Signal to Noise Ratios in Fiber Spectroscopy
a component of DocDB 392

Mike Lampton

December 18, 2013

1 Introduction

The recoverable signal-to-noise ratio (SNR) for a feature depends on its strength relative to the background, and on how well the feature is concentrated in the image. Accordingly I begin this short note with a discussion of the properties of the signal's image. My model for the DESI spectrographs is presented in section 3. The SNR discussion appears in section 4 below, and my computational results are presented in section 5. In section 6, I provide a quantitative update to the Exposure Time spreadsheet of DocDB 392.

2 Point Spread Function and MTF

The usual way to obtain an estimate of a system's point spread function (PSF) is to convolve the several processes that cause that spread. When these processes are all circularly symmetric, a proven method is to express each process in 1-D frequency space as a factor in the modulation transfer function (MTF), then multiply these 1-D spectra to get the net MTF, and finally inverse Bessel transform them to get the 2-D PSF and/or the 2-D encircled energy (EE). For details see Fischer and Tadic-Galeb *Optical System Design* McGraw-Hill 2000 Chapter 10 or Schroeder *Astronomical Optics* Academic Press 2000 Chapter 11.

We initially work at the camera pupil, where the object (the fiber face) appears at infinity. Coordinates on the object are therefore angles: a fiber whose linear radius is R will be assigned an angular radius $\theta = R/F_{coll}$ where F_{coll} is the focal length of the collimator. In this scene, spatial frequency f measures the number of cycles per radian of angle at the pupil.

Because pupil diffraction transmits no power above its cutoff angular frequency $f_c = D/\lambda$, it is customary to work in normalized dimensionless units. Here, normalized frequency $v = f/f_c = f\lambda/D$ where D is the pupil diameter and λ is the wavelength. MTFs are nonzero only in the finite interval $0 < v < 1$ and Bessel integrals of the MTF span only this interval.

Following Schroeder eqns 11.1.11 through 11.1.13, suppose we have a PSF described by its radial intensity distribution $I(\theta)$ where θ is the off axis angle. PSFs in real space are usually normalized so that the central peak is one: $I(0) = 1$. In scaled space, define $i(w)$

to be this same distribution but now scaled horizontally so that $w = \theta D/\lambda$. This puts $i(w)$ into normalized frequency units. Then for a filled pupil (Schroeder 11.1.11):

$$t(v) = \frac{\pi^2}{2} \int_0^\infty i(w) J_0(2\pi v w)\, w\, dw \tag{1}$$

MTF factors are scaled vertically such that $t(0) = 1$. For each of several kinds of point spread functions (Airy, Gaussian, etc) we evaluate the $t(v)$ at a sufficiently densely packed list of points $0 < v < 1$ and then multiply these contributors to obtain the system MTF, here abbreviated as $T(v) = t_1(v) \cdot t_2(v) \cdots$. From that product we evaluate the system PSF and EE using Schroeder 11.1.12 and 11.1.13:

$$PSF(w) = 8 \int_0^1 T(v) J_0(2\pi v w)\, v\, dv \tag{2}$$

$$EE(w) = 2\pi w \int_0^1 T(v) J_1(2\pi v w)\, dv \tag{3}$$

Again, w is the normalized angular offset from the optical axis. To convert this into the linear radial coordinate r on the camera focal plane, substitute $w = rD/(F_{cam}\lambda)$. As a practical matter, the frequency normalization places $T(v)$ into a unit box $0 < T(v) < 1$ and $0 < v < 1$, where the integration can be efficiently performed by iterative Simpson or Romberg integration.

2.1 Strehl ratio

An important concept is the *Strehl ratio,* which is the ratio of an aberrated system's peak central intensity compared to the pure diffraction peak. The unobstructed peak intensity is

$$PSF(0) = 8 \int_0^1 T(v)\, v\, dv \tag{4}$$

and for pure diffraction the integral evaluates to 1/8. When other MTF factors are multiplied into the integrand, they reduce the $PSF(0)$. In showing PSF plots for aberrated systems it is customary to divide out the Strehl ratio, thereby normalizing the peak PSF to unity for easy comparison of profile shapes.

