Draft version September 8, 2023
Typeset using LaTeX **twocolumn** style in AASTeX62

Survey Operations for the Dark Energy Spectroscopic Instrument

Edward F. Schlafly,[1] David Kirkby,[2] David J. Schlegel,[3] Adam D. Myers,[4] Anand Raichoor,[3] Kyle Dawson,[5] Jessica Aguilar,[3] Carlos Allende Prieto,[6,7] Stephen Bailey,[3] Segev BenZvi,[8] Jose Bermejo-Climent,[8] David Brooks,[9] Axel de la Macorra,[10] Arjun Dey,[11] Peter Doel,[9] Kevin Fanning,[12] Andreu Font-Ribera,[13] Jaime E. Forero-Romero,[14] Juan García-Bellido,[15] Satya Gontcho A Gontcho,[3] Julien Guy,[3] ChangHoon Hahn,[16] Klaus Honscheid,[17,18,12] Mustapha Ishak,[19] Stéphanie Juneau,[11] Robert Kehoe,[20] Theodore Kisner,[3] Anthony Kremin,[3] Martin Landriau,[3] Dustin A. Lang,[21,22] James Lasker,[20] Michael E. Levi,[3] Christophe Magneville,[23] Christopher J. Manser,[24,25] Paul Martini,[17,26] Aaron M. Meisner,[11] Ramon Miquel,[27,13] John Moustakas,[28] Jeffrey A. Newman,[29] Jundan Nie,[30] Nathalie. Palanque-Delabrouille,[23,3] Will J. Percival,[22,21,31] Claire Poppett,[3,32,33] Constance Rockosi,[34,35] Ashley J. Ross,[17] Graziano Rossi,[36] Gregory Tarlé,[37] Benjamin A. Weaver,[11] Christophe Yèche,[23] and Rongpu Zhou[3]

(DESI Collaboration)

[1] *Space Telescope Science Institute, 3700 San Martin Dr, Baltimore, MD 21218, USA*
[2] *Department of Physics and Astronomy, University of California, Frederick Reines Hall, Irvine, CA 92697, USA*
[3] *Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA*
[4] *Department of Physics and Astronomy, University of Wyoming, Laramie, WY 82071, USA*
[5] *Department of Physics and Astronomy, The University of Utah, 115 South 1400 East, Salt Lake City, UT 84112, USA*
[6] *Instituto de Astrofísica de Canarias, C/ Vía Láctea, s/n, E-38205 La Laguna, Tenerife, Spain*
[7] *Universidad de La Laguna, Dept. de Astrofísica, E-38206 La Laguna, Tenerife, Spain*
[8] *Department of Physics & Astronomy, University of Rochester, Bausch and Lomb Hall, Rochester, NY 14627-0171, USA*
[9] *Department of Physics & Astronomy, University College London, Gower Street, London, WC1E 6BT, UK*
[10] *Instituto de Física, Universidad Nacional Autónoma de México, Cd. de México C.P. 04510, México*
[11] *NSF's NOIRLab, 950 N. Cherry Ave., Tucson, AZ 85719, USA*
[12] *The Ohio State University, Columbus, 43210 OH, USA*
[13] *Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Campus UAB, 08193 Bellaterra Barcelona, Spain*
[14] *Departamento de Física, Universidad de los Andes, Cra. 1 No. 18A-10, Edificio Ip, CP 111711, Bogotá, Colombia*
[15] *Instituto de Física Teórica (IFT) UAM/CSIC, Universidad Autónoma de Madrid, Cantoblanco, E-28049, Madrid, Spain*
[16] *Department of Astrophysical Sciences, Princeton University, Princeton NJ 08544, USA*
[17] *Center for Cosmology and AstroParticle Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, USA*
[18] *Department of Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, USA*
[19] *University of Texas, Dallas, 800 W Campbell Rd, Richardson, TX 75080, USA*
[20] *Department of Physics, Southern Methodist University, 3215 Daniel Avenue, Dallas, TX 75275, USA*
[21] *Perimeter Institute for Theoretical Physics, 31 Caroline St. North, Waterloo, ON N2L 2Y5, Canada*
[22] *Department of Physics and Astronomy, University of Waterloo, 200 University Ave W, Waterloo, ON N2L 3G1, Canada*
[23] *IRFU, CEA, Université Paris-Saclay, F-91191 Gif-sur-Yvette, France*
[24] *Astrophysics Group, Department of Physics, Imperial College London, Prince Consort Rd, London, SW7 2AZ, UK*
[25] *Department of Physics, University of Warwick, Coventry, CV4 7AL, UK*
[26] *Department of Astronomy, The Ohio State University, 4055 McPherson Laboratory, 140 W 18th Avenue, Columbus, OH 43210, USA*
[27] *Institució Catalana de Recerca i Estudis Avançats, Passeig de Lluís Companys, 23, 08010 Barcelona, Spain*
[28] *Department of Physics and Astronomy, Siena College, 515 Loudon Road, Loudonville, NY 12211, USA*
[29] *Department of Physics & Astronomy and Pittsburgh Particle Physics, Astrophysics, and Cosmology Center (PITT PACC), University of Pittsburgh, 3941 O'Hara Street, Pittsburgh, PA 15260, USA*
[30] *National Astronomical Observatories, Chinese Academy of Sciences, A20 Datun Rd., Chaoyang District, Beijing, 100012, P.R. China*
[31] *Waterloo Centre for Astrophysics, University of Waterloo, 200 University Ave W, Waterloo, ON N2L 3G1, Canada*
[32] *Space Sciences Laboratory, University of California, Berkeley, 7 Gauss Way, Berkeley, CA 94720, USA*
[33] *University of California, Berkeley, 110 Sproul Hall #5800 Berkeley, CA 94720, USA*
[34] *Department of Astronomy and Astrophysics, University of California, Santa Cruz, 1156 High St., Santa Cruz, CA 95065, USA*
[35] *University of California Observatories, 1156 High St., Sana Cruz, CA 95065, USA*
[36] *Department of Physics and Astronomy, Sejong University, Seoul, 143-747, Korea*
[37] *University of Michigan, Ann Arbor, MI 48109, USA*

