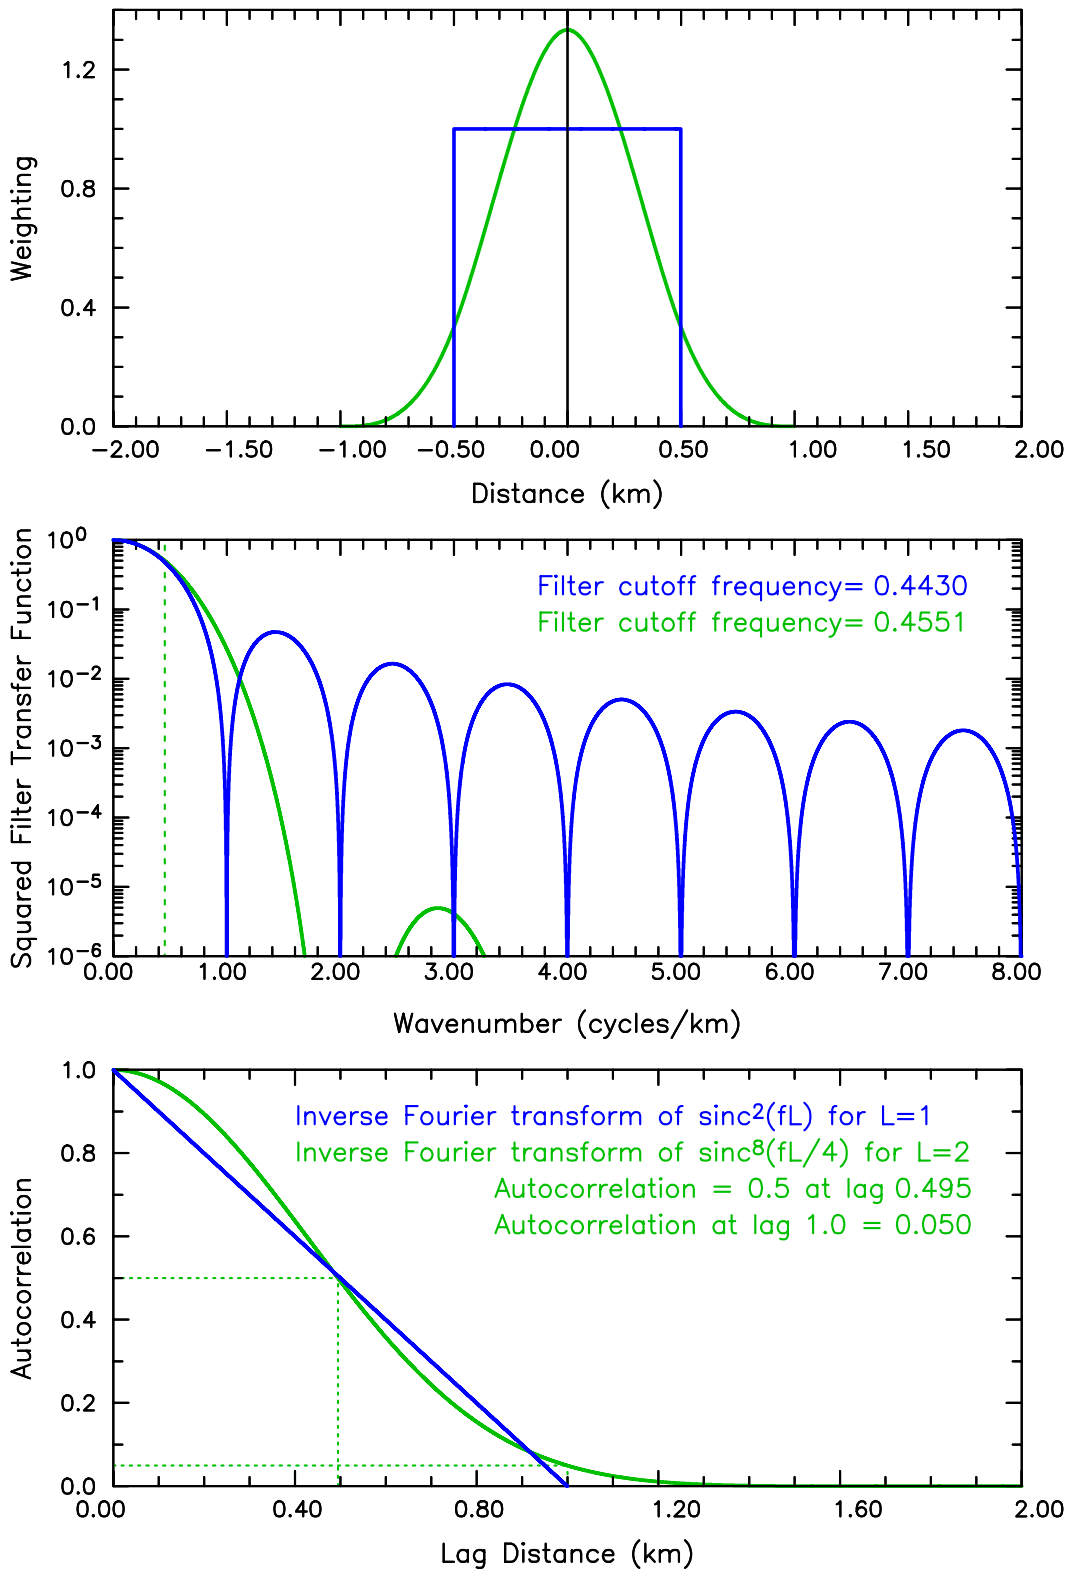

Figure 4. The weighting function (top) and squared filter transfer function (middle) for the cases of a uniform-weighted running average with a span of L_1=1 km (blue lines, same as in Figure 2) and a Parzen smoother