



The Stability of Fiber Spectrographs in the Faint-source Regime

Kevin Bundy^{1,2}, David Law³, Nick MacDonald¹, Kyle B. Westfall¹, T. Sivarani⁴, Devika Divakar⁴, Matthew Bershad^{5,6,7}, Meng Gu⁸, Renbin Yan⁹, Namrata Roy², Claire Poppett¹⁰, and Niv Drory¹¹

¹UC Observatories, MS: UCO Lick, UC Santa Cruz, 1156 High St, Santa Cruz, CA 95064, USA; kbundy@ucsc.edu

²Department of Astronomy & Astrophysics, UC Santa Cruz, 1156 High St, Santa Cruz, CA 95064, USA

³Space Telescope Science Institute, 3700 San Martin Dr, Baltimore, MD 21218, USA

⁴Indian Institute for Astronomy (IIA), Bangalore, India

⁵Department of Astronomy, University of Wisconsin–Madison, Madison, WI 53706, USA

⁶South African Astronomical Observatory, PO Box 9, Observatory 7935, Cape Town, South Africa

⁷Department of Astronomy, University of Cape Town, Private Bag X3, Rondebosch 7701, South Africa

⁸Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA

⁹Department of Physics and Astronomy, University of Kentucky, 505 Rose St, Lexington, KY 40506-0055, USA

¹⁰Space Sciences Laboratory, University of California, Berkeley, 7 Gauss Way, Berkeley, CA 94720, USA

¹¹McDonald Observatory, The University of Texas at Austin, 1 University Station, Austin, TX 78712, USA

Received 2021 June 30; revised 2022 May 26; accepted 2022 May 31; published 2022 August 11

Abstract

The use of optical fibers in astronomical instrumentation offers high-multiplex and light-gathering flexibility. However, with most previous fiber spectrographs optimized for large fields of view on modest-aperture telescopes, the performance of fibers in the context of faint targets on large telescopes remains largely untested. In this paper, we evaluate aspects of fiber stability, especially as they apply in the context of precision sky subtraction of faint sources at modest spectral resolution ($R \sim 3000$). After introducing a framework for describing potential systematic errors, we use publicly available data from existing instruments, including instrumentation used by the fourth-generation Sloan Digital Sky Survey’s MaNGA project (MaNGA: Mapping Nearby Galaxies at Apache Point Observatory) and the Very Large Telescope’s FLAMES: Fiber Large Array Multi Element Spectrograph. We isolate sources of fiber systematics and estimate the observed amplitude of persistent residuals as well as stochastic noise contributions resulting from changing fiber stresses. Comparing these levels against their impact on various sky subtraction schemes demonstrates that 0.1% precision sky subtraction with fiber instruments is possible. As a demonstration, we show that the MaNGA instrument can deliver 0.2% residuals on bright near-IR sky lines with nonlocal sky subtraction, if pseudo-slit limitations are addressed by allocating 50% of its fibers to sky. We further highlight recently published deep exposures that achieved a 1σ background level of 27.6 AB per square arc second, equivalent to a precision of 0.2% of the sky background continuum.

Unified Astronomy Thesaurus concepts: [Astronomical instrumentation \(799\)](#); [Spectrometers \(1554\)](#); [Sky noise \(1463\)](#)

1. Introduction

The ability of fiber-based instruments to decouple the collection of light at the telescope focal plane from the optics required to disperse that light has made them indispensable to high-multiplex multiobject spectroscopy over large ($>1^\circ$) fields of view. Such large fields are typically achieved on modest-sized telescopes (<4 m) where the physical scale of the focal plane can span 1000 mm or more. This makes focal plane mounted monolithic spectrographs extremely challenging because of cost and glass size limitations. Instruments that use pick-off mirrors require large numbers of expensive optical components. In this regime, fiber instruments provide better value, but they have historically been optimized to provide the widest fields possible in service of large-area sky surveys of relatively bright targets ($i_{AB} \lesssim 20$). Examples include the Sloan Digital Sky Surveys (York et al. 2000; Eisenstein et al. 2011; Blanton et al. 2017), the 2dF Galaxy Redshift Survey (Colless et al. 2001) and Galaxy and Mass Assembly (GAMA, Driver et al. 2011) on the Anglo-Australian Telescope, the Dark

Energy Spectroscopic Instrument (DESI) on the Mayall 4 m telescope (DESI Collaboration et al. 2016), the 4 m Multi-Object Spectroscopic Telescope (4MOST, de Jong et al. 2012), and the WHT Enhanced Area Velocity Explorer (WEAVE, Dalton et al. 2012).

Instruments on 8–10 m telescopes have aimed for smaller fields of view, typically of order $\sim 10'$. Given the long focal lengths of instrument ports near the primary mirror (e.g., Nasmyth, Cassegrain) on large telescopes, the associated focal planes still span ~ 500 mm as a result of finer plate scales (slower beam speeds). Imaging spectrographs, optimized for depth, typically sample only some fraction of this available field. In the case of Keck Observatory’s DEIMOS (DEep Imaging Multi-Object Spectrograph, Faber et al. 2003), for example, only one quarter of the telescope’s unvignetted field of view is captured by the instrument. The typical multiplex of ~ 100 at $i_{AB} \approx 24$ is also a result of the limitation that dispersed target light not overlap on the detector, a loss of efficiency that we return to below.

To achieve wider fields of view and greater multiplex on 8–10 m telescopes, fiber spectrographs are becoming more common. These instruments have fibers mounted at prime focus, such as Subaru’s Prime Focus Spectrograph (PFS, Tamura et al. 2016) and the Hobby–Eberly Telescope Dark

Energy Experiment (HETDEX, Hill et al. 2008), or at a Nasmyth port, such as the Very Large Telescope’s Multi-Object Optical and Near-infrared Spectrograph (MOONS, Cirasuolo et al. 2014) and Keck’s Fiber-Optic Broadband Optical Spectrograph (FOBOS, Bundy et al. 2020). Fiber-based instruments are further motivated by the advent of Extremely Large Telescopes (20–40 m), which have usable focal planes that can far exceed glass size limitations (see Hill 2014). The unvignetted Nasmyth field at the Thirty Meter Telescope (TMT), for instance, spans nearly 1.9 m. Fiber-based designs have been investigated for TMT’s Wide Field Optical Spectrograph (Fiber-WFOS), for the Multi-Object Spectrograph for Astrophysics, IGM, and Cosmology (MOSAIC, Hammer et al. 2014) on the European Extremely Large Telescope (E-ELT), and for the Many Instrument Fiber System (MANIFEST, Goodwin et al. 2012) on the Giant Magellan Telescope.

The fact that fiber instruments decouple the focal plane from the spectrograph optics not only allows for wider fields of view compared to monolithic “imaging spectrographs,” but also provides greater multiplex even over identical fields. The on-sky location of multiple slits or “slitlets” in an imaging spectrograph must be chosen so that they do not overlap in the dispersion direction. The multiplex is therefore determined by the linear extent of the instrument’s spatial dimension (the total “slit length”) divided by the desired per-object slit length. A fiber instrument enables a greater multiplex because targets can overlap in both on-sky dimensions. Spectral fiber traces never overlap because they are independent of on-sky location. A final benefit of this flexibility is the possibility of a more efficient sampling and modeling of the sky background (or technically the sky “foreground”), which can be collected by dedicated sky fibers distributed across the field. Depending on the required performance, adequate sky subtraction may be achieved with a smaller amount of sky-only measurements in a fiber instrument, compared to a multislit spectrograph in which the sky signal adjacent to each target is collected in every slit. This effectively frees up detector pixels for additional science targets.

Achieving this last benefit is challenging, however. The problem is that the sky foreground varies both spatially and temporally. The work by Yang et al. (2013) estimates that spatial variations on scales as small as $10''$ sampled over few-minute timescales reach 0.3%–0.7%. Although the authors suggest this may be an upper limit given potential systematics, their value is in line with a systematic floor of 0.6% reported in Sharp & Parkinson (2010). Meanwhile, the response of fiber instruments also varies spatially (i.e., across the field of view and fiber to fiber) and temporally, as a result of changing stresses and flexure as the instrument moves to track the target field. Accurate models of the sky foreground therefore require adequate sampling of both sets of variations as well as sufficient instrument calibrations.

For modest-sized telescopes where fiber instruments have been most prevalent, the precision and accuracy of subtracting the sky continuum in automated reductions of survey data typically reaches 0.5%–1% of the sky background level (e.g., Bolton et al. 2012; Law et al. 2016; Childress et al. 2017). This is on par with Poisson-limited performance in the continuum for few-hour integrations on 2–4 m telescopes, but for many-hour integrations on larger telescopes, a better performance is needed to ensure that sky subtraction systematics do not

dominate observations of faint sources. Theoretically, a 10 m telescope can reach a 1σ background level¹² of 0.3% in 16 hr of integration (in median dark conditions at Maunakea, corresponding to a V-band surface brightness of $20.7 \text{ mag arcsec}^{-2}$). For a 30 m telescope, the equivalent Poisson limit is 0.1%. The question is whether fiber instrumentation on such telescopes can deliver this level of performance.

Achieving Poisson-limited sky subtraction is even more challenging at wavelengths with strong sky-line emission. Intrinsic line variability and variations in instrument response can introduce subtle systematics that have a big impact because so much flux is concentrated in narrow sky lines. Slight errors in subtracted sky lines can leave residuals that are many factors larger than the sky continuum. This issue garnered significant attention in the MaNGA survey (MaNGA: Mapping Nearby Galaxies at Apache Point Observatory, Bundy et al. 2015), which innovated line modeling techniques that resulted in routine sky-line residuals that were only 1.2–1.3 times the theoretical Poisson error (Law et al. 2016), even among the bright lines present at wavelengths near 850 nm.

When designing future fiber instruments for faint object spectroscopy on large and extremely large telescopes, an important question therefore is how to use experience with existing instruments to ensure that requirements on sky subtraction performance can be met. This problem was addressed by Rodrigues et al. (2012) and Yang et al. (2013) using on-sky data from the VLT Fiber Large Array Multi Element Spectrograph (FLAMES, Pasquini et al. 2002). Meanwhile, sky subtraction performance results from the recently commissioned DESI (Besuner et al. 2021) are anticipated soon, with further insight to be gained from PFS and MOONS when those instruments go on-sky. In this paper, we analyze very high signal-to-noise ratio (S/N) data obtained with flat-field observations of the MaNGA instrument (Drory et al. 2015). These data allow us to explore potential fiber-based systematics at a level of precision relevant to faint object observations on larger telescopes. We also revisit the FLAMES data in the context of this MaNGA analysis.

It may seem surprising that sky subtraction performance on smaller telescopes can inform expectations and design choices for future instruments on larger telescopes. But this follows from the conservation of $A\Omega$, the product of the telescope’s collecting aperture, A , and its accessible solid angle, Ω , combined with the fact that the sky subtraction problem is fundamentally about sampling and modeling surface brightness, not discrete sources. Ignoring differences in the telescope’s designed $f/\#$, the same fiber on a 2 m telescope compared to a 10 m telescope will see a factor of 5^2 fewer photons per second (because of telescope aperture difference), but will collect those photons over a solid angle 5^2 wider (because of the difference in plate scale). The result is the same number of total sky photons collected by that fiber on either telescope. Thus, the impact of per-fiber systematics on the ability to model and subtract the sky foreground is the same, to first order.

In detail, large and small telescope designs, as well as deployed fiber diameters, are not identical. Still, the simple argument above helps explain why deep observations of low surface brightness extended sources using the MaNGA instrument on a 2.5 m telescope have achieved a continuum

¹² The background level is quoted as an error per wavelength pixel at 600 nm assuming Nyquist sampled $R \sim 5000$ spectroscopy.

sky subtraction precision of 0.2% (Gu et al. 2017), a similar level as required for the faintest discrete sources on 10 m telescopes. It also motivates our analysis of high-S/N calibration data from MaNGA.

Achieving precision at the subpercent level requires the careful mitigation and control of systematics from a variety of instrumental sources. These include contributions from the atmosphere and the telescope, the way in which light is collected at the focal plane, the fiber run itself and its output, the spectrometer, and the detector. We will refer to the combination of instrument modulations, excluding the telescope, as the “instrument response.” This response function can modulate the observed throughput as a function of wavelength and field position (and even flux) as well as the wavelength-dependent point-spread function (PSF). In the case of a spectrograph, the spectral component of the PSF determines the “line-spread function” (LSF), which can be thought of as the delivered shape of a spectral line that has zero intrinsic width (i.e., a delta function). While the purpose of instrument calibrations is to measure (and remove or model) the instrument response, the problem is that all aspects of the response function can vary with time during the course of observations. Law et al. (2021) has recently demonstrated that even with a flexing fiber instrument like MaNGA, it is possible to measure the width of the LSF to a subpercent accuracy.

Our particular aim here is to isolate the contribution of fibers to the potential systematic variations in throughput and sky-line characterization that can limit sky subtraction performance. While we attempt to focus on the contribution from fibers, we note that in practice, other terms (e.g., detector systematics and stray light) may ultimately be the limiting factors for many studies (e.g., Gu et al. 2017). Reaching 0.1% precision requires fine control of all aspects of astronomical instruments.

This paper is organized as follows. In Section 2 we introduce a formalism for defining the sky subtraction problem. To make our treatment general and tractable, we focus on local sky subtraction, though we hope our results can benefit investigations of nonlocal sky subtraction approaches as well. Section 3 highlights several major sources of systematic error introduced by the use of fibers in modest-resolution spectrographs. In this context, it also describes the specific instruments whose data will be analyzed in later sections. In Section 4, we analyze MaNGA calibration data to tease out the presence of fiber systematics in that instrument. In Section 5, we extend the analysis to on-sky data from both MaNGA and VLT-FLAMES. We present our conclusions in Section 6.

2. Systematic Errors in Sky Subtraction

At low flux or surface brightness limits, it is critically important to subtract the bright sky foreground with high fidelity in order to study the flux of much fainter sources beneath. In typical applications on large telescopes, interesting targets may be orders of magnitude fainter than the sky. Over a particular angular region of interest, i.e., the spatial extent of an astrophysical source, what is observed is a combination of the source and sky flux at that location: $F_{\text{obs}}(\lambda, t) = F_{\text{source}}(\lambda) + F_{\text{sky}}(\lambda, t)$. To be fully general, the source flux itself also depends on time, although we omit the time dependence here for convenience and on the assumption that the observer has chosen observational timescales that are appropriate given potential variations of the source. To obtain $F_{\text{source}}(\lambda)$ we must subtract a model of the sky foreground,

$M_{\text{sky}}(\lambda, t)$, but this model is almost always informed by sampling positions offset from the source.

The two basic approaches to building the sky model, $M_{\text{sky}}(\lambda, t)$, are usually denoted “local” and “nonlocal.” Local sky subtraction implies that separate source+sky and sky-only observations are made using the same physical parts of the instrument so that the light path traveled is identical in both cases. Nonlocal methods construct the sky model by sampling the sky at locations separate from the source using different parts of the instrument. The sky foreground at the source position is then determined from the model. Because the instrumental response as well as the sky foreground *both* vary in time and position (for the instrument this can mean variations along a slit, across a multislit mask, across the detector and among multiple fibers), the choice of local versus nonlocal sky subtraction involves several trade-offs.

A local approach addresses concerns about modeling the changing response function across different “parts” of the instrument or different regions of sky, but this benefit comes at a price. Over the course of observations, half of the instrument’s observing capacity must be spent on sky-only sampling. Local sky subtraction also makes assumptions for how the sky and instrument vary in time. Usually, the subtracted sky spectrum is built from averaging a sequence of high-cadence observations (as fast as 30 s for some near-IR applications; Steidel et al. 2014) on- and off-source. The assumption is made that the time variation between temporal samples is negligible. A fast cadence incurs an overhead penalty in observing efficiency. The approach is also intrinsically noisy since the local sky is collected in just one location for each source. In some cases, there may be a significant error contribution from readnoise. Mitigating readnoise motivates CCD charge shuffling (so called nod-and-shuffle observing strategies; see Glazebrook & Bland-Hawthorn 2001), which typically requires a further loss in half the instrumental multiplex to make room for detector storage zones.