2.2 MTF Contributors

Here I list a variety of PSF contributors and their associated MTFs. In section 3 below I adopt a subset of these to model the DESI spectrograph PSF and LSF. The PSF then allows the pixel signal harvest to be calculated, and the LSF allows the sky continuum to be estimated for these same pixels. These in turn allow the evaluation of the expected signal to noise ratio on a given emission line in the presence of night sky continuum foreground.

2.2.1 Diffraction

For a uniform pupil, the Airy pattern has an MTF contribution given by Schroeder Table 11.1 page 283 which also includes variants for centrally obstructed pupils (not addressed

here). This spectrum is the autocorrelation of a uniform disk of unit diameter, that is the overlap area of two equal circles:

$$t(v) = \frac{2}{\pi}\left[arccos(v) - v\sqrt{1-v^2}\right] \tag{5}$$

It is easy to show that for this function $\int_0^1 t(v)\, v\, dv = 1/8$.

2.2.2 Uniform circular fiber at infinity

for a fiber angular radius $\theta = R/F_{coll}$ the intensity $i(w)$ is a constant out to normalized radius $w = \theta D/\lambda$. The Bessel integral identity $\int J_0(x)xdx = xJ_1(x)$ along with equation 1 provides the $t(v)$ spectrum for this case:

$$t(v) = \frac{J_1(2\pi v\theta D/\lambda)}{\pi v\theta D/\lambda} \tag{6}$$

For small arguments, $J_1(x) \to x/2$ so that for small frequencies v or for tiny fibers θ we have $t(v) \to 1$ as it must.

Nomenclature: PSF usually describes an optical system imaging a point source. Here I generalize this concept to allow including the extent of the source to be part of the calculation. The mathematical formalism is unchanged, and the PSF formulas remain, but a better nomenclature would now be ESF for *extended source function.*

2.2.3 Gaussian

At the pupil, a circularly symmetric Gaussian angular density function can be written as $exp(-\theta^2/2\sigma^2)$ where θ is the off axis angle and σ is the standard deviation angle of the distribution. Its Bessel transform is again a Gaussian. When put into cutoff-normalized form, its spectrum is:

$$t(v) = exp\left[-2\pi^2\sigma^2 v^2 D^2/\lambda^2\right] \tag{7}$$

2.2.4 Wavefront Error

Aberrations in an optical system can be quantified by their pattern of disturbance of the planar wavefront at a system pupil. Although complete WFE descriptions can be prepared using Zernike polynomials, these polynomial coefficients usually follow a simpler common distribution known as the *statistical exponential Hopkins ratio* defined by several authors. Here we adopt the formulation of Fischer and Tadic-Galeb (2000) p.348. In addition to the RMS WFE ω, this distribution is characterized by a parameter n (typically n=3) that is the ratio of pupil size to dominant disturbance wavelength lateral size:

$$t(v) = exp\left[-\frac{4\pi^2\omega^2}{\lambda^2}\cdot\left[1 - exp(-2n^2v^2)\right]\right] \tag{8}$$

2.2.5 Pixellization

Although pixel blur is not quite circularly symmetric, there is nonetheless a popular MTF contribution formula based on the average blur due to square pixel averaging; see for example Schroeder eqn 17.1.5. For a linear square pixel size p :

$$t(v) = exp\left[-0.282\frac{\pi^2 v^2 p^2 D^2}{\lambda^2 F_{cam}^2}\right] \tag{9}$$

2.2.6 Negative exponential profiles

Galaxy light profiles are usually classified using a Sersic function (for example Baes and Gentile arxiv 1009.4713): $I(r) \sim exp\left[-(r/r_{core})^{1/n}\right]$. For emission line galaxies, $n = 1$. The Bessel transform of functions of this kind have been well studied; see for example http://MathWorld.wolfram.com/HankelTransform.html. Specifically we use the Hankel proportional transform pair

$$exp(-r/r_{core}) \Leftrightarrow (1 + 4\pi^2 f^2 r_{core}^2)^{-3/2} \tag{10}$$

and move into cutoff-frequency normalized coordinates to obtain

$$t(v) = \left(1 + 4\pi^2 r_{core}^2 D^2 v^2/\lambda^2\right)^{-3/2} \tag{11}$$

2.2.7 Moffat distributions

These are often used to characterize atmospheric seeing. Distributions of this kind have a central bulge and then fall off as a power law rather than as a faster exponential or Gaussian profile. Each is described by a parameter β that indicates the steepness of the wings beyond the central bulge: $I(\theta) \sim [1 + (\theta/\theta_{core})^2]^{-\beta}$. For integral convergence, $\beta \geq 2$. Again, there is a standard analytic Hankel (generalized Bessel) proportional transform pair for functions of this kind with integer values of β as, for example, at MathWorld.com:

$$\frac{1}{(1+x^2)^\beta} \Leftrightarrow \frac{f^{\beta-1}}{2^{\beta-2}\cdot(\beta-2)!}\cdot K_{\beta-1}(2\pi f) \tag{12}$$

where K_n is the modified Bessel function of integer order n. These diverge at the origin, but the prefactor cancels the divergence and normalizes the right hand expression to unity at $f = 0$, as required for MTF service. To scale this relationship for use with a core radius and with frequencies normalized to the cutoff frequency, we need only substitute $v = f\lambda/D\theta_{core}$ and so we have for integer values of β:

$$t(\beta, v) = \frac{(\theta_{core} v D/\lambda)^{\beta-1}}{2^{\beta-2}\cdot(\beta-2)!}\cdot K_{\beta-1}(2\pi\theta_{core} v D/\lambda) \tag{13}$$

For values of β that are not integers, we use the geometric mean of the surrounding pair of $t(\beta, v)$ curves. Let $m = floor(\beta)$ and $n = ceil(\beta)$:

$$t(\beta, v) = t(m, v)^{n-\beta}\cdot t(n, v)^{\beta-m} \tag{14}$$

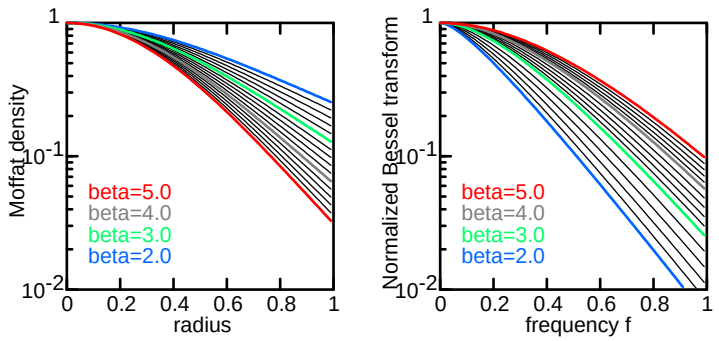

Figure 1: Left and right hand sides of Hankel-transform relationship 12 for various values of β. Source: /Java/Bessel/BesselPlot.java

3 DESI Spectrograph Model:

Here I combine a circular fiber, pupil diffraction, and an overall Gaussian to cover everything else, for a quick estimate of a NIR-band DESI spectrograph fiber image at 800nm. The spread of the Gaussian was chosen to mimic the detailed PSF calculations by Jelinsky using Zemax, reported in DocDB 334 v1 Figs 4 and 5: 2D RMS $\approx 7\mu m$ and hence 1D RMS $\approx 5\mu m$. This figure is in good agreement with Jelinsky's point-source images. One of his PSFs is shown below with linear stretch, zoomed to reveal its one-micron computational pixels. The measured FWHM here is 12 microns and therefore its RMS (1D) is 5 microns.

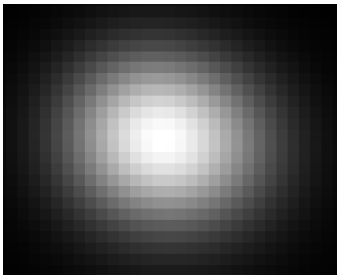

Figure 2: Image at 793nm for a point source object being imaged by the NIR spectrograph. FWHM=12 microns; RMS=5 microns. From Jelinsky DocDB 334, $NIR\_Delta\_793\_250.FITS$

The Gaussian aberration model combined with a uniform fiber disk object is probably sufficient for evaluating the signal to noise ratio of a spectroscopic singlet in the presence of continuum night sky irradiance. The singlet itself will be modelled as a narrow-wavelength emission feature broadened only by the fiber size and the spectrograph diffraction and aberrations, and will be distributed on our detector by the two dimensional $PSF(r)$. In contrast, the night sky continuum will be modelled as a stripe, uniform in the y direction and modelled by the one-dimensional $LSF(x)$. This contrast is seen in the following figure comparing the flat-topped PSF with the cylindrical bulge of the LSF.