with a span of L_4=2 km (green lines). The bottom panel shows the associated lagged autocorrelation functions (see Section 2c).

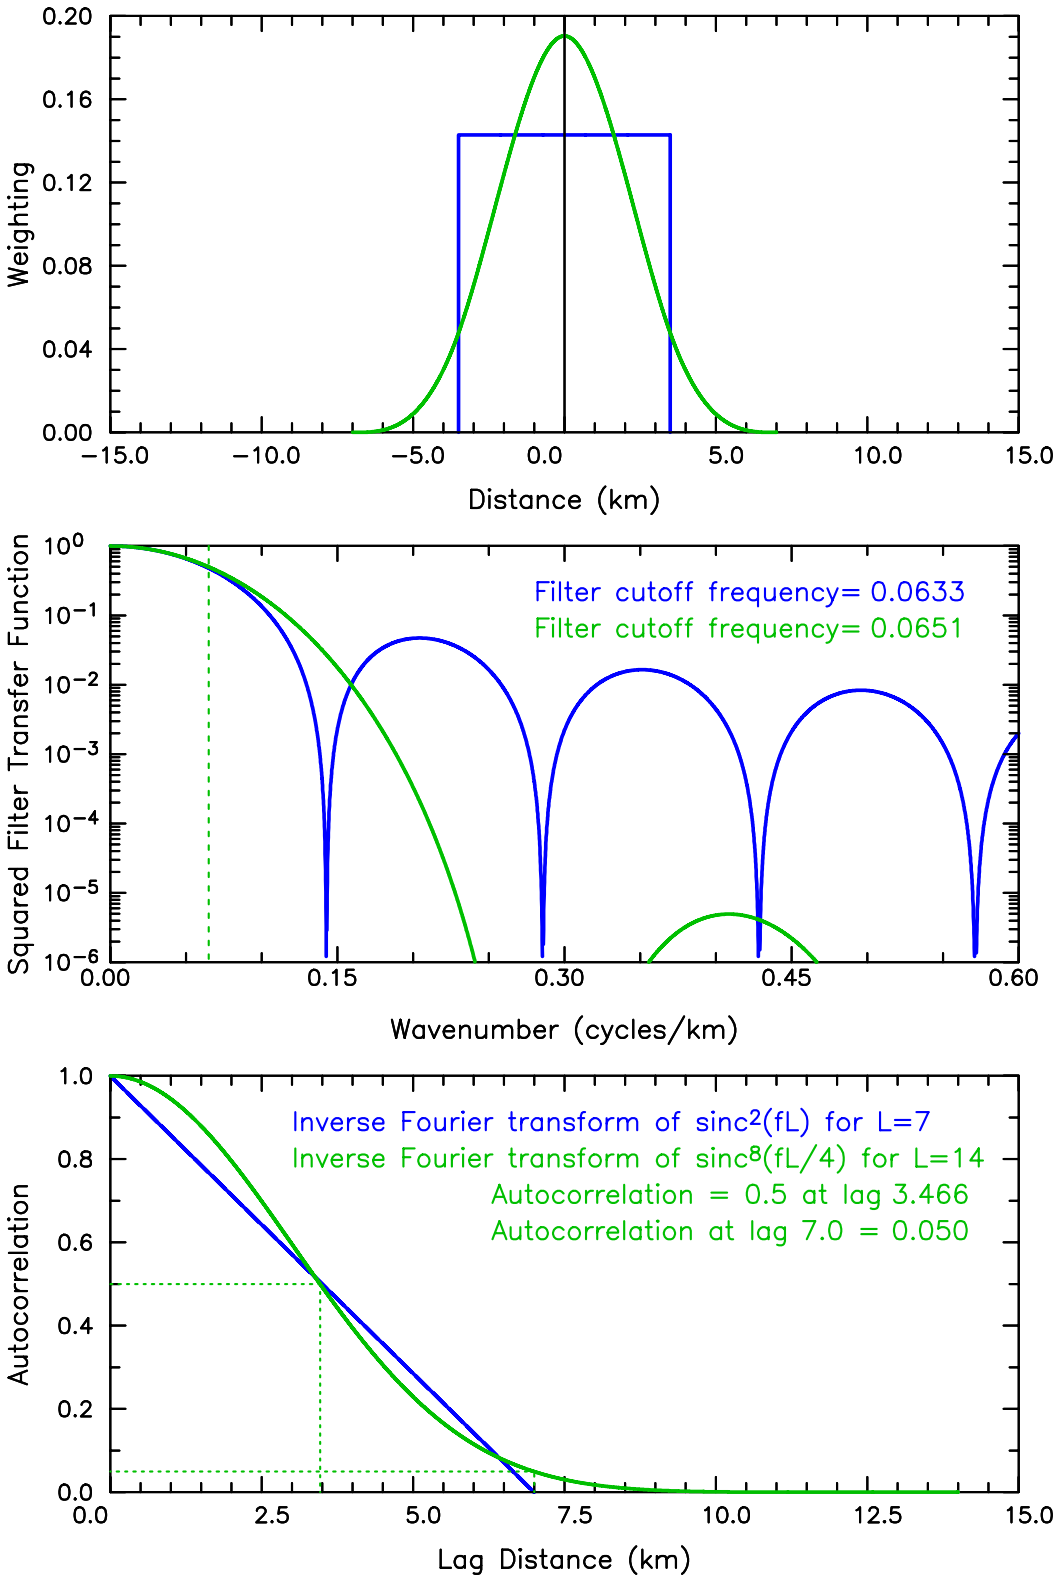

Figure 5. The same as Figure 4, except for the cases of a uniform-weighted running average with a span of L_1=7 km (blue lines) and a Parzen smoother with a span of L_4=14 km (green lines). The half-power filter cutoff wavenumbers for both of these smoothers are approximately k_c=1/15 cpkm, shown by the vertical dashed line in the middle panel.