The steep price of local sky subtraction motivates nonlocal techniques. Here, there is no need to model time variation in the sky because the sky and science samples are obtained contemporaneously. Additionally, the inclusion of many sky spectra taken across the field reduces the statistical noise in the final sky spectrum model that is subtracted from each source spectrum. Instead, the challenge with nonlocal sky subtraction is modeling the spatial variations across the sky as well as the response variations across different parts of the instrument. There is an additional challenge if the instrument response changes differentially with time. Various algorithms have been developed for nonlocal sky subtraction in fiber spectrographs (e.g., Sharp & Parkinson 2010; Bolton et al. 2012; Law et al. 2016; Childress et al. 2017) that achieve a typical precision of 1% the sky background by allocating $\sim 10\%$ of available fibers to sky-only positions. In Section 5.2, we describe deep MaNGA observations that predominantly used nonlocal sky subtraction combined with all-sky nods to achieve a precision equivalent to a background level of 27.6 AB arcsec⁻² (Gu et al. 2017).

2.1. Local Sky Subtraction Systematics

In what follows, we will focus on local sky subtraction techniques in order to develop intuition about the possible impact of fiber systematics, which are easier to formulate¹³ in the case of local sky subtraction.

¹³ By contrast, nonlocal performance depends more critically on the specific algorithms that are used in its implementation.

2.1.1. Sky Nodding

We begin with simple “sky nodding,” a local sky subtraction technique in which fibers (or slits) are first allocated to sources and then the telescope is offset or nodded a short distance away (e.g., $\sim 10''$), causing the same parts of the instrument that earlier collected $F_{\text{source}}(\lambda) + F_{\text{sky}}(\lambda)$ to ideally¹⁴ see flux from the sky only. We will refer to different field positions in the instrument frame (i.e., different fibers or slit positions and their mapping through the optical path to different detector locations) as A and B . We will refer to different on-sky positions as α and β . For incident flux collected by the instrument and measured at the detector (F_{det}), we wish to describe the instrument response through a transfer function, $T(\lambda, t)$, such that the object-sky nodding pair can be represented by,

$$\begin{aligned} F_{\text{det},A,t_1} &= T_A(F_{\text{source}} + F_{\text{sky},\alpha,t_1}, t_1) \\ F_{\text{det},A,t_2} &= T_A(F_{\text{sky},\beta,t_2}, t_2) \end{aligned}$$

with t_1 and t_2 representing the time of each observation (note that both the sky foreground and T depend on time). We will assume that the transfer function can be divided into separable LSF-scale and bandpass-scale components, where LSF-scale refers to wavelength separations on the order of the instrumental LSF width, σ_{inst} . The bandpass scale refers to wavelength features $\gg \sigma_{\text{inst}}$ that are apparent across wavelength differences that approach the spectral bandpass. The LSF-scale component will be described with a convolution kernel in wavelength space, $k(\lambda, t)$. The bandpass-scale component will be further separated into a multiplicative term, $g(F, \lambda, t)$, and an additive term, $h(\lambda, t)$. The transfer functions can then be approximated as

$$\begin{aligned} &T_A(F_{\text{source}} + F_{\text{sky},\alpha,t_1}, t_1) \\ &\approx g_A(\lambda, t_1)(F_{\text{source}} + F_{\text{sky},\alpha,t_1}) + h_A(\lambda, t_1) \\ &+ [k_A(\lambda, t_1) \otimes (F_{\text{source}} + F_{\text{sky},\alpha,t_1})] \\ &T_A(F_{\text{sky},\beta,t_2}, t_2) \\ &\approx g_A(\lambda, t_2)(F_{\text{sky},\beta,t_2}) + h_A(\lambda, t_2) \\ &+ [k_A(\lambda, t_2) \otimes (F_{\text{sky},\beta,t_2})]. \end{aligned} \quad (1)$$

This functional form for $T(F, \lambda, t)$ is motivated by expectations regarding how spectrographs behave. The multiplicative $g(\lambda, t)$ term is meant to capture variations in throughput that might arise from, e.g., beam vignetting, pupil illumination coupled with nonuniform dispersing media or optical aberrations, and detector characteristics. The additive $h(\lambda, t)$ term accounts for stray and scattered light in the optical system under the assumption that the dominant contribution is from sources far from the target location in the instrument plane. The $k(\lambda, t)$ term describes the wavelength calibration and LSF. In the following discussion, we will drop the explicit λ dependence when denoting these functions.

The goal of sky subtraction is to extract an accurate and precise measure of F_{source} . In the case of a single object-sky pair, our estimator of F_{source} is given by subtracting F_{det,A,t_2} from F_{det,A,t_1} , but it is clear that even in this case, several terms may contribute to a systematic residual, $R(\lambda, i)$, where i denotes

the specific sky-object pair. Let us assume that the sky nod distance is short enough (e.g., of order $10''$) such that spatial variations in the sky are negligible (i.e., $F_{\text{sky},\alpha} = F_{\text{sky},\beta}$). It should be noted here that Yang et al. (2013) report 0.3%–0.7% variations even on this scale, although they caution that these variations may result from systematics instead of intrinsic variation in the sky background. We will also assume that the exposure cadence is fast enough that the time variation in the sky foreground between exposures is negligible (in practice, $t_2 - t_1 \lesssim 60$ s). Then, the observed residuals for the i th sky-object pair about the true F_{source} can be expressed as

$$\begin{aligned} R(\lambda, i) &= F_{\text{source}}[g_A(t_1) - 1] + F_{\text{sky}}[g_A(t_1) - g_A(t_2)] \\ &+ [(h_A(t_1) - h_A(t_2))] \\ &+ [k_A(t_1) \otimes (F_{\text{source}} + F_{\text{sky}}) - k_A(t_2) \otimes F_{\text{sky}}]. \end{aligned} \quad (2)$$

In the low flux, low surface brightness regime, $F_{\text{sky}} \gg F_{\text{source}}$, and so the amplitude of residuals is often compared to the theoretical Poisson noise or to the sky foreground itself, in which case,

$$\begin{aligned} R(\lambda, i)/F_{\text{sky}} &\approx [g_A(t_1) - g_A(t_2)] \\ &+ [(h_A(t_1) - h_A(t_2))/F_{\text{sky}}] \\ &+ [k_A(t_1) \otimes F_{\text{sky}} - k_A(t_2) \otimes F_{\text{sky}}]/F_{\text{sky}} \end{aligned} \quad (3)$$

where we have assumed that $g_A \approx 1$. Note that each of the remaining terms in the sky-nodding residual involve time-dependent variations for the same part of the instrument. As we will see when exploring the behavior of these terms in the context of fiber-induced systematics, the relative amplitude of residuals from sky nodding can beat down through the combination of many sky-object pairs. First, however, it is worth exploring the corresponding expressions for a second local sky subtraction technique known as “cross-beam switching.”

2.1.2. Cross-beam Switching

Imagine the case in which some component of the response function for a specific path through the instrument drifts smoothly and monotonically with time. With an exposure cadence set by t_{exp} , this might manifest as, for example, $h_A(t) - h_A(t + t_{\text{exp}}) \approx \text{constant}$. Equation (3) warns us that in this case, no amount of coadding will improve the residual contamination of the signal, which will hit a “systematic floor.” Cross-beam switching is a variant on sky nodding that helps to address this potential problem by defining paired collecting locations (e.g., fibers) in the instrument associated with each target. In the first exposure, the source is allocated to location A while location B samples only the sky, some short distance away. In the second exposure, the source is nodded to location B , so that A samples only the sky at a third location on the opposite side of the source ($F_{\text{sky},\gamma}$). A sequence of observations alternating between these two configurations is then obtained. Assuming no spatial variation in the sky foreground over short nod distances ($F_{\text{sky},\alpha} = F_{\text{sky},\beta} = F_{\text{sky},\gamma}$, although see Yang et al. 2013) the estimator for F_{source} in the first exposure is given by

$$T_A(F_{\text{source}} + F_{\text{sky},t_1}, t_1) - T_B(F_{\text{sky},t_1}, t_1)$$

while in the second configuration, after executing the nod, we have

$$T_B(F_{\text{source}} + F_{\text{sky},t_2}, t_2) - T_A(F_{\text{sky},t_2}, t_2).$$

¹⁴ Of course, the fact that the same nod adjusts positions across the field implies that a fraction of positions will be contaminated by serendipitous sources at the “sky-only” position.

The best estimate of F_{source} comes from averaging these two. Subtracting the true F_{source} yields residual terms as follows:

$$\begin{aligned}
R(\lambda, i) = & F_{\text{source}} [g_A(t_1) + g_B(t_2) - 2] \\
& + F_{\text{sky},t_1} [g_A(t_1) - g_B(t_1)] \\
& + F_{\text{sky},t_2} [g_B(t_2) - g_A(t_2)] \\
& + [(h_A(t_1) - h_A(t_2)) + h_B(t_2) - h_B(t_1)] \\
& + [k_A(t_1) \otimes (F_{\text{source}} + F_{\text{sky},t_1}) - k_B(t_1) \otimes F_{\text{sky},t_1}] \\
& + [k_B(t_2) \otimes (F_{\text{source}} + F_{\text{sky},t_2}) - k_A(t_2) \otimes F_{\text{sky},t_2}]. \quad (4)
\end{aligned}$$

Assuming again that $F_{\text{sky}} \gg F_{\text{source}}$ and $g_A \approx g_B \approx 1$, we have

$$\begin{aligned}
R(\lambda, i) \approx & F_{\text{sky},t_1} [g_A(t_1) - g_B(t_1)] \\
& + F_{\text{sky},t_2} [g_B(t_2) - g_A(t_2)] \\
& + [(h_A(t_1) - h_A(t_2)) - (h_B(t_1) - h_B(t_2))] \\
& + [k_A(t_1) \otimes F_{\text{sky},t_1} - k_A(t_2) \otimes F_{\text{sky},t_2}] \\
& - [k_B(t_1) \otimes F_{\text{sky},t_1} - k_B(t_2) \otimes F_{\text{sky},t_2}]. \quad (5)
\end{aligned}$$

The potential of cross-beam switching to mitigate certain residual terms now becomes apparent. If a similar level of drift coherently affects either the global instrument or widespread regions of it, then $(h_A(t_1) - h_A(t_2))$ will cancel $(h_B(t_1) - h_B(t_2))$ and the additive systematic term effectively drops out. Meanwhile, a similar logic applies to the kernel term, especially because we expect the velocity structure (which determines line shape) of the sky foreground not to vary nearly as much as the amplitude of the sky continuum and line emission features. For the purposes of the convolution, $\text{LSF}_{\text{sky},t_1} \approx \text{LSF}_{\text{sky},t_2}$ to a good approximation. Coherent drift in the k term therefore largely cancels in the case of cross-beam switching.

Cross-beam switching, however, is less effective at reducing residuals from the multiplicative term if temporal sky variations are not adequately sampled. Only when the sky spectrum remains approximately constant between exposures ($F_{\text{sky},t_1} \approx F_{\text{sky},t_2}$) do the systematic multiplicative differences between the A and B instrument location paths cancel. The benefits of cross-beam switching also diminish when the changing instrument response is not coherent between A and B . If the instrument response at A and B is drifting, but in different directions, some amount of random canceling is still possible. If the response functions at A and B are independent and stochastic (as one might expect for fiber systematics alone; see below), cross-beam switching performs no better than sky nodding.

From a survey efficiency standpoint, there is little difference between sky nodding and cross-beam switching when the goal is to observe a fixed number of sources. With sky nodding, 100% of the available instrument is targeting sources 50% of the time and only observing the sky the other 50%. With cross-beam switching, 50% of the available instrument is allocated to sources (and 50% to sky) for 100% of the time. Instead, the challenge with cross-beam switching is striking a balance in an instrument design that both enables efficient target allocation and maintains flexibility for programs that do not require cross-beam switching. For example, robotic positioners that each deploy a sky-source pair of fibers (e.g., Cirasuolo et al. 2014) limits programs that desire greater multiplex without the performance that beam switching offers.

2.2. The Impact of Systematics

After combining many sky+object and sky pairs in local sky subtraction (or building a sky model based on many sky samples in nonlocal sky subtraction), what is the nature of the systematic errors that impact estimators of F_{source} ? Here, we define “systematic errors” broadly. They may manifest as flux residuals that degrade the *accuracy* of our measurement of F_{source} . They can also manifest as additional contributions to measurement uncertainty, thereby degrading the *precision* of F_{source} .

In the case of local sky subtraction techniques that execute many iterations of a recurring observing sequence, the nature of the systematics depends on whether residual terms like $g_A(t_1) - g_A(t_2)$, $h_A(t_1) - h_A(t_2)$, and $k_A(t_1) - k_A(t_2)$ are persistent—do they hover around a constant value?—or stochastic with time—are they effectively random if the sequence is long enough? Similarly, with nonlocal sky subtraction, we care about whether changes in the instrument response across the field are correlated (leading to persistent errors in accuracy) or independent (giving a stochastic term that degrades precision). Persistent terms will introduce final flux residuals and biases, while stochastic terms will cancel and beat down with time, adding only to the final uncertainty.

Among possible persistent terms, the most important may be an additive term, $h_A(t_1) - h_A(t_2)$, that accounts for scattered light. For example, with local sky subtraction, nodding from source to sky will alternate patterns of scattered light illumination in a systematic, repeating fashion. Calibration stars or other bright targets might shift back and forth in the instrument field and on the detector, leading to a persistent residual. In nonlocal sky subtraction, scattered light resulting from a few bright sources among the allocated targets could bias the sky model. In both cases, the scattered light problem will be related to the instrument’s optical design and the observational strategy. It is a generic problem common to all instruments and not unique to the use of optical fibers, although a well-designed fiber feed allows for estimates of scattered light (e.g., via choices in fiber and block spacing) that are well defined.

Our attention will therefore focus on the throughput, $g_A(t_1) - g_A(t_2)$, and LSF, $k_A(t_1) - k_A(t_2)$, systematic terms because both of these can be affected by how light rays travel through an optical fiber under changing degrees of stress. As we describe below, the complexity of these effects as well as the nature of many fiber instruments argues that their associated systematics should be mostly stochastic in nature.

3. Fiber Systematics

At low-to-moderate spectral resolution, fiber systematics ultimately derive from variations in the output fiber beam as a result of variable input illumination, stress, and fiber motion. Before we examine the potential magnitude and impact of these variations, we review basic concepts for fiber instrument designs, including the instruments used in our analysis, and we consider problematic features that affect fiber stability. We then summarize the physical basis of observed instabilities in the fiber output using broadband lab measurements as an illustration. In Sections 4 and 5, we use on-telescope observations to estimate the expected contribution to sky subtraction systematics from variations in the fiber response and demonstrate achieved performance.

3.1. Moderate-resolution Fiber Instruments

3.1.1. The MaNGA Instrument

Much of our analysis in this paper relies on data from the MaNGA instrument (Drory et al. 2015), which has completed a large, spatially resolved spectroscopic survey of nearby galaxies (Bundy et al. 2015). MaNGA utilizes the 2.5 m Sloan Telescope (Gunn et al. 2006) to deploy an array of fiber collectors at the 3° diameter focal plane. There is no atmospheric dispersion corrector. Accounting for a total of 1423 fibers, these collectors include 17 fiber bundles for science targets (ranging from 19 to 127 fibers each), 12 seven fiber bundles for instantaneous flux calibration (Yan et al. 2015), and 92 single fibers for sky sampling. All fibers have a $120\ \mu\text{m}$ core diameter and feed the two, dual-channel BOSS spectrographs (Smee et al. 2013), which provide continuous wavelength coverage over 360–1030 nm at $R \sim 2000$.

The way in which the MaNGA hardware is attached to the telescope is important for the analysis that follows. As with other SDSS programs, MaNGA makes use of a number of independent “cartridges,” each with its own, full complement of fiber bundles and associated assemblies. During the day, aluminum plug-plates with target holes drilled for a specific sky location are installed into each cartridge, and the associated fiber bundles in that cartridge are manually plugged into the plate. Random telecentricity errors are important and arise from the interface of a given bundle’s ferrule housing to the plug-plate hole and the local shape of the plug-plate, which is stressed to match the focal plane curvature (Drory et al. 2015). The downstream fiber cables are routed through the bottom of the cartridge and directed to one of two “pseudo-slits” that fix the output ends of the fibers in place. When a cartridge is mounted onto the Cassegrain focus of the telescope, the two pseudo-slits insert into the telescope-mounted spectrographs where they direct the output light into on-axis collimators.