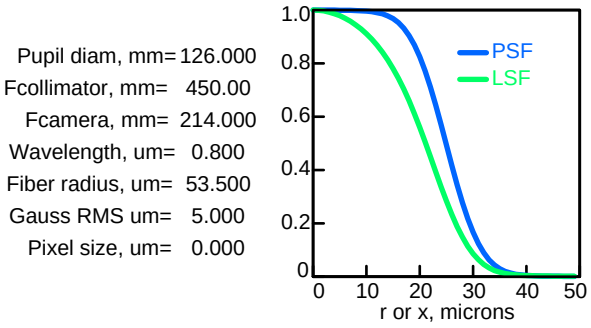

Figure 3: Plot of the PSF and the LSF for a fiber imaging system with the parameters listed. Here *fiber radius* is defined in the fiber plane, while r, x, and the Gaussian RMS are determined in the focal plane. The LSF two-sided integral is 42.24 microns; the PSF integral weighted by $2\pi r$ is 2051 square microns. Source: /Java/PSFspectro/Spectro.java

4 Signal and Noise

When a field of pixels is illuminated by signal and sky, the resulting signal to noise ratio can be recovered in the following way. Define a two-dimensional grid of pixels indexed by i. The corresponding background-subtracted noisy pixel data are denoted by d_i, and let our model be the product of the coefficient a multiplying the two-D model profile function f_i where the f are normalized to make $\sum f_i = 1$. Let the RMS error in each pixel be denoted σ_i. Then the discrepancy of the fit is

$$\chi^2 = \sum \left(\frac{af_i - d_i}{\sigma_i} \right)^2 \tag{15}$$

We suppose that the shape and location of f are correct but its coefficient a is to be determined from the data. The best fitting a is obtained by adjusting it to minimize χ^2. For the simplest case where all σ_i are equal and are not affected by changes in a, this best fit point and its variance can be determined analytically from equation 1; $\hat{a}$ is the best estimator of a in the sense that it minimizes χ^2 and $var(a)$ is the square of the change in a from $\hat{a}$ that causes χ^2 to increase by 1:

$$\hat{a} = \frac{\sum f_i d_i}{\sum f_i^2} = N_{eff} \cdot \sum f_i d_i \tag{16}$$

$$var(a) = \frac{\sigma_i^2}{\sum f_i^2} = N_{eff} \cdot \sigma_i^2 \tag{17}$$

Here I define the effective number of pixels $N_{eff} \equiv 1/\sum f_i^2$ which is the multiplier of single-pixel variance that gives the estimator variance. N_{eff} reminds us that the more widely spread a signal becomes, the more variance is gathered from its background.

For a slightly more general case that the σ_i vary from one pixel to another, but do not depend on a, we can still analytically minimize χ^2 with respect to a by expanding the

quadratic form to get:

$$\chi^2 = a^2 \sum \frac{f_i^2}{\sigma_i^2} - 2a \sum \frac{f_i d_i}{\sigma_i^2} + \sum \frac{d_i^2}{\sigma_i^2} \quad (18)$$

$$\hat{a} = \frac{\sum f_i d_i/\sigma_i^2}{\sum f_i^2/\sigma_i^2} \quad (19)$$

$$var(a) = \frac{1}{\sum f_i^2/\sigma_i^2} \quad (20)$$

An even more general case is to allow the σ_i depend on a as they would if the signal contributes to its measurement error. Again the best estimator $\hat{a}$ comes from minimizing χ^2 and obtaining its variance from the size of the region within which χ^2 increases by one. For this general case this requires a nonlinear solver.

For the present task, however, we are interested in detecting a weak emission feature, not measuring a strong feature accurately. Accordingly, our null hypothesis is that the emission feature is absent, and the distribution of estimators is the distribution in $\hat{a}$ with no signal being present. This corresponds to the faint signal case whose statistics are wholly driven by the no-signal σ_i.

Proceeding with this model, the σ_i will be evaluated over a large grid using the LSF spatial profile, and the f_i will be evaluated over that same grid using the PSF two dimensional profile.