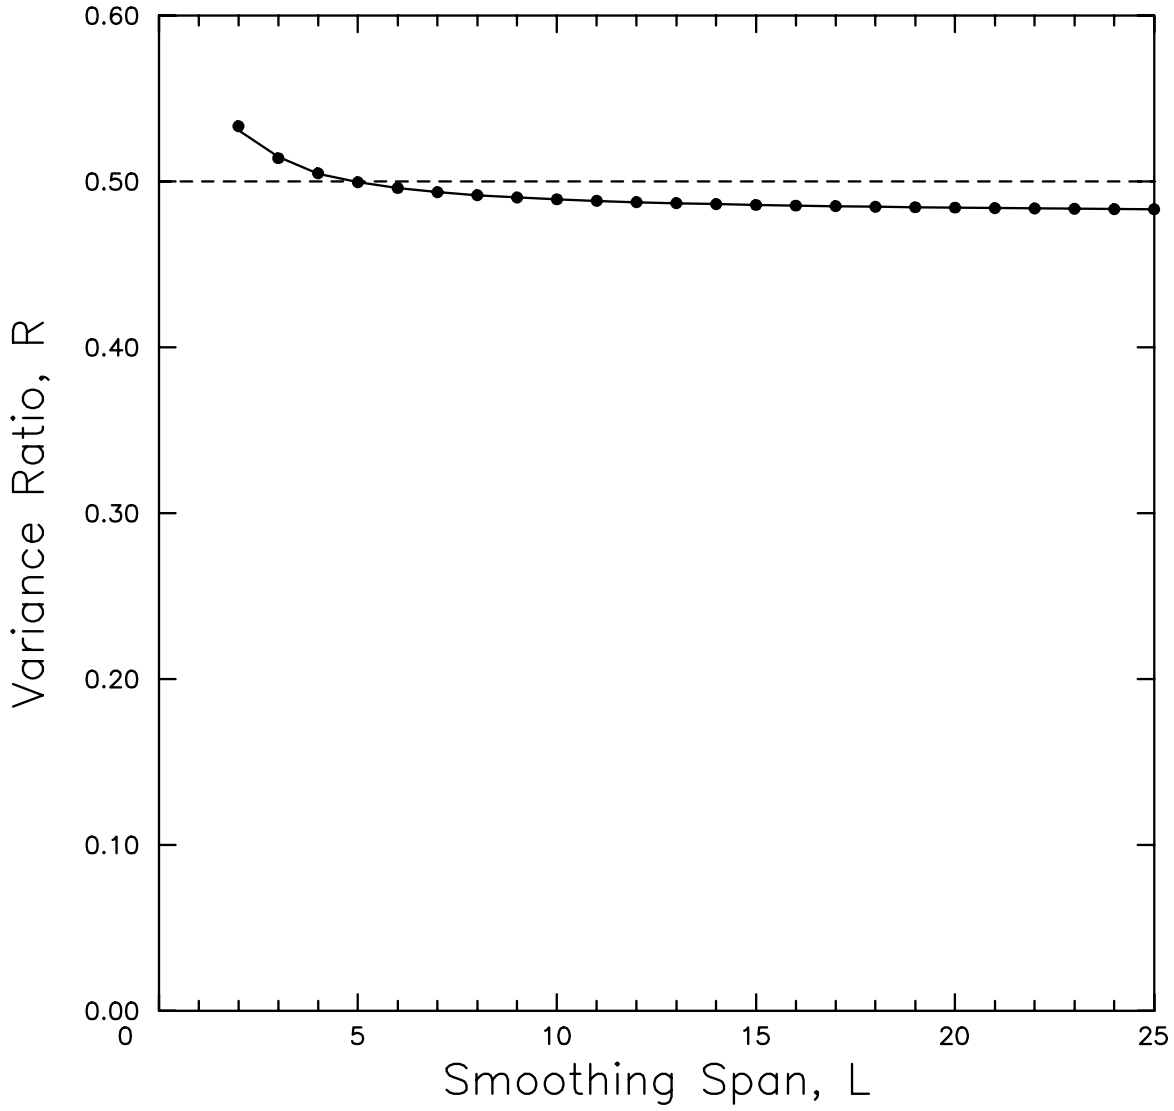

Figure 6. The ratio $R = \sigma_4^2/\sigma_1^2$ of the residual variance σ_4^2 after filtering uncorrelated white noise using a Parzen smoother with a span of $L_4 = 4L$ and the residual variance σ_1^2 after filtering with a uniform-weighted running average with the same span of $L_1 = L$. The dots correspond to the analytical approximation equation (2.28) and the solid line is the exact solution obtained by numerical integration of equation (2.25).