There are two important consequences of this design. First, because the spectrographs are mounted on the telescope, there is significant spectrograph flexure during MaNGA observations. This flexure dominates MaNGA’s systematic errors. Second, the fiber run can be kept short. Not only does this improve throughput at the shortest wavelengths, it also helps reduce fiber motion. Indeed, the stiff cabling inside the cartridges prevents fiber motion beyond a few millimeters during the course of typical few-hour observations during which the cartridge rotates and the telescope slews in order to track fields across the sky. To measure the impact of more dramatic changes in fiber stress, we make use of observations using the same fiber hardware but plugged into different plates. Because crowding from the many installed assemblies makes it difficult to fully plug an installed plate, extremely tight bend radii (~ 10 cm) often result. Therefore, comparing observations from the same fiber across different pluggings captures the maximum variation in fiber stress states that MaNGA encounters.

Much of the data we analyze come from MaNGA calibration exposures. In SDSS, these calibrations are obtained by closing flat-field petals at the top of the telescope and illuminating them from underneath with flat-field and arc line lamps. The telescope pupil illumination for these calibrations adequately approximates the sky and shows little spatial and temporal variation (after a sufficient warm-up time), as we demonstrate below.

3.1.2. Instruments with Telescope-mounted Fibers and Fixed Spectrographs

Mounting the spectrographs in a separate, fixed location (e.g., on the dome floor instead of on the telescope) becomes necessary as larger spectrograph arrays exceed weight limits; it also dramatically limits the impact of spectrograph flexure. In many instruments, the fibers feeding the spectrographs are deployed at prime focus in order to maximize field of view. This obviously means that the fiber focal plane system must be mounted to the telescope and requires a long fiber run (often 20–40 m) in order to connect the fiber output to the spectrograph array.

For such instruments, motion of the telescope necessarily translates into motion in the fibers. At a minimum, there is the motion and induced stresses related to the smooth slewing and rotation involved in tracking a field. For an altitude–azimuth (Alt–Az) telescope, the magnitude of the slewing vector should not exceed the sidereal rate of 15 minutes of arc ($0^\circ 25'$) per clock minute. For a (large) telescope with prime focus 20 m above the elevation axis, fibers at prime focus would move $87\ \text{mm}\ \text{minute}^{-1}$. This motion is small compared to length of the fiber run, which experiences a maximum angular change in the gravity vector of $0^\circ 25'\ \text{minute}^{-1}$, also small. The velocity of instrument rotation depends on the field position (and formally becomes infinite at the meridian crossing). For the vast majority of field locations, the rotation rate rarely exceeds $1^\circ\ \text{minute}^{-1}$ and is more typically $0.2^\circ\ \text{minute}^{-1}$. Note that these estimates do not consider the effect of cable wrapping systems, which can impart rapid temporal variations, especially around median transit.

Over the minute-scale time frames associated with local sky subtraction (sky nodding), these rates of motion are small. Nonlocal sky subtraction with these instruments is likely to be more challenging, however, because stresses can build over time as fiber systems wrap and unwrap and thereby enter different stress states. The response function of different fibers can drift with time. What is critical for nonlocal sky subtraction is a calibration scheme that can capture the same stress states that are operative during observations.

Finally, we consider the oscillating motion induced in fibers mounted at prime focus from telescope nodding on scales of $10''$. The corresponding fiber motion at prime focus for a 20 m long telescope is 1 mm, while the angular change in the gravity vector is vanishingly small, suggesting negligible stress variation. Although this repeating motion could be a source of a persistent systematic residual term (the overall instrument response may oscillate between two states associated with the two nod positions), it is likely to be overwhelmed by other factors. Also, because the overall fiber motion and stress state is dominated by the telescope slewing and instrument rotation, the specific manifestation of systematics from nodding the fiber system is likely to change with time. It may therefore behave more like a stochastic systematic term and beat down over time.

3.1.3. Nasmyth Fiber Instruments

Fiber spectrographs deployed at the Nasmyth port become compelling on large telescopes with space or weight limitations at prime focus. Keck’s FOBOS concept (Fiber Optic Broad-band Optical Spectrograph, Bundy et al. 2019), for example, is designed to observe faint targets with high on-sky densities.¹⁵

¹⁵ Galaxies with $r_{\text{AB}} \sim 21$, for example, have an on-sky density of $1\text{--}2\ \text{arcminute}^{-2}$, providing hundreds of targets for typical $10'$ Nasmyth fields of view. At $r_{\text{AB}} \sim 21$, the on-sky density surpasses $40\ \text{arcminute}^{-2}$.

An example of a current Nasmyth fiber instrument on the Very Large Telescope is described below.

Compared to a prime focus fiber instrument, the degree of fiber motion and stress is reduced for a Nasmyth instrument because the location of the focal plane remains fixed, although the focal plane must rotate as the telescope tracks the field. The fiber run does not have to accommodate telescope slewing and the distance between the focal plane and spectrograph array can be minimized to the extent possible. As above, the corresponding rotation rate is typically below $1^\circ \text{ minute}^{-1}$ and more often $0.2^\circ \text{ minute}^{-1}$. This rate also describes the change in velocity vector, assuming the focal plane is vertical. As a result of this rotation, a fiber at the perimeter of a 1 m diameter focal plane will rarely move faster than 9 mm minute^{-1} and, more typically, the motion is 2 mm minute^{-1} .

3.1.4. VLT-GIRAFFE

The Fiber Large Array Multi Element Spectrograph (FLAMES) is mounted at the Nasmyth-A platform of VLT-UT2 (Pasquini et al. 2002). The FLAMES fiber run can feed either the GIRAFFE or UVES (D’Odorico et al. 2000) spectrographs. In Section 5 we present data from GIRAFFE at $R \sim 6500$ in which the “Medusa” mode (there is also an IFU mode) was used to collect light from up to 132 objects.

The primary components of FLAMES are the corrector, fiber positioner, fiber run, and the spectrographs. The corrector employs a lens system, 900 mm in diameter, in order to provide improved image quality over the $25'$ field of view and to locate the pupil at the center of curvature of the focal plate. The corrector is mounted with a cross-support structure to the Nasmyth rotator. This mounting system holds two metallic positioning plates that are used by the “OzPoz” pick-and-place magnetic button fiber positioning system. Fibers on both plates feed the available spectrographs, allowing the buttons on one plate to be positioned while the second plate is used to observe. These plates have a shape that follows the field curvature of a radius of 3950 mm. The OzPoz positioning accuracy is $50 \mu\text{m}$ ($0''.1$) and the minimum distance between positioners is $10''$.

FLAMES fibers feature micro-optics couplings with a $1''.2$ on-sky diameter. These convert the $f/15$ VLT Nasmyth beam to the $f/5$ GIRAFFE acceptance angle. The 13 m fiber cables are held at constant tension by a retractor. The positioning robot is composed of an $R-\theta$ arm and gripper. Back-illumination metrology provides positioning accuracy. Four of the available magnetic buttons are equipped with 19 fiber bundles that are used to target reference guide stars. There is no atmospheric dispersion corrector.

The Medusa fibers have a core/cladding/buffer diameter of $230/253/280 \mu\text{m}$. Pseudo-slits are of two types: (a) eight of the pseudo-slits are for single-fiber targets and each hold nine fibers; (b) five pseudo-slits contain 13 fibers, 12 for objects plus one additional calibration fiber. The pseudo-slits are curved to match the GIRAFFE collimator, and the center-to-center fiber density on the slit is 2.26 times the fiber core diameter (this ensures a fiber-to-fiber contamination below 0.5%). Fibers are sampled on the detector by 4.3 pixels.

3.2. Variations in Fiber Output Transmission

Optical fibers guide light via “total internal reflection,” i.e., reflections at the interface between the inner fiber “core” and an outer sheath known as the “cladding,” which is manufactured

to have a slightly lower index of refraction compared to the core. Incoming light rays with an angle of incidence shallower than the fiber’s critical angle can be guided through the fiber, while steeper incident rays dissipate through the cladding and are lost. The acceptance half-angle, θ , is often specified by the fiber’s “numerical aperture,” $\text{NA} = n \sin \theta$, where n is the index of refraction of the medium outside the fiber (e.g., air). A typical value is $\text{NA} = 0.22$, which corresponds to a full acceptance angle of $12^\circ.4$ or equivalently $f/\# = 2.27$.

For most of the applications relevant here, the fiber core diameter is $\sim 100 \mu\text{m}$ or larger, many times greater than the wavelength of the light being transmitted. This is the “multimode” regime in which hundreds or thousands of solutions (i.e., propagating modes) exist to Maxwell’s Equations given the boundary conditions of the fiber geometry and dielectric composition. As the core size is decreased relative to the transmission wavelength, a limit is reached in which only a single mode satisfies the boundary conditions and can propagate through the fiber. This is the single mode regime.

When describing the output light intensity distribution from the sum of all propagating fiber modes, two limits are considered. The distribution pattern on the output fiber face is called the “near field” and includes complex optical effects that make simplifying ray-optics assumptions invalid. In a fiber spectrograph, the near field plays an analogous role as the slit image in a traditional imaging spectrograph. The “far-field” limit is valid at a back-focal distance that is much larger the multimode fiber diameter. The far-field distribution can be thought of as the angular distribution of rays exiting the fiber. It is analogous to the spectrograph pupil. In what follows, we will be interested in how both the near- and far-field fiber output changes with time and the impact of these changes on the instrument response and resulting systematics.

3.2.1. Stress-induced and Geometric Focal Ratio Degradation

A generic feature of real-world multimode fibers is Focal Ratio Degradation (FRD), that is, the angular width of the output beam is broader compared to the input width. With the fiber output directed to the spectrograph’s collimator, if the collimator accepts light with a beam angle (or beam “speed”) similar to what is input into the fiber, FRD can send light to wide angles that never reach the collimator and are lost from the system. This loss can be mitigated by oversizing the collimator in anticipation of FRD. The relative impact of the FRD loss is also lower if the input beam angle can be made as large as possible given constraints in the optical design and fiber acceptance angle.

A number of physical effects are responsible for FRD including imperfections in the fiber material and stresses along the fiber length that slightly alter the interior geometry in complex ways (Gloge 1972). The theoretical path of internally reflected optical rays is then distorted. Equivalently, the stresses excite or favor different propagating modes. FRD is the first-order description of how these subtle physical effects combine to distort the output beam as measured at a sufficient distance from the output face of the fiber (i.e., in the far field).

We will refer to the FRD caused by impurities and stresses as “stress-induced FRD” in order to distinguish it from a second source of degradation: “geometric FRD.” Geometric FRD results from a misalignment between the axis of the input light cone and the normal to the fiber surface. Because most designs desire an input cone that is narrower than the fiber

acceptance angle, a nontelecentric beam effectively fills the fiber at a wider, but still accepted, beam angle. This results in an effectively larger output cone, which is made roughly azimuthally symmetric thanks to the tendency of circular fibers to smear (or “scramble”) input inhomogeneities azimuthally (see below).

Drory et al. (2015) highlight the importance of nontelecentricity and geometric FRD in contributing a $\sim 5\%$ throughput loss in the MaNGA instrument. Under the expectation that the magnitude of FRD can vary with time, it is an important source of potential sky subtraction systematics. We pursue this and expand on the Drory et al. (2015) analysis below to argue that geometric FRD in MaNGA strongly dominates over stress-induced FRD in affecting the instrument response. We will also show that to first order, FRD is independent of wavelength.

3.2.2. Near-field and Far-field Distortions

If FRD represents a first-order description in the far field of how the fiber output beam varies (i.e., broadens) under varying conditions, higher-order descriptions would address changes in the angular distribution pattern of light *within* the far-field beam or across the image of the near-field fiber face. We will refer to these as distortions in the fiber output light distribution, and while they may be linked to FRD, these distortions can occur and impact the instrument response even if they do not give rise to measurable FRD. One example critical to some instruments (especially planet-hunting spectrographs) occurs when the illumination structure incident on the input fiber changes with time.

In the case of the near field, the fiber image at a specific wavelength is reimaged by the spectrograph to a specific spatial and spectral location on the detector. Integrating over a continuous wavelength range, and after extracting a 1D spectrum by summing over or modeling across the detector’s spatial dimension, distortions in the shape of the fiber image will therefore manifest as distortions in the LSF. In terms of the systematic error formalism in Section 2, the high-frequency kernel term, $k(\lambda, t)$, is most relevant.

In the far field, variable distortions imply nonuniform and changing illumination patterns on dispersing elements (e.g., Volume Phase Holographic, or VPH, gratings) and other optical components in the collimation system and camera. A grating may have surface imperfections that yield changes in efficiency when different parts of the grating are illuminated unevenly. Vignetting functions and imperfections in other optical components may similarly drive wavelength-dependent variations in instrument response under changing far-field illumination. These systematics would predominantly impact the $g(\lambda, t)$ terms in Section 2. However, because spectrograph optics usually suffer aberrations that can depend on the path that light rays take through the system, the resulting PSF, or LSF as described by $k(\lambda, t)$, can also be affected by far-field distortions.

The lab measurements presented in Smith et al. (2016) and reproduced in Figure 1 illustrate the nature of near- and far-field distortions that arise from changing the input location of a focused spot (top row). This experiment represents an extreme in which the spatial structure of a target is much smaller than the fiber diameter. While some instruments, like the SDSS-V Local Volume Mapper and the Keck Planet Finder, fall in this regime, Figure 1 is useful because it illustrates generally the kinds of distortions that arise from fibers. For both the output




Figure 1. Illustration adapted from Smith et al. (2016) demonstrating how changing the location of an input illumination spot on a circular $200\ \mu\text{m}$ core fiber (top row) results in different near-field (middle row) and far-field (bottom row) output patterns.

near-field image (middle row) and far-field pattern (bottom row), the input spot is significantly scrambled in azimuth but not significantly in radius. The scrambling is not uniform, however, leading to distortions in light intensity that reach up to 40% and clearly vary as the input spot location changes. This behavior can result from both scrambling within the fiber itself and the impact of stresses from the way the fiber is mounted in its ferrule.

For the purposes of sky subtraction and the instruments we described above, the sky foreground flux fills the fiber input. Any variation in sky flux across a fiber of diameter $\sim 1''$ would be vanishingly small. In the applications that concern us here, the sky flux dominates over the source and the atmospheric PSF is typically of order the fiber diameter. With guiding requirements that ensure that source miscentering is minimal and calibration sources that also fill the fiber, most instruments do not have to worry about input flux variations as extreme as those in Figure 1. Still, it should be noted that the sky, target, and calibration flux can each illuminate the fiber in subtly different ways. For example, calibration and sky sources may both uniformly illuminate the fiber, but might inject that illumination with different angular distributions because either the calibration system does not adequately reproduce the telescope pupil or the angular pointing of fibers with respect to the chief ray differ as they are moved to different field locations. Compared to slit (imaging) spectrographs that face similar challenges, the scrambling properties of fibers can be used in fiber-fed spectrographs to address such concerns by improving the overall stability of the instrument response (e.g., Smith et al. 2016).

Moving beyond the impact of fiber illumination and injection, stresses in the fiber system from bends, cabling, gluing, etc., can impart fiber-to-fiber output distortions similar to those in Figure 1. Using a slightly underfilled but centered input source, we show in Figure 2 how the far-field output patterns vary among a set of fibers glued into the 127 fiber MaNGA bundle designated MA 176, a replacement IFU assembly that was tested in 2017 December by one of us (M. Bershady) using one of the two “SDSS Test Stands.” The

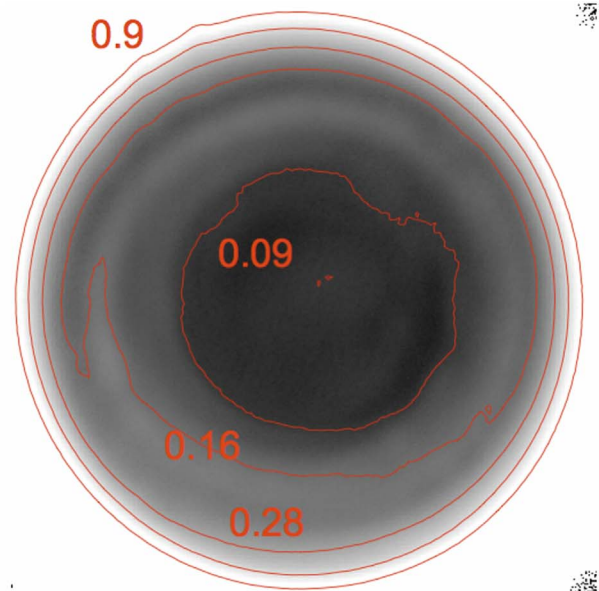



Figure 2. Far-field intensity distributions for fibers associated with a 127 fiber MaNGA bundle, MA 176. The left panel shows far-field patterns stretched to reveal core structure. The first eight fibers (from the top right) are free-roaming single sky fibers, while the remainder are members of the bundle and as such are subject to varying stresses. The exact location of an underfilled illumination source also varied from fiber to fiber in these lab measurements and is responsible for some degree of output variation. The image on the right shows the standard deviation of the intensity variation across this sample, divided by the mean intensity profile. The innermost contour indicates that variations within the fiber cores across this set of fibers is about 9%.

structural variation in far-field output across the entire set as measured within the core (right panel) is within 10% and results from a combination of changing stresses induced by the bundle assembly process and slight miscentering errors.