4.1 Night Sky Continuum

From the worksheet in DocDB 392, the product of the night sky continuum intensity $0.041 photons/m^2 s\mathring{A}$ and the throughput $5.44 m^2 a s^2$ yield $223 e/ksec\mathring{A}$) along our dispersion direction in the focal plane. Since each $\mathring{A}$ spans 1.5 pixels or 22.5 μm, this is just 10 e/sec per kilosecond per μm along the ridge. In the spatial direction, the effective width of the LSF is 42.24 μm so the 1-D profile will be that of an LSF whose scaled ridge peak is 0.24 e/ksec per μm^2. We assume that owing to the vast number of samples of this ridge, its location and intensity are perfectly well known, such that its systematic subtraction error is negligible. This ridge then appears only in our field of σ_i values with the variances being given by that scaled LSF profile intensity. Other sources of noise, e.g. read noise, can be included at this point in the development.

4.2 Dark Noise

Again from the worksheet, we expect very little dark current $\lesssim$ 1 e/ksec per pixel, but anticipate seeing a read noise level of 2 e-rms per pixel per readout. The variance of this pixel dark noise is then $1+2^2$=5 e/ksec. On an area basis, this is 0.024 per square micron, which we apply to the entire grid area as a characteristic of the noise level to be seen in the valleys between night sky continuum strips.

4.3 Galaxy Singlet Signal

Again from the worksheet a nominal singlet of 4E-17 erg/cm2 sec will deliver 195 e/ksec to the focal plane. Spread over the $\sim$ 2000 square microns area of the PSF, this signal amounts

to ~ 0.1 electron per square micron. Roughly then, for a fiber image having ~ 2000 subpixels of 1 μm^2 each, we have 2000 roughly equal terms, with $f_i \sim 1/2000$, $d_i \sim 0.1$ and $\sigma_i^2 \sim 0.24$, and so we expect...

$$\hat{a} \;=\; \frac{2000 \times 0.1/(2000 \times 0.24)}{2000/(2000^2 \times 0.24)} \sim 200 \tag{21}$$

in good agreement with the assumed 195 electrons. Similarly the uncertainty in a is...

$$\sigma_a \;=\; \sqrt{var(a)} \sim \frac{1}{\sqrt{2000/(2000^2 * 0.24)}} \sim 20. \tag{22}$$

This indicates that if we had one micron size subpixels and no read noise, we might achieve a SNR $\sim 200/20 = 10$.

4.4 Applying the LSF and the PSF

More accurate work requires that real pixels be sized and their f_i, σ_i and d_i be evaluated by summing the contributions from the LSF and the PSF. This can most easily be carried out by adding the fine micron-sized grid into a coarse grid of realistic pixels. Then, a simple rectangular box of these, or a more appropriately weighted sum of these, can give estimates of the recovered signal and noise.

5 Computational Results

Here, I explore the available singlet signal-to-noise ratio in three illustrative situations. First, there is the no-pixels situation where individual one-micron grid points are evaluated for their f_i, d_i, and σ_i values, and are summed over a large (200x200 micron) postage stamp field. The noise model is an LSF stripe, appropriate for the dispersed night sky continuum. This approach would be correct if we had individual noiseless pixels on one micron centers and could weight them optimally. Second, there is the multipixel situation, where each 15x15 micron pixel adds up its own f_i, d_i, and σ_i, from which the SNR is obtained from the pixel-level $\hat{a}$ and $var(a)$ statistics. The exact result here depends on just how the pixel grid aligns with the centroid of the singlet feature. So, for these, I average the SNR over a series of one-micron pixel grid offsets from the fixed fiber centroid center. I find that very little information or SNR is given up by pixellization on this scale. Third, there is the crude *single resel* model, where a rectangular group of $m \times n$ fluxes is coadded into a single zone whose signal and noise are evaluated, again with dithering to average out the exact location of this giant resel. This single resel model is commonly employed in spreadsheets and quick estimates of SNR because it requires only three numbers: the source flux, the fraction (typically half) of flux in the resel, and the uniform background flux. The model is inaccurate when the background is spatially variable, as it is in our DESI spectroscopy application. In section 6, I develop correction factors for our exposure time spreadsheet based on both the multiple-pixel picture and the single resel picture.