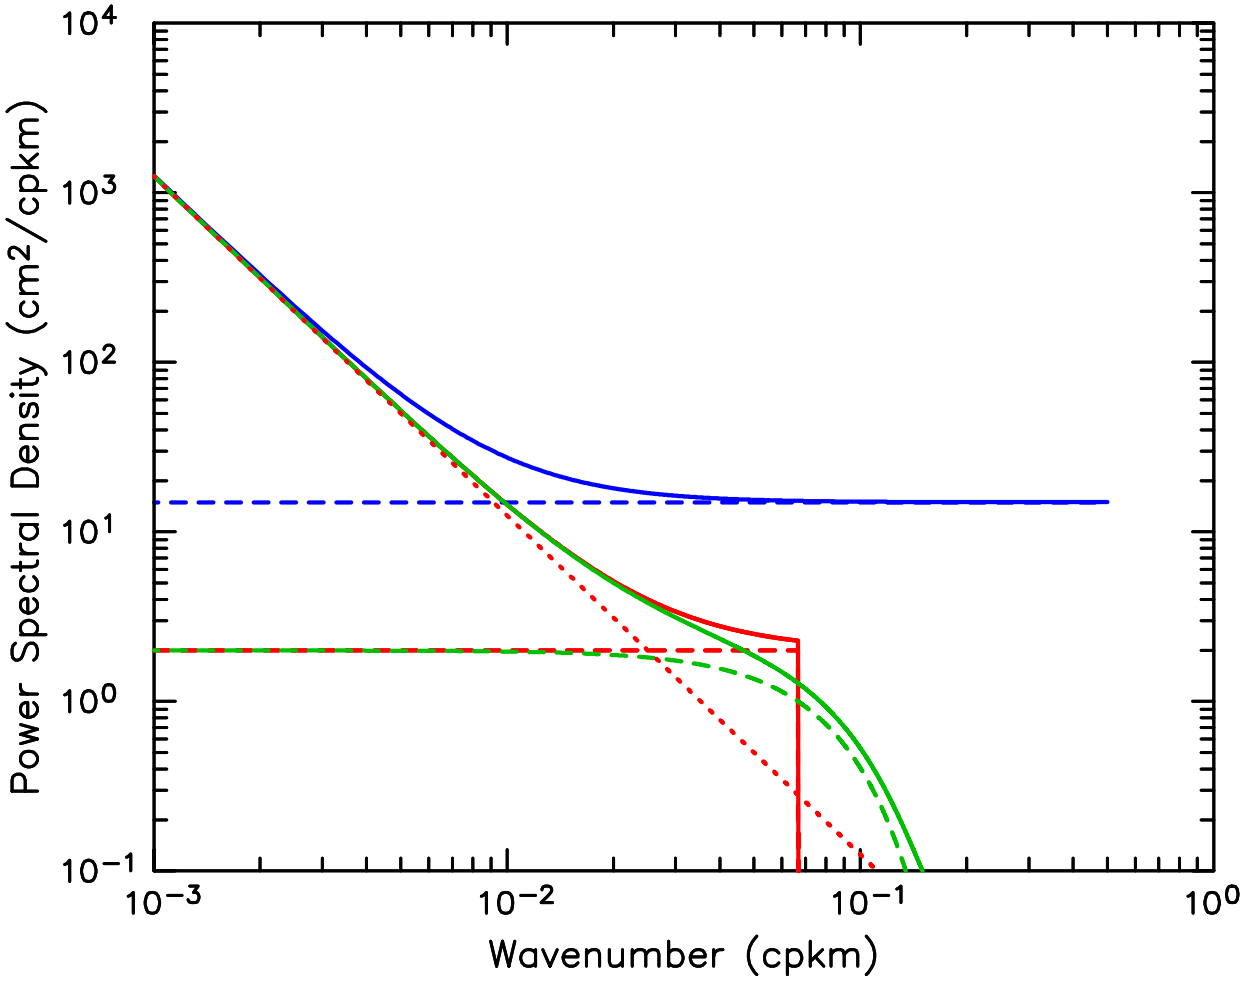

Figure 7. The 1-sided, 1-dimensional power spectral density of the science requirements for measurement noise of SWOT data before and after smoothing 2-dimensionally to attenuate variability with wavelengths shorter than a filter cutoff wavelength of λ_c=15 km. The dotted red line is the requirement for red noise from orbit errors and long-wavelength measurement errors. The dashed red line is the requirement for residual uncorrelated errors after smoothing with an idealized filter that has a magnitude of 1 for wavelengths longer than λ_c and 0 for shorter wavelengths. The solid red line is the sum of the power spectral densities of the red noise and the idealized low-pass filtered uncorrelated errors. The dashed and solid green lines are the analogous spectra for SWOT data filtered with a (realizable) Parzen smoother with a filter cutoff wavelength of λ_c. The blue dashed and solid lines are the analogous spectra for the onboard pre-processed SWOT estimates of SSH with no smoothing that is derived in Section 3d.