In what follows, we will be interested in estimating the impact of time-dependent *changes* in the near- and far-field fiber output. While lab experiments such as those behind Figures 1–2 use dedicated instrumentation setups, precise alignment and measurement tools, and robust procedures to characterize fiber output, it is challenging to connect those measurements first of all to the instrument performance on-sky, and second to their impact on the final data products, e.g., sky-subtracted 1D spectra. The details depend on the optical design (e.g., are micro-optics used that invert the near and far field?), the spectrograph design and as-built performance (e.g., grating uniformity), and the observational setup and data reduction approach. With these specifics in mind, the approach in the next section is to work backwards from obtained, on-telescope data to isolate the impact of the distortions described above.

4. Constraints on Fiber Systematics from MaNGA Calibration Data

MaNGA’s observing strategy (see Law et al. 2015) includes taking arc and flat-field calibration exposures every time a new cartridge is installed on the telescope (see Section 3.1) and again at the end of the observing sequence, typically 1–2 hr later. In addition, a number of defined tests, involving specific sequences of instrument movement followed by calibration exposures, were carried out (usually on cloudy nights) to study MaNGA instrument stability, especially under the influence of flexure. These calibration data teach us how the response function varies in the MaNGA instrument and allow us to isolate systematic contributions associated with changes in the condition—especially the stress state—of the fibers themselves.

In this way, we can relate the performance of MaNGA fibers to expectations for other instruments. All MaNGA data described below are publicly available from sdss.org.

We will examine data from MaNGA calibrations obtained as described below. While different experiments have been performed using different cartridges, within each experiment, we investigate exposures taken *with the same cartridge*, meaning that the same fibers, bundles, and other hardware are held fixed when looking for variations between exposures.

1. *Instrument Fixed at Zenith:* The telescope is pointed at zenith and the instrument is held stationary. The purpose of this experiment was to determine a base level of instrument stability under very modest changes in ambient conditions (e.g., temperature). Arcs and flat-field exposures were taken every 10 minutes over the course of one hour with the dome closed on 2015 June 15 (MJD 57,188). Plate 8603 was mounted in Cartridge #5 for this test.
2. *Telescope at Zenith, Instrument Rotated:* The telescope is pointed at zenith and the instrument (i.e., the Cassegrain-mounted cartridge and telescope-mounted spectrographs) is rotated 20° every 10 minutes for one hour with arcs and flat-field exposures taken at each rotator position. This experiment was carried out with the dome closed on 2015 May 24 (MJD 57,167) using Cartridge #1 with plate 8314 mounted.
3. *Simulated Observing Track:* At a declination of $+42^\circ$, the telescope tracks an imaginary field from -2 hr to $+2$ hr in hour angle (reaching a maximum altitude of 80°). We took nine arc and flat-field exposures during this track. At the end, calibrations at every other position along the track were taken again. Finally, the cartridge (Cartridge #3 with plate 8606) was unmounted and remounted to the telescope, and calibrations were taken again at every

other position along the original track. This experiment was carried out on 21 May 2015 (SJD 57,164).

4. *Same Cartridge, Different Pluggings*: From all exposures taken during regular MaNGA operations up to 2015 July, we compiled sets of calibration exposures obtained for the same cartridge but observed on different nights and with different plates. Each exposure therefore captures MaNGA fibers in dramatically different stress states (as well as at different temperatures, humidity levels, times of night, and seasons) since they were plugged into different holes on different plates.

4.1. Variation in the Wavelength-dependent Throughput Response

In this section, we analyze MaNGA flat-field exposures from the tests above to examine the possibly wavelength-dependent throughput stability of the instrument as would be described by the $g(\lambda, t)$ systematic terms in Section 2.

With integration times of 30 s, the flat-field exposures exceed counts of 20,000 e^- per pixel at nearly all wavelengths, thus providing a high-S/N characterization of the relative throughput. The flat-field MGFLAT frames store the extracted and wavelength-rectified 1D spectra of every fiber, normalized by a superflat composed of all fiber flats in the exposure.

4.1.1. Instrument Fixed at Zenith

Figure 3 shows the normalized fiber flats from two arbitrary fibers for five flat-field exposures taken over one hour with the telescope pointing at zenith. This experiment represents a floor in the MaNGA instrument stability since nothing is moving during this experiment. A boxcar smoothing of 50 pixels (about 55 Å) was applied and each fiber flat was normalized by an initial fiber-flat spectrum to highlight variations across the exposure sequence. We see maximum variations of $\sim 0.5\%$ among these exposures in the blue channels for these two fibers. Their variations in the red channel are even smaller. These variations likely owe to minor temperature or detector amplifier changes over the 1 hr exposure sequence. It is unclear if they are related to variable output among the flat-field lamps. The spectral signatures do not correlate between the two spectrographs, but the fibers on each spectrograph map to different portions of the field. So, these signatures may be explained by wavelength-dependent differences in the flat-field illumination pattern.

Over the observing sequence in this experiment, we compare the 1D flat-field spectrum obtained for each fiber from a given exposure to that in the previous exposure by computing the fractional difference spectrum, $(F_{i+1}/F_i - 1)$, where i denotes the exposure number. We compute the mean and standard deviation of the fractional difference spectrum over a restricted wavelength range defined for each channel in order to avoid low-noise behavior at the channel edges. For the blue channels, $\Delta\lambda_c = [4600, 5400]$ Å, and for the red channels, $\Delta\lambda_c = [7500, 9000]$ Å. This approach allows us to visualize the statistical behavior of all 1423 fibers over the duration of the exposure sequence.

The wavelength-averaged mean differences between fiber flats in the “Instrument Fixed at Zenith” experiment are shown in Figure 4. The deviations are typically 0.2%–0.3%, but we see that groups of fibers with similar fiber numbers often show coherent behavior within the same spectrograph (i.e., the same fiber number groupings deviate similarly comparing b1 to r1, or

comparing b2 to r2). Because the fiber number increases continuously along the pseudo-slit, we know that these groupings correspond to v-groove blocks associated with specific MaNGA bundles. The coherent behavior might result from changes within the fiber run or subtle alignment variations, perhaps as a result of changing temperature, in either the telecentricity of specific bundles at the focal plane or with the angular alignment of v-groove blocks into the collimator. Although, the observed changes do not seem to drift smoothly as might be expected with a temperature dependence. Detector variations cannot be the culprit because the detectors are different between each spectrograph’s two color channels. Subtle variations in the flat-field illumination pattern may also play a role, an explanation we return to below when examining field-dependent behavior.

We now turn to the scatter about the mean. We measure the width of the distribution of flux differences among all wavelength pixels within $\Delta\lambda_c$ between subsequent fiber-flat spectra. The statistic is meant to capture wavelength-dependent transmission variations after removal of a mean (IFU-dependent) throughput offset. We see in Figure 5 that when the MaNGA instrument and telescope is held fixed with minimal changes in ambient conditions, the 1σ scatter in throughput differences about a mean offset is roughly 0.3% (per 1.1 Å pixel).

4.1.2. Tests with Instrument and Telescope Motion

Having established a baseline for the stability of the MaNGA instrument, we now explore the impact of different types of instrument and telescope motion to help isolate fiber-based systematics. We first consider the “Telescope at Zenith, Instrument Rotated” experiment in which the telescope is held fixed but the instrument is rotated, thus sweeping the fibers around the focal plane. Note the gravity vector remains fixed—no fiber motion or stresses are induced in this experiment. For the b1 channel (other channels show similar behavior), Figure 6 reveals a somewhat larger variation in the mean fractional difference between fiber flats, 0.3%–0.5%, roughly twice as large as was observed in the sequence with the instrument and telescope held fixed. The 1σ scatter under instrument rotation is nearly identical to the fixed-instrument experiment and so we do not show it here.

After subtracting the fixed-instrument mean variation, we attribute the remaining 0.2%–0.3% deviations under instrument rotation to a combination of variations in the flat-field surface flux across the focal plane (i.e., nonflatness) and nontelecentricity (i.e., geometric FRD) caused by misalignment between the fibers at the focal plane and the calibration beam delivered by the flat-field system. Drory et al. (2015) calculate that the ferrule-plate interface leads to a geometric FRD effect in the MaNGA instrument that can account for a 4%–5% throughput loss, much more significant than observed here. One explanation for the difference is that the angular distribution of rays from the flat field is extremely uniform across the field. In this case, the random angular misalignment between the chief ray and a plugged fiber will remain constant even as the instrument is rotated. The resulting geometric FRD variation will be small, as we observe, while the 4%–5% losses reported by Drory et al. (2015) result from far larger angular misalignments from one fiber plugging to the next.

Figure 7 shows the field dependence of fiber-flat spectral differences averaged over each IFU at several rotator angles from this experiment. Some spatial coherence in the sign and amplitude of flux offsets across the plate is apparent, consistent




Figure 3. Fiber-flat spectra for FiberID #100 in both MaNGA spectrographs as observed in a series of five exposures taken over 1 hr with the telescope pointing at zenith and the instrument fixed. A boxcar smoothing of 50 pixels ($\sim 55 \text{ \AA}$) has been applied. The variation may be driven by flexure from subtle temperature changes or subtle changes in the output intensity of the calibration lamps.

with the idea that the flat-field pattern or varying angular distribution of flat-field flux across the plate contributes to the variation in Figure 6.

We now investigate the “Simulated Observing Track” experiment in which telescope motion, instrument rotation, and therefore flexure, all vary between exposures. The expected amount of motion and changing stress on the fibers remains small, however, because of the stiffness of the short (1 m) fiber cabling in the MaNGA cartridges. We estimate that fiber motion is limited to $\sim 3 \text{ cm}$ at a bend radius of 100 mm. This is on par with expectations for changing stresses on Nasmyth fiber instruments (Section 3.1).

Figure 8 again plots the mean difference between fiber flats for the b1 channel (other channels are similar) for the observing

sequence across the simulated observing track while Figure 9 plots the standard deviation. The mean variations now reach 1% and, because they again manifest as coherent deviations within fiber blocks, we can attribute the variability to flat-field illumination and instrument flexure affecting alignment angles at the pseudo-slit and focal plane. This interpretation is strengthened by Figure 10, which compares the rms variation in throughput offsets among IFUs at different rotator angle positions for both the telescope fixed, instrument rotated experiment (left panel) and the simulated observing track (right panel). A strong dependence on rotator angle is seen in both cases. Given the spatial coherence in mean offsets seen across the field (Figure 7), angular misalignments or spatial structure in the flat field may sweep across the plugged fibers as the




Figure 4. Wavelength-averaged mean differences between subsequent fiber-flat spectra for all MaNGA fibers (numbered arbitrarily) as observed in a sequence of exposures with the telescope and instrument fixed. Each exposure pair is represented by a different color. The deviations are typically 0.2%–0.3%. Coherent behavior at the same fiber number locations may be caused by flexure variations of specific v-groove blocks along the pseudo-slit or subtle changes in the flat-field illumination pattern.

instrument rotates, explaining the bulk of block-coherent variation in Figure 8.

However, note that compared to Figure 6, Figure 8 shows evidence for a slight increase in the fiber-to-fiber scatter *within* a block, at the level of 0.05%. This is an indication of stress-induced FRD. The wavelength-dependent scatter, however, remains nearly identical to what was observed when the instrument and telescope were held fixed (Figure 9).

The three tests analyzed so far indicate that changes in the overall throughput normalization is the dominant type of long-

wavelength-mode instrument response variation in MaNGA. Even accounting for instrument flexure, the level of variation reaches a maximum of $\sim 1\%$ and there is little evidence for any wavelength dependence (as would be revealed in the scatter about the mean offset; Figure 9).

4.1.3. Testing Induced Fiber Stresses

Our final test involves comparing different pluggings of the same fibers over many different nights. This test should reveal

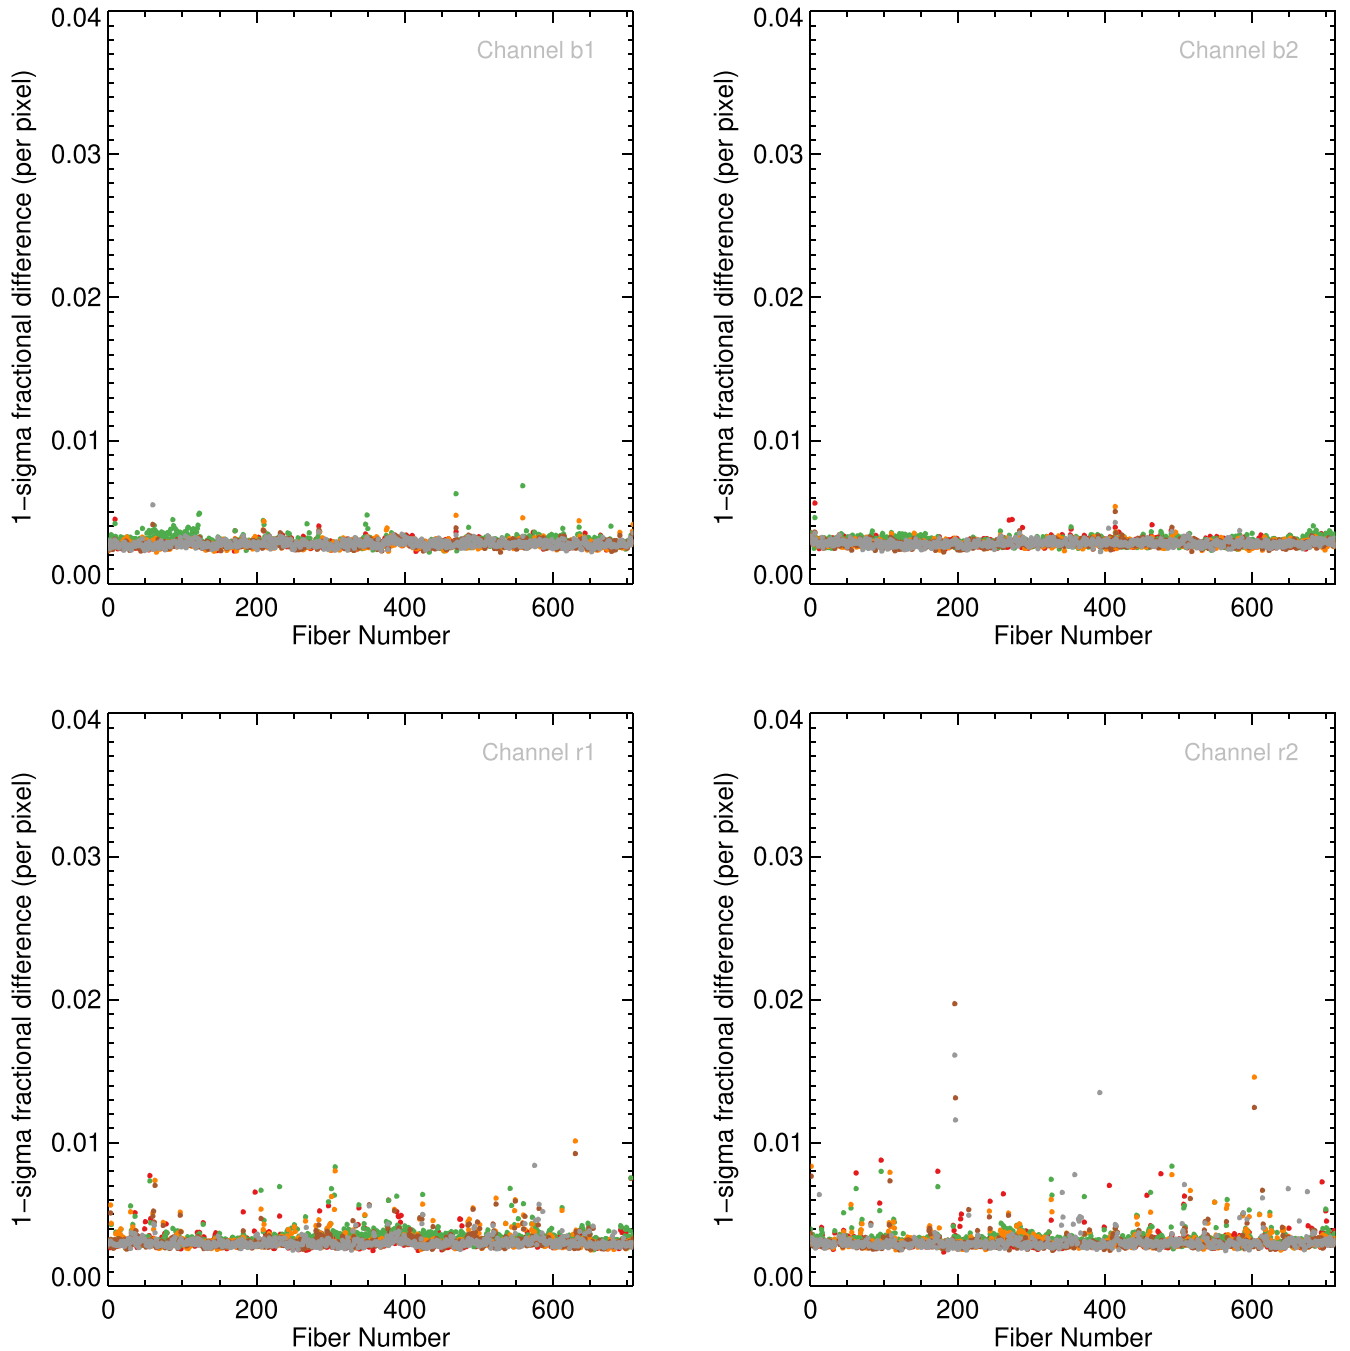


Figure 5. Standard deviation of difference spectra between subsequent fiber flats as computed for all wavelength pixels within $\Delta\lambda_c$ for all MaNGA fibers in the “Instrument Fixed at Zenith” experiment. The typical 1σ scatter is 0.3%.

the maximum throughput variations possible. Across the compiled sequence of flat-field exposures, the ambient conditions are different, the telescope position and gravity vectors are different, and, important for our purposes, the fibers are plugged in different locations on the focal plane, often with dramatically different bend radii and stress states.