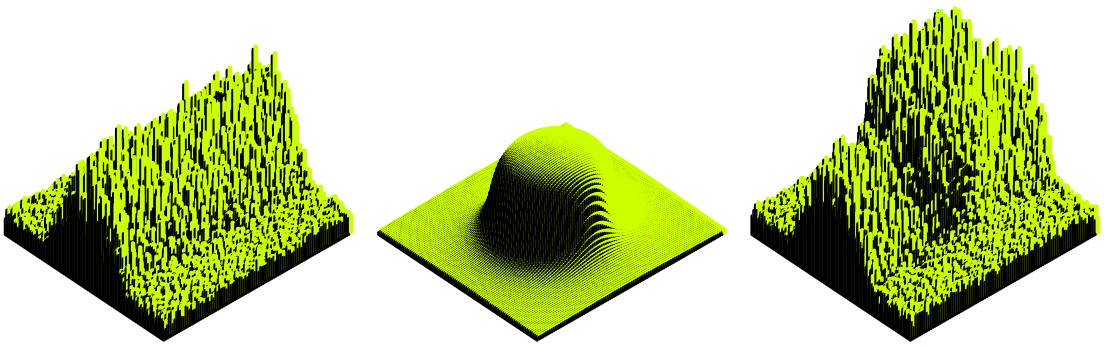

Figure 4: Pre-pixellization thumbnails, 100x100 micron square field. Left: uniform sensor noise plus night sky continuum flux and noise. Center: noiseless aberrated image of fiber. Right: sum of these two. Source: /Java/Oblique/ThreeObliques.java

5.1 SNR on a One-Micron 2D Grid

This case is of mathematical interest only. Each one square micron flux estimate has its own weighting determined by its location-dependent local noise level. Values for $\hat{a}$, $var(a)$, and SNR are obtained as per Section 4. Using Gridder.java, the results for a 200x200 grid of one micron are:

```
Given NS peak intensity, e/ksec sqmicron:  0.240000
Given dark variance, e/ksec sqmicron:  0.024000
Given signal intensity, e/ksec total:  200.000000
aHat = 200.000000
vara = 573.360441
RMS = 23.944946
SNR = 8.352493
```

5.2 SNR with Optimally-Weighted 15 micron Pixels

This is a realistic case where the individual one square micron grid of flux estimates are added up into 15x15 micron pixels. Each pixel thus contains some signal from the d_i, and some noise from the σ_i^2. The pixels are then weighted and combined into a SNR estimate for the gridding. This process is repeated, and SNRs are averaged, for a succession of pixel-grid offsets from the signal peak location:

```
Given NS peak intensity, e/ksec sqmicron:  0.240000
Given dark variance, e/ksec sqmicron:  0.024000
Given signal intensity, e/ksec total:  200.000000
SNR average = 8.095889
```

5.3 SNR with a Single Resel, Various Size Resels

This situation resembles aperture spectroscopy: an aperture is chosen that is large enough to encompass most of the signal energy. Signal and noise within this aperture are added and evaluated. Everything outside this aperture is discarded. No weighting is required, but of course the aperture size should bear some resemblance to the PSF size. Here I evaluate the SNR for 25 different rectangular apertures whose dimensions are multiples of 15 microns, i.e. they can be realized by combining square pixels of that size. Each aperture's SNR has been averaged over $\pm$ 7 μm phase with respect to the PSF peak location. Source: PSFspectro/Gridder.java.

Table 1: SNRs with NS peakVar=0.24, Dark Var=0.024, Signal=200

SNR	dy=15	30	45	60	75	90	105 μm
dx=15	2.684	3.851	4.366	4.351	4.038	3.675	3.420
30	3.963	5.623	6.277	6.167	5.703	5.188	4.828
45	4.726	6.602	7.243	7.030	6.483	5.896	5.487
60	5.127	7.063	7.655	7.375	6.790	6.174	5.745
75	5.158	7.080	7.651	7.359	6.773	6.159	5.731
90	5.087	6.979	7.540	7.252	6.674	6.069	5.647
105μm	5.019	6.885	7.438	7.153	6.583	5.986	5.570

The best such rectangle is 45 μm in the spectral direction and 60 μm in the spatial direction, and delivers a SNR=7.655.