Figure 11 shows the mean offsets between fiber-flat pairs taken from Cart #4 across eight different plates that were plugged and calibrated over the course of approximately six months (from MJD 57,133 to MJD 57,336). Deviations are now many percent instead of many tenths of a percent as we observed for the simulated observing track. While less distinct, block-to-block patterns are clearly visible and consistent across

channels on the same spectrograph. We can attribute some of this variation to dust buildup on the face of the fiber bundles, but dust contamination in MaNGA typically causes dramatic throughput reductions that are concentrated among small sets of fibers within specific IFUs. Building on our intuition from the previous experiments, a significant source of the variation may be deviations from telecentricity and angular misalignment as fiber bundles get plugged into different holes on different plates and focal plane locations. This conclusion quantitatively supports the geometric FRD estimates (several percent throughput variation) from Drory et al. (2015).

The fact that mean flux offsets of individual blocks show more scatter *within* a block (compare Figure 11 to Figure 8) is




Figure 6. Wavelength-averaged mean differences between subsequent fiber-flat spectra for all MaNGA fibers as observed in a sequence of exposures with the telescope pointed at zenith and instrument rotated 20° between each exposure over the course of an hour. Each exposure pair is represented by a different color. Only the b1 channel is shown because the other spectrograph channels show similar behavior. The variations are again coherent across fiber blocks but appear somewhat stronger, 0.3%–0.5%, compared to the experiment with no instrument rotation, which exhibited 0.2%–0.3% variations (Figure 4). The larger deviations may result from changing angles between the fiber input and the flat-field beam as rotation sweeps the fibers around the focal plane.

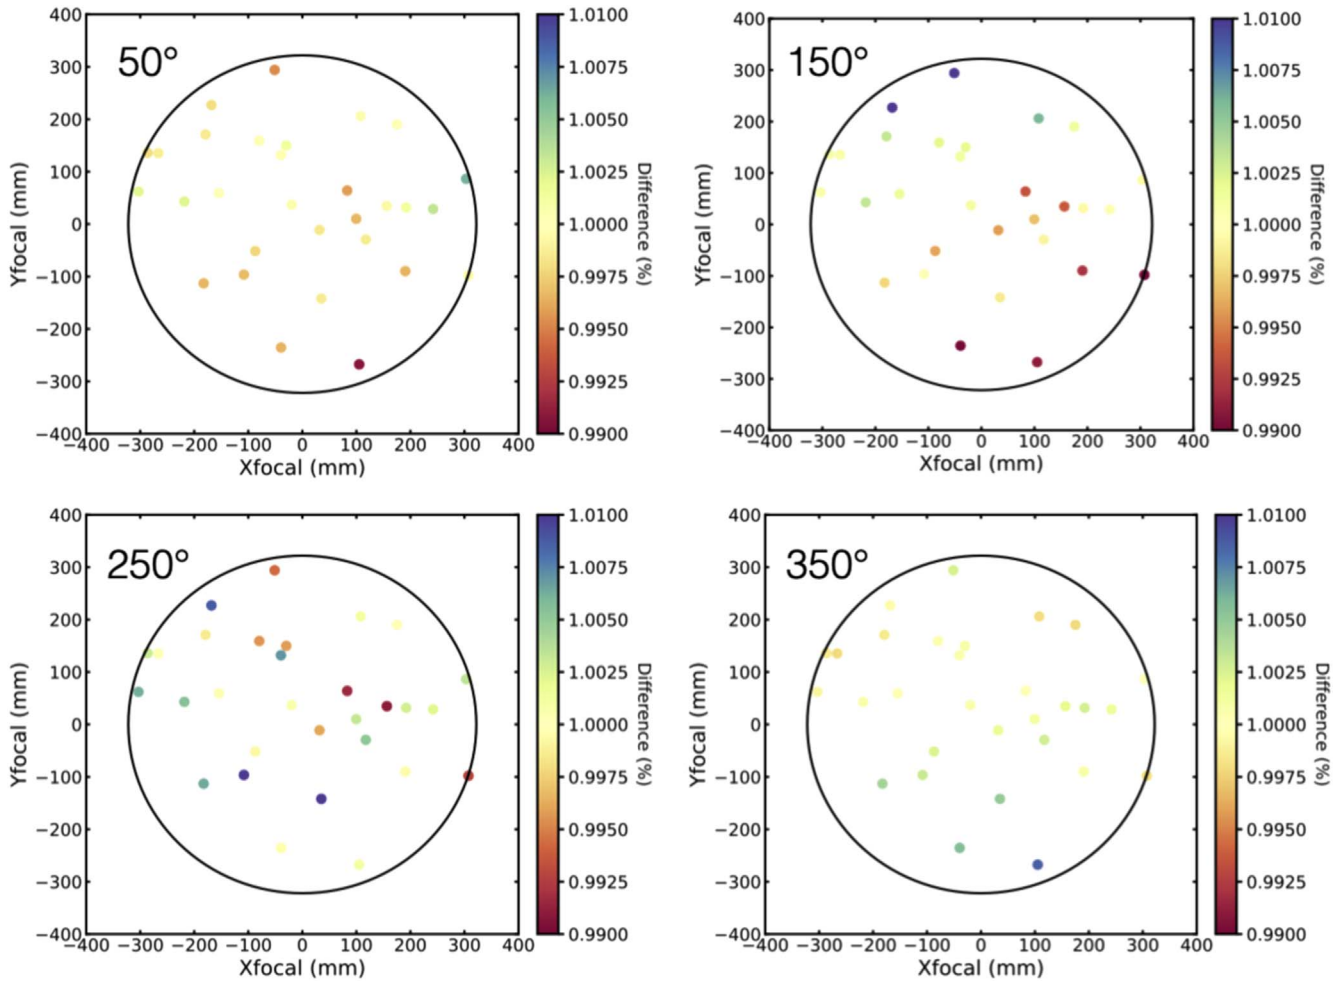


Figure 7. The relative mean difference compared to the first exposure in the average fiber-flat spectra measured within each science IFU and plotted against their location in the focal plane. Four specific rotator angles from the “Telescope at Zenith, Instrument Rotated” experiment are shown. A degree of coherent spatial variations across the plate is apparent, suggesting that differences in the flat-field illumination and angular distribution pattern as seen by different IFUs plays a role.

again a sign of stress-induced FRD affecting different fibers within the same bundle to varying degrees. We estimate this differential FRD effect on the throughput normalization at 1%–2% and note that it is random in nature, at least as observed when fibers are stressed in dramatically different ways from one measurement to the next.

Finally we look for wavelength-dependent variations by studying the scatter about these mean throughput differences in Figure 12 as we did for Figure 5. For the first time, this plot reveals a noticeable increase in scatter compared to the baseline value with the instrument in a stationary position. Most exposure pairs carry a $\sim 0.4\%$ scatter term on the blue channel with few




Figure 8. Wavelength-averaged mean fiber-flat differences across a simulated observing track with both telescope motion and instrument rotation. Each exposure pair is represented by a different color. Only the b1 channel is shown because the other spectrograph channels show similar behavior. Coherent block-to-block variations now reach 1%. Fiber motion and induced stresses are negligible because cabling in plugged cartridges limit fiber motion.




Figure 9. Standard deviation of the difference spectrum between pairs of fiber flats observed during a simulated observing track. Each exposure pair is represented by a different color (although hard to distinguish since most points overlap) and only the b1 channel is shown since the other channels are similar. Despite the significant importance of flexure in this experiment, the 1σ scatter is nearly identical to what was observed when the instrument and telescope were stationary (Figure 5).

exceeding 1%. The red channel performs better. Examples of the relative wavelength variations between exposures are shown in Figure 13. The stochastic wavelength-dependent behavior may be the result of changing far-field output of fibers under different stress states.

4.1.4. Summary of MaNGA Constraints on Fiber Systematics

In the previous sections we have used MaNGA calibrations obtained during various instrument tests to isolate the contribution from fibers to the throughput ($g(\lambda, t)$ terms) systematic response terms. Keeping in mind that the calibration light source fills the MaNGA fibers (at a beam angle that mimics the telescope pupil) and that the observed instrument response depends on MaNGA’s particular spectrograph design and performance, we want to emphasize the following points:

1. With minimal fiber motion between exposures, stress-induced FRD introduces a maximum 0.1% level stochastic variation in the throughput normalization. We find no evidence for increased wavelength-dependent scatter in throughput. These statements apply to fiber motion estimated at $\sim 1^\circ$ with a 100 mm bend radius, similar to what may be expected for telescope- and Nasmyth-mounted fiber instruments with carefully designed fiber runs.
2. Across maximum differences in fiber stress states—many tens of degrees at ~ 10 cm bend radii—stress-induced FRD may contribute overall throughput variations at the 1%–2% level.
3. Maximum fiber stresses are also associated with additional wavelength-dependent throughput scatter, which exhibits a combined level of 0.4% but occasionally reaches 1% (this depends on the spectrograph color channel). Stochastic in nature, this additional scatter may result from far-field variations as they propagate through the spectrograph.

5. On-sky Demonstrations

In the subsections that follow, we present several on-sky demonstrations of sky subtraction performance achieved with existing fiber instruments, including MaNGA on a 2.5 m telescope and FLAMES on an 8 m telescope. As discussed in Section 1, these demonstrations are relevant to future instrument designs for telescopes of all sizes because $\Delta\Omega$ is roughly conserved in all designs at fixed $f/\#$. Larger telescopes obviously collect more sky photons per unit area, but because of their finer plate scales, those photons are distributed over larger physical scales at the focal plane. Fibers with the same physical diameter collect the same sky flux at a given $f/\#$, independent of telescope diameter.

5.1. Pseudo-local Subtraction of Sky Lines in MaNGA

Given the significant effort to test and understand the MaNGA instrument limitations (see Drory et al. 2015; Yan et al. 2016), this instrument presents a valuable comparison point for sky subtraction performance with fibers, but with the important caveat that MaNGA was not designed to achieve a subpercent precision. Standard MaNGA galaxy targets are relatively bright compared to the sky foreground. In addition, flexure of its telescope-mounted spectrographs has been a long-known limitation in subtracting sky lines in MaNGA data. In




Figure 10. Comparison of the clear rotator angle dependence of the rms scatter in fiber flux offsets among IFUs in the experiment where the telescope is fixed and the instrument rotated (left panel) and the one meant to simulate an observing track (right panel).

Section 5.2, we exploit very long MaNGA integrations to place constraints on continuum residuals. In this section, we use all-sky data to show that even bright sky lines can be subtracted to a high accuracy if adjacent MaNGA fibers are used to sample the sky.

Our approach is motivated by the fact that executing sky nodding or cross-beam switching with MaNGA is problematic because (1) exposures less than a few minutes long are readnoise dominated on the Sloan 2.5 m telescope (see Figure 4 in Law et al. 2015), and (2) spectrograph flexure would require LSF modeling between exposures. Instead, in this section we perform a nonlocal but highly sampled sky subtraction approach. It is the equivalent of allocating 50% of fibers to sky regions located very near the science targets but maintaining *different* fibers for sampling sky and source.

We analyze 15 minute exposures on Plate 8069, which was observed on MJD 56,901. With the exception of guide stars and flux calibration standards, all fibers (both fibers contained within MaNGA bundles and also free-roaming “sky” fibers) are allocated to blank sky regions on this special plate. The standard MaNGA data reduction pipeline (v2_1_2) is run on the data through the 2D extraction stage under the assumption that the 17 science IFU bundles have been assigned to actual galaxy targets. A sky model for these reductions is constructed using nonlocal sky spectra from the 92 sky fibers alone. These sky fibers are distributed across the plate, but plugged near the IFUs. They are also colocated to the IFU fibers on the pseudo-slit. The resulting “fiber spectra” are flux calibrated, wavelength rectified, and sky subtracted, and they were used by the survey team to evaluate the MaNGA sky subtraction performance. By stacking these 1D fiber spectra vertically along the y-axis, we can build a 2D visual representation of sky residuals that is reminiscent of a 2D extracted and rectified slit spectrum. Examples of standard MaNGA nonlocal sky-subtracted fiber spectra of the sky in Plate 8069 in two red wavelength ranges (and incorporating fibers from both spectrographs) is shown in the upper panels of Figure 14.

The sky-line residuals apparent in the upper panels are representative of the standard MaNGA pipeline performance, although during normal survey operations, the automated pipeline yields a fraction of plates ($\sim 10\%$) with significantly worse or better performance (see Yan et al. 2016). Law et al. (2016) show that near the brightest sky lines (which for MaNGA occur near 9400 Å; right panels in Figure 14) the error

associated with sky subtraction residuals is 20%–25% greater than what would be expected from Poisson statistics alone (Figure 15). This level of subtraction, obtained without sky nodding and with just 6% of available fibers placed on sky, exceeds MaNGA’s science requirements (Yan et al. 2016), although a number of potential avenues have been identified for improving on it (Law et al. 2016) and were explored in work by Gu and collaborators.

We now apply a postprocessing step on the reduced Plate 8069 output to simulate how the sky-line subtraction performance improves when the sky sampling rate is increased from 6% of available fibers to 50%. From the reduced spectrum of each fiber along the pseudo-slit, we subtract the mean of the spectra associated with the two fibers on either side. Effectively, this step subtracts a sky model from locally adjacent fibers but is still “nonlocal” in the sense that different fibers are used to sample the source versus the sky. The near-perfect subtraction that results (second row of panels in Figure 14) demonstrates that the vast majority of MaNGA sky-line residuals are strongly and locally correlated along the pseudo-slit.

Our pseudo-local sky-line model is only applied to wavelength pixels near sky lines (identified by a threshold in inverse variance) and, to avoid confusion from correlated noise, we only display spectra corrected by a unique set of adjacent sky fibers (i.e., every third spectrum is shown). This allows a comparison of the average residuals from the standard pipeline versus pseudo-local sky subtraction (gray versus green curves in the bottom panels of Figure 14). The pseudo-local correction dramatically improves the worst offenders by removing the known LSF differences that are inherent in standard MaNGA reductions (due to both trace curvature effects and v-groove alignment; see Law et al. 2021), leading to a mean residual among this set of nearly 300 spectra that is at most 0.2% of the sky-line flux.

This result is reassuring in several ways. First, it demonstrates that a fiber instrument can achieve vanishing sky-line residuals even at red wavelengths dominated by OH lines. This was achieved *without* sky nodding. Over the 15 minute exposure window, the MaNGA pipeline’s instrument characterization (via calibration frames) and modeling of the LSF variation remained accurate enough to yield excellent performance even when different fibers were used to sample the sky versus the “source.” Most of the remaining MaNGA sky-line

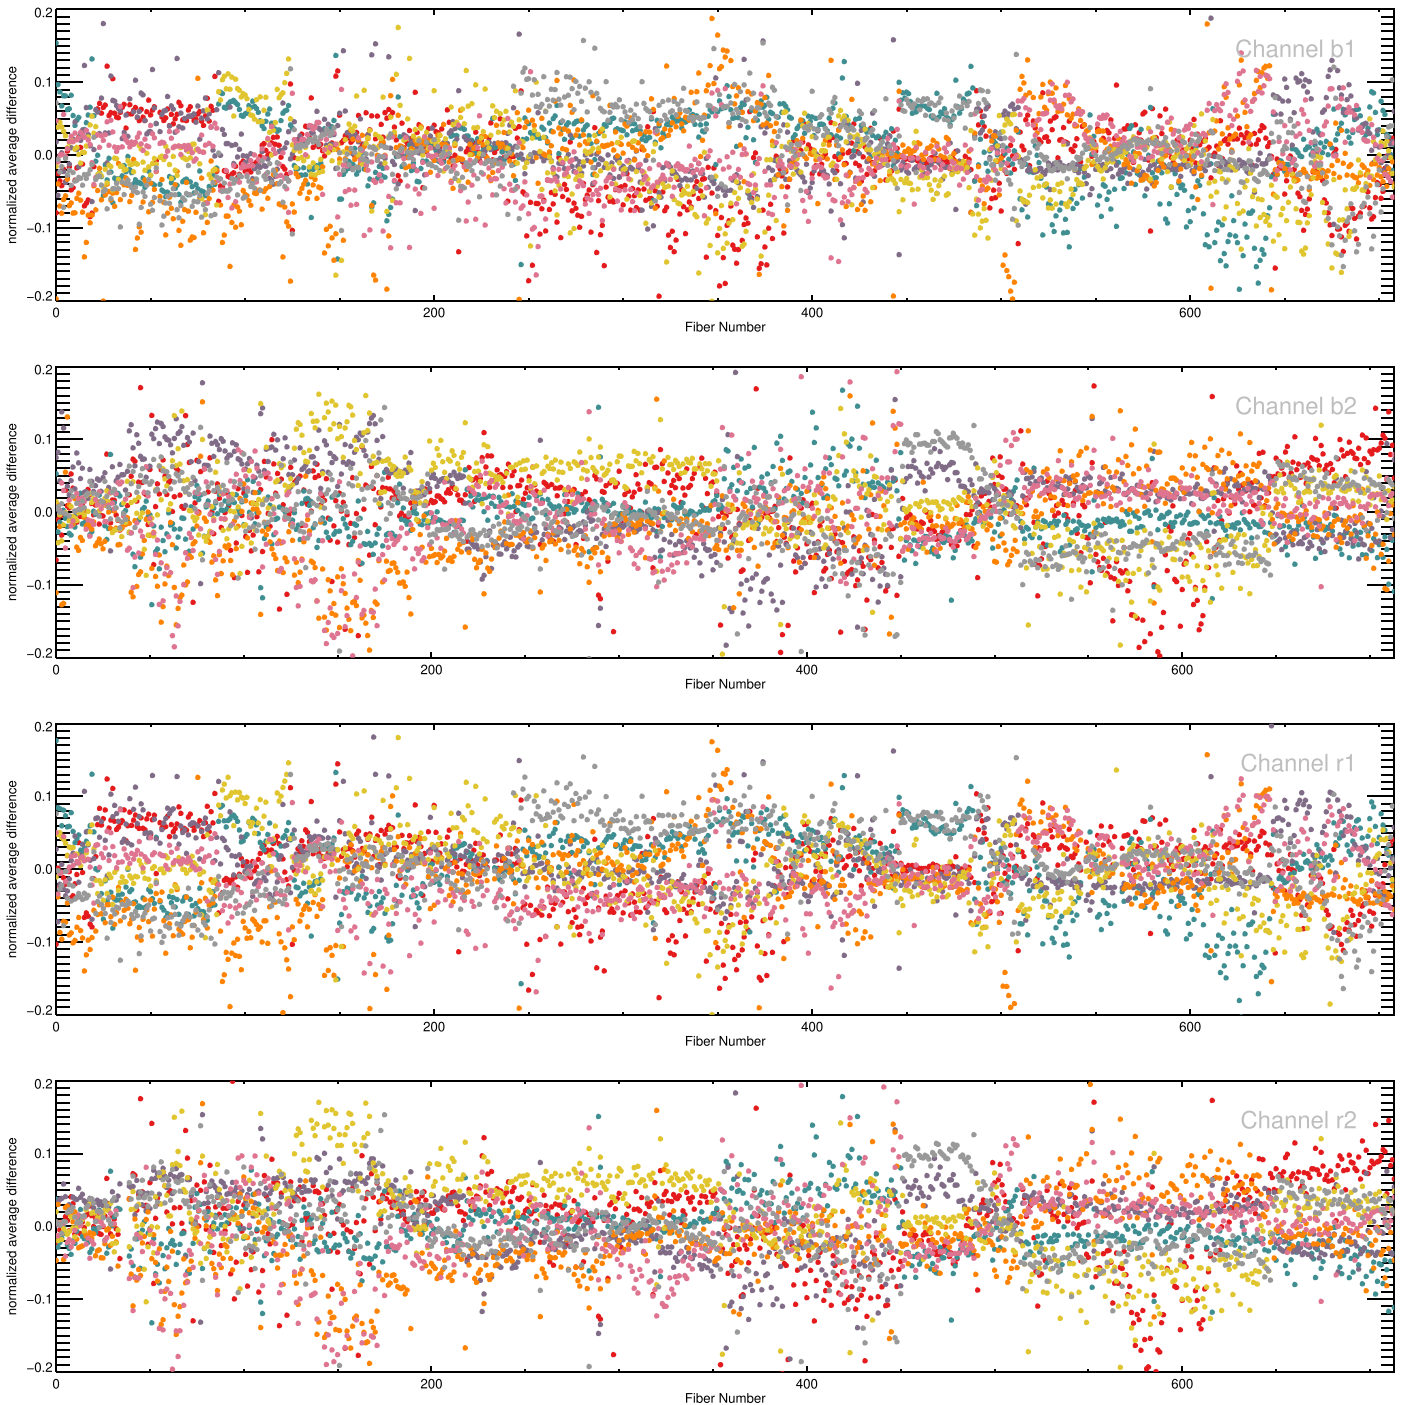


Figure 11. Wavelength-averaged mean differences between subsequent fiber-flat spectra for all MaNGA fibers as observed in a sequence of exposures compiled across many different pluggings and observed on different nights. Each exposure pair is represented by a different color. This is the maximum variation that would be expected. Note the order-of-magnitude larger scale on the y-axis. Instead of tenths of a percent, the throughput normalization is seen to vary at the several percent level.

residuals in the standard pipeline are therefore the result of inadequacies in the nonlocal sky model which does not fully account for changes in the instrument response among groups of fibers along the pseudo-slit. Furthermore, Law et al. (2021) showed that even the LSF resulting from the flexing MaNGA instrument can be calibrated to subpercent accuracy. Taken together, this suggests there are avenues for improvements to nonlocal fiber sky subtraction in future instruments that, by better sampling the slit function, can achieve near-perfect sky-

line performance without sky nodding and with only a modest sky fiber budget.

5.2. Deep MaNGA Coma Observations

With ~ 2.25 hr of exposure time, median MaNGA observations reach a per-fiber S/N of 5.1 \AA^{-1} at an r -band surface brightness of $22.5 \text{ AB arcsec}^{-2}$ (Yan et al. 2016). Scaling the S/N to 1.0, we can define MaNGA’s r -band surface brightness limit as $23.4 \text{ AB arcsec}^{-2}$. In the g -band, according to fitting

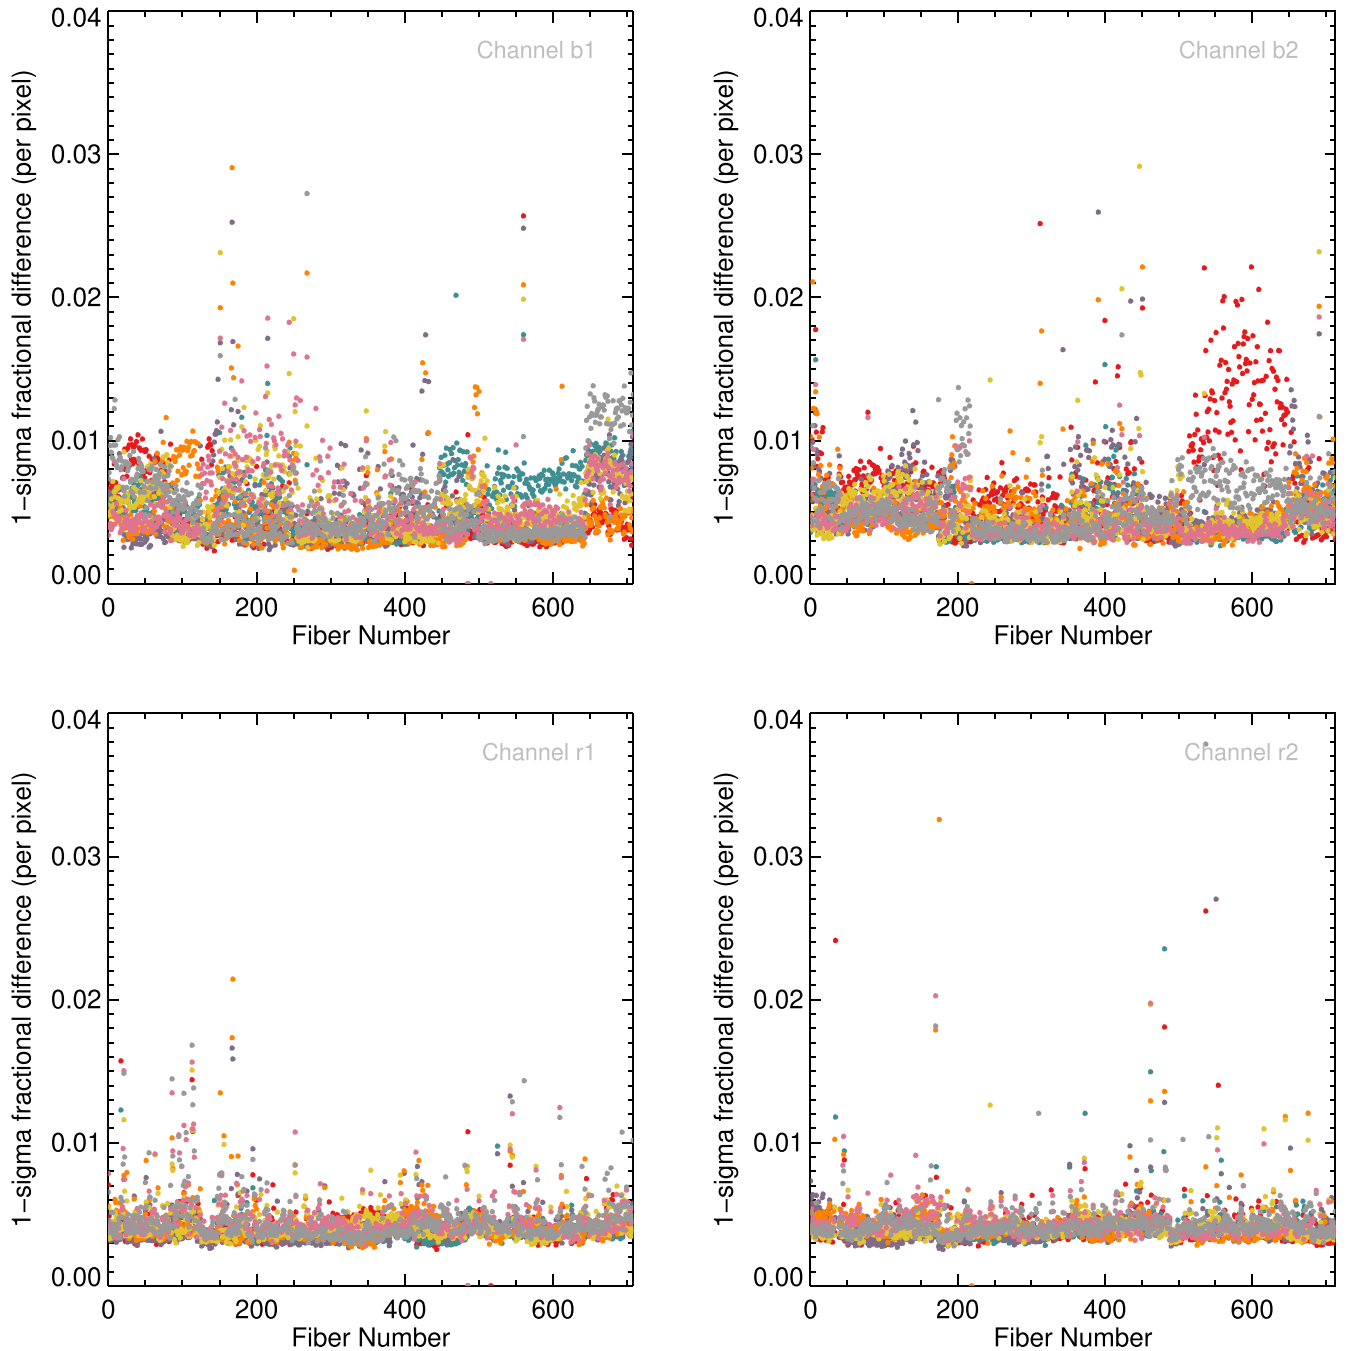


Figure 12. Standard deviation in the wavelength-dependent difference spectrum between pairs of fiber flats observed with fibers plugged into different plates and under different stress states and conditions. Each exposure pair is represented by a different color. Compared to the stationary instrument test (Figure 5), a greater degree of scatter is observed.

functions provided in Yan et al. (2016), the corresponding limit is $24.3 \text{ AB arcsec}^{-2}$. An ancillary MaNGA program targeting the Coma Cluster attempted to push far beyond this limit by exposing for longer and averaging multiple fibers within IFU bundles (see Gu et al. 2017). Because each fiber has a core diameter of $2''$, the total collecting solid angle of sets of, for example, 19, 37, and 127 fibers is equal to 59.7, 116.2, and 399.0 arcsec^2 , respectively. This “light-bucket” mode was used to observe diffuse low surface brightness ($24\text{--}26 \text{ AB arcsec}^{-2}$) sources in Coma including ultradiffuse dwarf galaxies (UDGs) and regions where the intracluster light (ICL) was detected in deep imaging observations (Abraham & van Dokkum 2014).

In the absence of systematic errors, one would expect that stacking individual spectra obtained with bundles of 19, 37, or 127 fibers would increase¹⁶ MaNGA’s median g -band surface brightness limit to 25.9, 26.3, and $26.9 \text{ AB arcsec}^{-2}$, respectively. Additionally, lengthening the exposure time by a factor of 6.0 (the Coma data reported in Gu et al. 2018, 2020 are based on 13.5 hr of exposure time) would yield expected median limits of $\mu_g = 26.9, 27.2, \text{ and } 27.9 \text{ AB arcsec}^{-2}$. By utilizing the darkest and clearest conditions and making use of

¹⁶ An exercise like this was carried out in Law et al. (2016). That analysis shows that in practice, the S/N increases more slowly than $\sqrt{N_{\text{fib}}}$ because of the intrinsic noise in the sky model.