5.4 Fraction of Monochromatic Fiber Energy within one Rectangle

This situation again resembles aperture photometry. Here, over the one-micron grid, I evaluate the monochromatic fiber energy fraction captured by various size rectangular apertures on the focal plane. Owing to the circular symmetry of the fiber image, the dx and dy dimensions are equivalent. Again, I average over a sequence of one-micron offsets between the rectangle center and the image center.

Table 2: Signal Capture Fraction for Rectangular Apertures

Frac	dy=15	30	45	60	75	90	105 μm
dx=15	0.095	0.200	0.274	0.319	0.329	0.330	0.331
30	0.200	0.415	0.561	0.644	0.661	0.664	0.664
45	0.274	0.561	0.746	0.845	0.866	0.868	0.868
60	0.319	0.644	0.845	0.951	0.972	0.975	0.975
75	0.329	0.661	0.866	0.972	0.994	0.997	0.997
90	0.330	0.664	0.868	0.975	0.997	0.999	1.000
105 μm	0.331	0.664	0.868	0.975	0.997	1.000	1.000

These findings will be useful in evaluating the correct signal capture fractions for the resel definition adopted in the Exposure Time spreadsheet. All three evaluations (spatial, LSF; spectral, PSF; and 2D) are available from this table, and are discussed in the section below.

6 Updates to the Exposure Time Spreadsheet

The DESI Exposure Time spreadsheet (in DocDB 392) has already seen a series of quantitative improvements during 2013, including the adoption of a *pivot redshift*, an improved throughput estimate from Sholl's DocDB 347, and the most recent detector read noise figure. The 12 November 2013 edition however still contains placeholders that represent early guesses as to optimum resel dimensions (4x3 pixels), the fraction of the spatial LSF that 4 pixels would capture (cell D29; 80%) and the fraction of the spectral PSF that 3 pixels would capture (cell D30, 70%).

The results of section 5 above can be used to test these three assumptions, and indeed to abandon the resel concept entirely, replacing it with a pixel-weighted SNR estimate, which should be a closer approximation to what the DESI data analysis plan offers.

6.1 Optimum Resel Size

This can be determined from section 5.3, Table 1. From among the pixel groupings explored, the optimum is 4 pixels spatially by 3 pixels spectrally. This agrees with the assumption in the 12 November edition of the spreadsheet. However the SNR shown there is 7.655, considerably higher than the spreadsheet shows, although not as high as the 8.1 found for optimally weighted pixels shown in section 5.2.

6.2 Spatial Grasp, 4 pixels

The 12 Nov spreadsheet assumed this is 80%. But from Table 3, with dx=60μm and dy=105μm, the fraction of the light captured is 97% not 80%.

6.3 Spectral Grasp, 3 pixels

The 12 Nov spreadsheet assumed this is 70%. But from Table 3, with dx=105μm and dy=45μm, the fraction of a spectral feature captured is 87% not 70%. Consequently the product of spatial and spectral grasps is 84% not the 56% from the spreadsheet assumptions. This factor can also be seen in Table 2, with dx=60μm and dy=45μm. This is a factor of 1.5 improvement in resel galaxy efficiency with no impact on resel noise, and by itself improves the spreadsheet SNR for 1ksec exposures to 7.65 compared to the earlier 5.0, or equivalently cuts the needed exposure times about in half to reach a nominal flux level at 5 sigma. The spreadsheet excludes the profile of the LSF, and so does not recognize the further slight improvement achieved along the ridge margins where the night sky contribution is reduced; this improvement is however fully included in Table 1 which predicts the higher SNR=8.1 for this case.

6.4 Optimum Pixel Weighting

The optimal pixel weighting case was presented in section 5.2 where a SNR=8.1 was predicted.

6.5 Concluding Remarks

I caution that the optimal pixel weighting case is a theoretical ideal in which each pixel's statistical weight has been chosen to minimize the variance of the flux estimator. It might not be entirely achievable in practice, where a variety of complications may interfere with achieving that optimization. In contrast, the 4x3 resel estimates are trivial to compute, since they depend only on the average LSF and PSF profiles, which are stable and deliver predictions nearly as high (here, 7.65, compared to the optimum 8.1). I include a reference to these calculations in the 1 Dec 2013 edition of the spreadsheet.

As always, it is vital to compare mathematical estimates of SNR against real world signal recovery results.