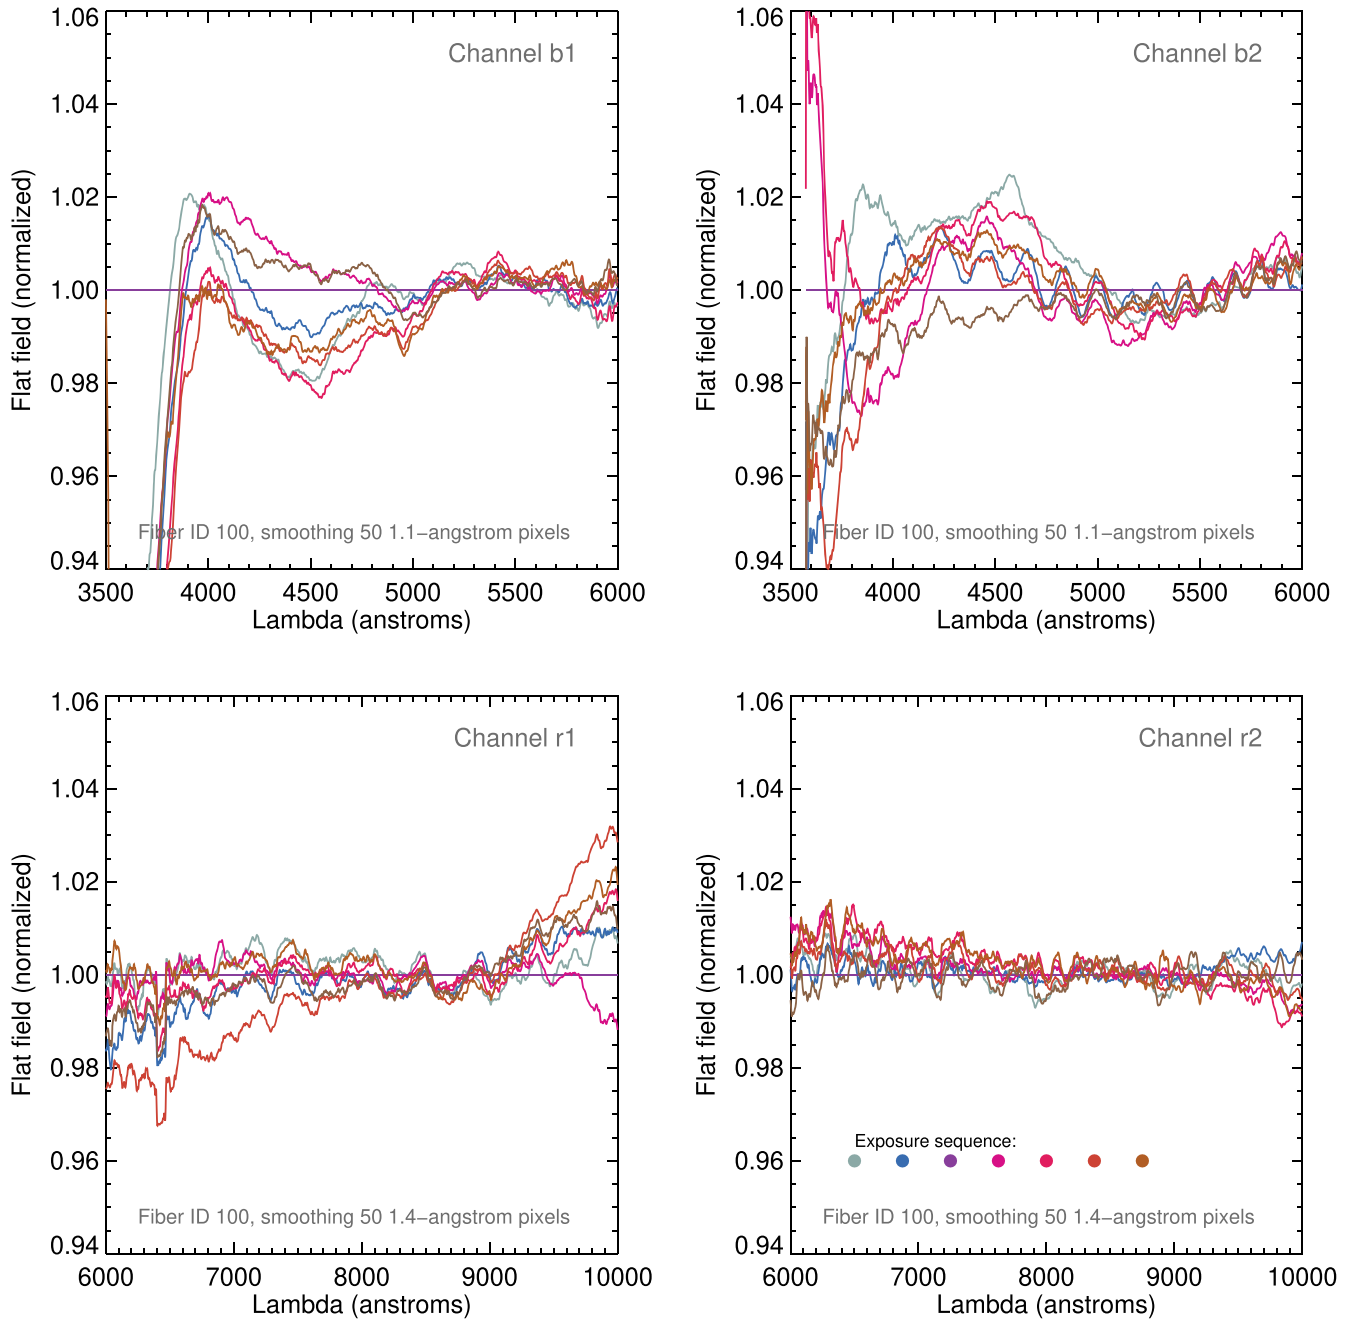


Figure 13. Relative throughput for two arbitrary MaNGA fibers from a sequence of flat-field exposures in which the fibers are plugged in different cartridges with very different stress states. An overall normalization term (illustrated in Figure 11) has been divided out. The stochastic wavelength-dependent response likely results from variations in the far-field illumination output by the fiber.

additional nod exposures (see below), Gu et al. (2018) demonstrate a surface brightness limit in stacks of 19 fibers of $\mu_g = 27.6 \text{ AB arcsec}^{-2}$ as determined by the $1\sigma \text{ \AA}^{-1}$ scatter in the sky-line-free 4500–5000 Å range. Because the surface brightness achieved actually exceeds expectations from scaling *median* MaNGA observations, it is clear that careful control of systematics was maintained.

Here, we highlight the MaNGA Coma observations to demonstrate the kind of sky subtraction performance that is possible with fiber instruments. The reduced fiber spectra are publicly available¹⁷ and we additionally provide stacked data products here. The first Coma plates designed and observed

were 8479 and 8480, but the resulting data are challenging to work with because the mix of bright and faint targets led to contamination from the wings of the detector PSF. In subsequent plate designs (PlateIDs 8953, 9051), bright targets (e.g., outer regions of central galaxies in Coma) were relegated to one spectrograph and faint targets (UDGs and the ICL) to the other, dramatically improving scattered light contamination among the faint sources.

The MaNGA instrument was never designed to achieve a subpercent precision. In addition to the attention paid to plate design, target-spectrograph allocation, and the limited use of sky nods, a number of custom adjustments were additionally required. On top of the standard 92 sky fibers, an additional 75 were used to sample the sky. Flat-field fiber tracing algorithms

¹⁷ At www.sdss.org.

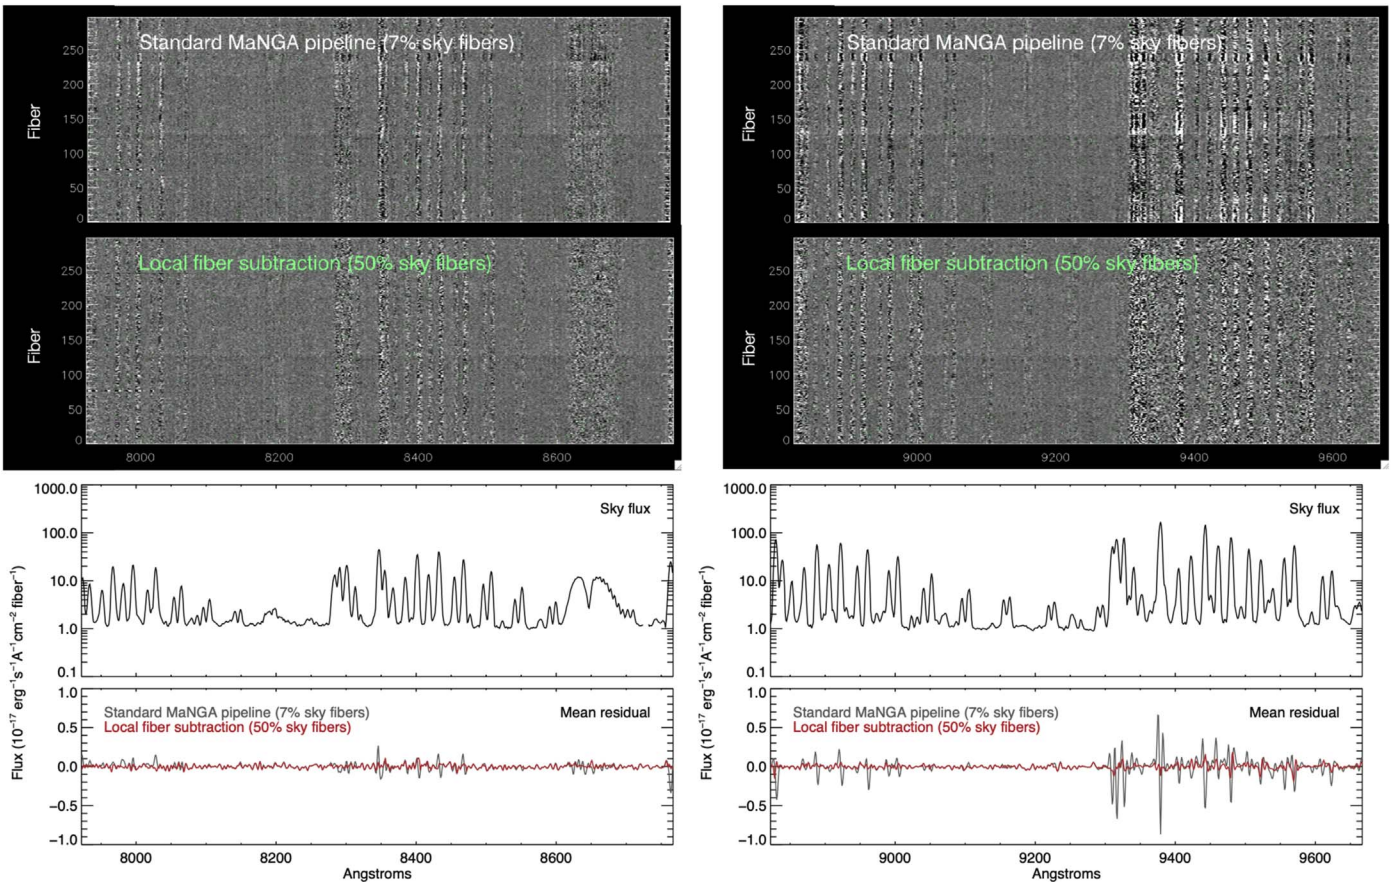


Figure 14. Pseudo-local subtraction of sky lines in MaNGA. The “images” in the upper panels are stacks of reduced and wavelength-rectified fiber spectra, where each row displays the 1D spectrum from a different fiber. The result in two red wavelength regions (left and right columns) for the standard pipeline output is shown in the upper panels. Systematic residuals near bright sky lines are apparent. A pseudo-local approach (second row of images) where the nonlocal sky sampling is increased to 50% of available fibers shows significant improvement. The mean sky flux and mean residual flux among this set of nearly 300 fibers is shown in the bottom panels.

were improved to remove biases at the 0.05 pixel level and the inverse variance determination and treatment of scattered light modeling were enhanced. Subtle detector issues including quadrant-dependent bias, a discovered seasonal bias drift, and 0.1% changes in amplifier gain between exposures all had to be addressed. These issues would affect any spectrograph (fiber or imaging). Different cartridges were used to minimize systematics associated with a particular cartridge’s pseudo-slit.

The Gu et al. (2018) final spectrum for UDG DF7 is reproduced in Figure 16, while the Coma ICL1 measurement (Gu et al. 2020) is shown in Figure 17. These figures present the spectra at wavelengths redward of 5700 Å, which are subject to brighter sky-line contamination and were not modeled by Gu and collaborators.

5.3. FLAMES-GIRAFFE Sky Subtraction Experiments

Observations to investigate sky subtraction performance with VLT’s FLAMES-GIRAFFE instrument were taken on 2012 March 8, and results were presented in Rodrigues et al. (2012) and Yang et al. (2013). In clear conditions with 0.9 seeing, the MEDUSA mode was used to deploy fibers over a 20′ region in the COSMOS field (see Capak et al. 2007). Seventy fibers were distributed in pairs separated by 12″ and oriented in the north–south direction. In three of these pairs, both fibers were placed on blank sky as determined by deep COSMOS photometry. GIRAFFE was configured to cover

820–940 nm with a spectral resolution of $R = 6500$. The target field was observed at low airmass (<1.2). During the observations, the moon angle was $\sim 28^\circ$ from the target field, with an estimated contribution of 50% to the sky continuum background flux.

The observations were carried out using cross-beam switching in which the telescope was nodded by 12″ between “A” and “B” positions oriented along the north–south axis. At any nod position, the object is always observed within one of two fibers in the pair with the sky observed by the other. The telescope guiding was switched off during all the dithered B exposures. In principle, the pointing error during offsets is better than 0.2″, which is much smaller than the diameter of the fibers (1.2″). However, in order to prevent risk of significant misalignments between the objects and the fibers during the B exposures, an AB–AB–AB sequence was preferred over an AB–BA–AB sequence. At either A or B positions, the integration time was set to ten minutes. The three consecutive AB sequences obtained represent a total effective exposure time of one hour. Immediately after observations, flat-field exposures were acquired.

5.3.1. FLAMES Data Reduction and Analysis

We have reduced the FLAMES data using custom IRAF and Python procedures. Raw frames were overscan corrected and bias subtracted. A scattered light term was removed along the




Figure 15. Automated MaNGA pipeline sky subtraction residual errors relative to Poisson expectations. Errors are determined from the measured variance in a set of blank, sky-subtracted fiber spectra taken for the all-sky Plate 8069. The worst sky lines are associated with residual systematic errors of 20%–25% above theoretical predictions. Figure adapted from Law et al. (2016).

dispersion and spatial directions. Individual fiber tracers were then extracted and flat fielded. For the present analysis, we focus on the “pure sky” observations. Note that we are not modeling second-order effects from atmospheric dispersion because these should largely cancel across the small fiber aperture over which the sky background can be assumed to be uniform. Time variations will be small given the experimental duration of ~ 1 hr.

We focus on data from Pair #1: (fibers 46 and 47), Pair #2: (fibers 134 and 135), and Pair #3: (fibers 22 and 43). These fibers span different areas of the available field. Example fiber sky spectra are shown in Figure 18, which indicates the wavelength regions we define for our analysis of continuum (marked in red) and sky-line (marked in blue) sky subtraction residuals.

The dithered exposures in a nodding sequence will be referred to as “A” and “B.” Within a fiber pair, the northern fiber will be denoted “fiber₁” and the southern fiber as “fiber₂.” The exposure associated with the northern fiber in the A position will be denoted “fiber_{1,A}.”

5.4. Offset-fiber Subtraction

We first consider “offset-fiber” subtraction, a basic implementation of nonlocal sky subtraction in which fiber₂ is assumed to represent the sky foreground in fiber₁. Sky flux in this case is observed during the same time as the object, but with a different fiber in a spatial location separated by $12''$. Following the same nomenclature of Equation (3), we define the sky-subtracted residual of the i th pair at position A or B as:

$$R(\lambda, A, i) = (\text{fiber}_{1,A}) - (\text{fiber}_{2,A}) \quad (6)$$

$$R(\lambda, B, i) = (\text{fiber}_{1,B}) - (\text{fiber}_{2,B}). \quad (7)$$

Residuals were measured separately for each of three exposures at each of the A and B positions. We take the average of the residuals measured at each position and divide by the median sky background. The results, which represent the combination of 30 minutes of integration time, are plotted in Figures 19 and 20 and listed in Table 1. Biases in the continuum and sky-line regions range from 3% to 10%. The rms is $\sim 11\%$ in the continuum and $\sim 25\%$ within sky lines. The

residuals quoted here include both systematics and Poisson errors.

5.5. Cross-beam Switching

To implement cross-beam switching with these data, we imagine that the source in position A is targeted by fiber₁ while fiber₂ collects only flux from the sky. After a telescope nod to position B, the source is moved to fiber₂ and fiber₁ is placed on sky. Because the estimator of F_{source} in position A is (fiber₁ – fiber₂), while for position B it is (fiber₂ – fiber₁), the sign of introduced systematics inverts between A and B. If we combine A and B, e.g., by taking the average of F_{source} , we expect the systematics to largely cancel, as described in Section 2.1.2.

We perform cross-beam switching for each fiber pair across the three sets of AB exposures. The results for each pair are shown in Figure 21 and tabulated in Table 1. As anticipated, cross-beam switching improves the quality of the sky subtraction, reducing the residual biases in accuracy to approximately 1% in both continuum and sky-line regions. The corresponding rms drops to 8%–10%. We can approximate the result of a longer set of exposures by further averaging all three fiber pairs and including the alternate case in which the imaginary source in position A is first assigned to fiber₂ instead of fiber₁ (i.e., reversing the sign of all operations). This experiment yields a final residual bias level of 0.1%–0.2% and an rms of $\sim 5\%$ in both the continuum and sky-line regions for an effective exposure time of 3 hr. These results are largely consistent with the findings reported in Rodrigues et al. (2012) and Yang et al. (2013).

6. Conclusions: Impact on Future Designs

Our analysis of MaNGA calibration data as well as on-sky observations taken with both MaNGA and VLT-FLAMES allows us to characterize the fiber contribution to systematic errors in these instruments. With careful attention paid to other nonfiber error sources (e.g., detector gain variations), both instruments have shown that a sky subtraction precision at the 0.1%–0.2% level can be achieved on sky. While both instruments image the telescope focus onto their fibers, designs employing pupil injection may still benefit by considering the systematics we study after inversion of the near field and far field.

The MaNGA calibration data allow us to investigate specific aspects of fiber systematics that can help identify design improvements in future instruments. When MaNGA fibers experience minimal changes in stress states, for example, $\sim 1^\circ$ angle variations at bend radii of 100 mm—values that may be expected during several-minute exposures in most fiber instruments—the data show 0.1% stochastic changes to the throughput normalization. This variation is independent of wavelength.

If the fibers are stressed more dramatically, i.e., with bend variations of several tens of degrees, we see larger systematic differences. Although the details will depend on the instrument, we can roughly associate these larger fiber stress changes to the worst-case scenario of what might be expected when comparing fiber states from different target configurations or telescope pointings in a fiber instrument. The changing stress states yield overall throughput variations at the 1%–2% level in MaNGA. In addition, they also impart a scatter term whose amplitude




Figure 16. Final 19 fiber spectrum of Coma Cluster UDG DF7 obtained using 13.5 on-source hours of MaNGA data by Gu et al. (2018). Here, we also show data redward of 5700 Å, which were not analyzed in Gu et al. (2018).




Figure 17. Final Coma intercluster light spectrum (target “ICL1”) obtained using 13.5 on-source hours of MaNGA data and combining 127 fibers. The observations are described and presented in Gu et al. (2020). Here, we also show data redward of 5700 Å, which were not analyzed in Gu et al. (2020).

can be as large as 1%. This may be the result of far-field variations propagating through the BOSS spectrograph.

These worst-case systematics are dominated by nonfiber effects (e.g., block alignment, flat-field regularity, etc.).

Controlling these is important in combination with sufficient calibrations and instrument designs that minimize fiber motion. Systematic errors can also be addressed by the design of the observations and calibrations. For example, fiber systematics in




Figure 18. Example FLAMES sky spectrum indicating wavelength regions used for the analysis of sky -subtraction residuals. The top panel indicates “continuum” regions in red, while the bottom panel indicates “sky-line” regions in blue.

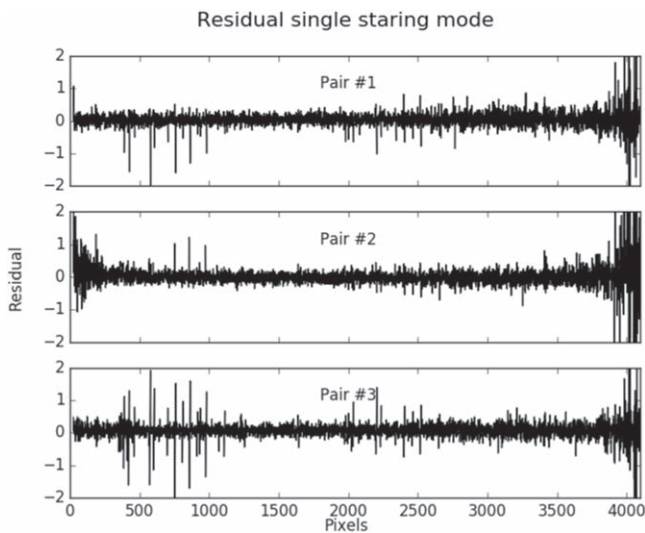


Figure 19. Average fractional residuals using “single stare subtraction” of fiber pairs observed in the A position. The total exposure time represented by this plot is 30 minutes.




Figure 20. Average fractional residuals using “single stare subtraction” of fiber pairs observed in the B position. The total exposure time represented by this plot is 30 minutes.

Table 1
FLAMES Sky Subtraction Results

| Fiber pair | Mode | Continuum | | Sky lines | | Effective t_{exp} (hr) |
|--------------------------|-----------------|-----------|-------|-----------|-------|---------------------------------|
| | | Accuracy | rms | Accuracy | rms | |
| Offset-fiber subtraction | | | | | | |
| Pair 1 | Offset-fiber | 0.030 | 0.114 | 0.025 | 0.227 | 0.5 |
| Pair 2 | Offset-fiber | 0.028 | 0.113 | -0.002 | 0.259 | 0.5 |
| Pair 3 | Offset-fiber | 0.066 | 0.107 | 0.062 | 0.286 | 0.5 |
| Cross-beam switching | | | | | | |
| Pair 1 | R_crossAB1 | 0.020 | 0.111 | 0.002 | 0.120 | 0.5 |
| Pair 2 | R_crossAB2 | 0.008 | 0.111 | 0.010 | 0.114 | 0.5 |
| Pair 3 | R_crossAB3 | 0.003 | 0.122 | 0.008 | 0.117 | 0.5 |
| All three pairs | R_CrossABBA_all | 0.001 | 0.054 | 0.002 | 0.058 | 3.0 |

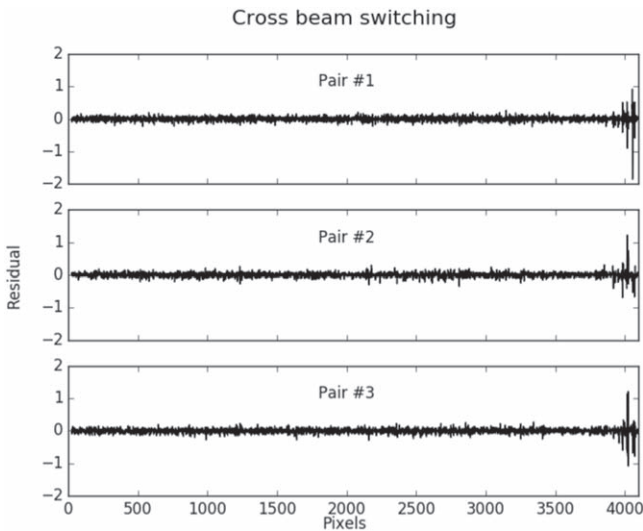


Figure 21. Residuals for each of three fiber pairs after applying cross-beam switching to a set of three AB exposure sequences. The amplitude of residuals is reduced compared to single stare sky subtraction results in Figure 19. Continuum and sky-line regions indicate accuracy biases of $\sim 1\%$, while the rms is approximately 9% for both the continuum and sky lines. These numbers reflect both instrument systematics and Poisson statistics.

MaNGA appear to vary stochastically. So, while these systematic errors will degrade both local and nonlocal sky subtraction performance, they will continue to average down over time if observations can be made to sample a range of stress states. Cross-beam switching provides a means to do this, but with a large penalty in sky fibers and overheads. Instead, it is worth emphasizing that the largest systematics result from the largest changes to the fiber stress state. Assuming these effects are repeatable, nonlocal sky subtraction techniques should greatly benefit from regular, sub 1% calibrations *at the stress states matched to the observations*. With such calibrations available, this work shows why achieving 0.1% level sky subtraction, even with nonlocal approaches, is possible. Although, it should be acknowledged that intrinsic spatial variations may be a limiting factor (see Yang et al. 2013).

Indeed, the fact that we measure fiber systematics to be minimal (at the 0.1% level) over the course of an observation is borne out by on-sky data both with MaNGA and FLAMES. Even with nonlocal sky subtraction and the inclusion of other important sources of error, Gu et al. (2017) used MaNGA to reach 0.2% level sky subtraction in long integrations. We have emphasized that because the sky foreground is a surface flux,

$A\Omega$ tells us that MaNGA’s success with large Ω fibers on a small telescope applies to small Ω fibers on a large telescope. Indeed, we confirm and build upon the earlier analyses by Rodrigues et al. (2012) and Yang et al. (2013) to show that the FLAMES data on the 8 m VLT demonstrate a 0.1%–0.2% level of sky subtraction accuracy.




In terms of sky *line* subtraction, the ability to measure and control the LSF is critical. Law et al. (2021) have demonstrated that the width of the final MaNGA LSF can be calibrated to a subpercent accuracy. Given flexure, focus, and other error sources, this indicates that even for a flexing instrument like MaNGA, fiber-based contributions to uncertainties in the LSF (the $k(\lambda, t)$ terms) can be controlled at the $\sim 0.2\%$ level.

We demonstrated one way of benefiting from this level of LSF stability using MaNGA, whose pseudo-slit design was never intended for such precision. We achieved Poisson-level nonlocal subtraction of the brightest sky lines by increasing the fraction of instrument fibers allocated to sky to 50%. Future instruments with more precisely formed pseudo-slits and precisely polished or cleaved fiber facets can expect a more uniform LSF and better sky-line subtraction.

This work was supported in part by a University of California Observatories Mini-Grant. Funding for the Sloan Digital Sky Survey IV has been provided by the Alfred P. Sloan Foundation, the U.S. Department of Energy Office of Science, and the Participating Institutions. SDSS acknowledges support and resources from the Center for High-Performance Computing at the University of Utah. The SDSS website is www.sdss.org. SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS Collaboration including the Brazilian Participation Group, the Carnegie Institution for Science, Carnegie Mellon University, the Chilean Participation Group, the French Participation Group, Harvard-Smithsonian Center for Astrophysics, Instituto de Astrofísica de Canarias, The Johns Hopkins University, Kavli Institute for the Physics and Mathematics of the Universe (IPMU) / University of Tokyo, Lawrence Berkeley National Laboratory, Leibniz Institut für Astrophysik Potsdam (AIP), Max-Planck-Institut für Astronomie (MPIA Heidelberg), Max-Planck-Institut für Astrophysik (MPA Garching), Max-Planck-Institut für Extraterrestrische Physik (MPE), National Astronomical Observatories of China, New Mexico State University, New York University, University of Notre Dame, Observatório Nacional / MCTI, The Ohio State University, Pennsylvania State University, Shanghai

Astronomical Observatory, United Kingdom Participation Group, Universidad Nacional Autónoma de México, University of Arizona, University of Colorado Boulder, University of Oxford, University of Portsmouth, University of Utah, University of Virginia, University of Washington, University of Wisconsin, Vanderbilt University, and Yale University.

ORCID iDs

Kevin Bundy  <https://orcid.org/0000-0001-9742-3138>
 David Law  <https://orcid.org/0000-0002-9402-186X>
 Kyle B. Westfall  <https://orcid.org/0000-0003-1809-6920>
 T. Sivarani  <https://orcid.org/0000-0003-0891-8994>
 Matthew Bershady  <https://orcid.org/0000-0002-3131-4374>
 Meng Gu  <https://orcid.org/0000-0002-4267-9344>
 Renbin Yan  <https://orcid.org/0000-0003-1025-1711>
 Namrata Roy  <https://orcid.org/0000-0002-4430-8846>
 Niv Drory  <https://orcid.org/0000-0002-7339-3170>

References

- Abraham, R. G., & van Dokkum, P. G. 2014, *PASP*, 126, 55
 Besuner, R., Allen, L., Baltay, C., et al. 2021, arXiv:2101.11794
 Blanton, M. R., Bershady, M. A., Abolfathi, B., et al. 2017, *AJ*, 154, 28
 Bolton, A. S., Schlegel, D. J., Aubourg, É., et al. 2012, *AJ*, 144, 144
 Bundy, K., Bershady, M. A., Law, D. R., et al. 2015, *ApJ*, 798, 7
 Bundy, K., Westfall, K., MacDonald, N., et al. 2019, *BAAS*, 51, 198
 Bundy, K., Westfall, K. B., MacDonald, N., et al. 2020, *Proc. SPIE*, 11447, 114471D
 Capak, P., Aussel, H., Ajiki, M., et al. 2007, *ApJS*, 172, 99
 Childress, M. J., Lidman, C., Davis, T. M., et al. 2017, *MNRAS*, 472, 273
 Cirasuolo, M., Afonso, J., Carollo, M., et al. 2014, *Proc. SPIE*, 9147, 91470N
 Colless, M., Dalton, G., Maddox, S., et al. 2001, *MNRAS*, 328, 1039
 Dalton, G., Trager, S. C., Abrams, D. C., et al. 2012, *Proc. SPIE*, 8446, 84460P
 de Jong, R. S., Bellido-Tirado, O., Chiappini, C., et al. 2012, *Proc. SPIE*, 8446, 84460T
 DESI Collaboration, Aghamousa, A., Aguilar, J., et al. 2016, arXiv:1611.00036
 D’Odorico, S., Cristiani, S., Dekker, H., et al. 2000, *Proc. SPIE*, 4005, 121
 Driver, S. P., Hill, D. T., Kelvin, L. S., et al. 2011, *MNRAS*, 413, 971
 Drory, N., MacDonald, N., Bershady, M. A., et al. 2015, *AJ*, 149, 77
 Eisenstein, D. J., Weinberg, D. H., Agol, E., et al. 2011, *AJ*, 142, 72
 Faber, S. M., Phillips, A. C., Kibrick, R. I., et al. 2003, *Proc. SPIE*, 4841, 1657
 Glazebrook, K., & Bland-Hawthorn, J. 2001, *PASP*, 113, 197
 Gloge, D. 1972, *BSTJ*, 51, 1767
 Goodwin, M., Brzeski, J., Case, S., et al. 2012, *Proc. SPIE*, 8446, 84467I
 Gu, M., Conroy, C., Law, D., et al. 2017, arxiv:1709.07003
 Gu, M., Conroy, C., Law, D., et al. 2018, *ApJ*, 859, 37
 Gu, M., Conroy, C., Law, D., et al. 2020, *ApJ*, 894, 32
 Gunn, J. E., Siegmund, W. A., Mannery, E. J., et al. 2006, *AJ*, 131, 2332
 Hammer, F., Barbuy, B., Cuby, J. G., et al. 2014, *Proc. SPIE*, 9147, 914727
 Hill, G. J. 2014, *AdOT*, 3, 265
 Hill, G. J., Gebhardt, K., Komatsu, E., et al. 2008, in ASP Conf. Ser. 399, *Panoramic Views of Galaxy Formation and Evolution*, ed. T. Kodama, T. Yamada, & K. Aoki (San Francisco, CA: ASP), 115
 Law, D. R., Yan, R., Bershady, M. A., et al. 2015, *AJ*, 150, 19
 Law, D. R., Cherinka, B., Yan, R., et al. 2016, *AJ*, 152, 83
 Law, D. R., Westfall, K. B., Bershady, M. A., et al. 2021, *AJ*, 161, 52
 Pasquini, L., Avila, G., Blecha, A., et al. 2002, *Msngr*, 110, 1
 Rodrigues, M., Cirasuolo, M., Hammer, F., et al. 2012, *Proc. SPIE*, 8450, 84503H
 Sharp, R., & Parkinson, H. 2010, *MNRAS*, 408, 2495
 Smee, S. A., Gunn, J. E., Uomoto, A., et al. 2013, *AJ*, 146, 32
 Smith, M. P., Wolf, M. J., Bershady, M. A., Fu, G., & Bing, L. 2016, *Proc. SPIE*, 9912, 99125B
 Steidel, C. C., Rudie, G. C., Strom, A. L., et al. 2014, *ApJ*, 795, 165
 Tamura, N., Takato, N., Shimono, A., et al. 2016, *Proc. SPIE*, 9908, 99081M
 Yan, R., Tremonti, C., Bershady, M. A., et al. 2015, arxiv:1511.01496
 Yan, R., Bundy, K., Law, D. R., et al. 2016, *AJ*, 152, 197
 Yang, Y., Rodrigues, M., Puech, M., et al. 2013, *Msngr*, 151, 10
 York, D. G., Adelman, J., Anderson Jr, J. E., et al. 2000, *AJ*, 120, 1579