ABSTRACT

The Dark Energy Spectroscopic Instrument (DESI) survey is a spectroscopic survey of tens of millions of galaxies at $0 < z < 3.5$ covering 14,000 sq. deg. of the sky. In its first 1.1 years of survey operations, it has observed more than 14 million galaxies and 4 million stars. We describe the processes that govern DESI's observations of the 15,000 fields composing the survey. This includes the planning of each night's observations in the afternoon; automatic selection of fields to observe during the night; real-time assessment of field completeness on the basis of observing conditions during each exposure; reduction, redshifting, and quality assurance of each field of targets in the morning following observation; and updates to the list of future targets to observe on the basis of these results. We also compare the performance of the survey with historical expectations and find good agreement. Simulations of the weather and of DESI observations using the real field-selection algorithm show good agreement with the actual observations. After accounting for major unplanned shutdowns, the dark time survey is progressing about 7% faster than forecast, which is good agreement given approximations made in the simulations.

Keywords: Redshift surveys (1358), Spectroscopy (1558), Observatories (1147), Telescopes (1689), Cosmology (343)

1. INTRODUCTION

The Dark Energy Spectroscopic Instrument (DESI) began a five year survey to measure redshifts of tens of millions of galaxies and quasars on May 14, 2021. Galaxies and quasars are selected to cover $0 < z < 3.5$ over 14,000 sq. deg. of the sky. The resulting redshifts will be used to measure the expansion history of the universe and the growth of structure to better understand the nature of dark energy (DESI Collaboration et al. 2016a).

The DESI survey consists of three programs. The dark program targets luminous red galaxies, emission line galaxies, and quasars, and covers $0.4 < z < 3.5$ (Zhou et al. 2023; Raichoor et al. 2023; Chaussidon et al. 2023). Dark program fields are observed whenever conditions are good and represent 90% of DESI's effective observing time. The bright program targets a magnitude-limited sample of bright galaxies with $0 < z < 0.4$, as well as Milky Way stars, and is observed when conditions are not good enough to observe dark fields (Hahn et al. 2022; Cooper et al. 2022). The combination of the dark program and the bright program are called the "main survey." Finally, a backup program observes bright stars and is only observed when conditions are too poor to observe bright program fields.

These programs consist of a number of "tiles," which are the combination of a location on the sky and an assignment of fibers to locations in the field. The aim of operations is to observe these fields as efficiently as possible. Two strategic goals drive many of the choices made in the DESI operations. First, we intend to observe in a "depth-first" mode, where we observe a given part of the sky to completion and never return to it, rather than a "breadth-first" mode where observations are spread over the full footprint each year. Second, we aim to ~~identify~~ observe $z > 2.1$ quasars ~~in initial observations and prioritize them for re-observation in subsequent exposures covering the same area~~ four times each to improve the signal-to-noise ratio in the Ly-α forest, which enters into the DESI spectral coverage for redshifts $z > 2.1$ (DESI Collaboration et al. 2016a). This choice means that no observations may overlap a past observation until the $z > 2.1$ quasars have been identified, placing pressure on the survey to rapidly and robustly deliver quasar redshifts. These two goals are in tension with one another—the depth-first goal means that we intend to make overlapping observations quickly to finish parts of the sky, while the goal of identifying $z > 2.1$ quasars means that we must complete analysis of observations before we can make overlapping observations.

Reconciling these goals means bringing together a large number of different processes and analyses ~~together~~ on a daily basis to execute the survey. We focus in this paper on the survey in the time frame from 2021–05–14, the first day of the main survey, to 2022–06–14, when the Contreras wildfire temporarily shut down the survey. Figure 1 shows the area of sky observed by DESI in the dark and bright programs during this period. We describe the DESI instrument in §2, and elaborate on this broad survey strategy in §4. We then describe the different observational and analysis processes that take place on a near-daily basis in order to enable the survey strategy in §5. The "merged target list", which plays a central role in tracking the current state of DESI observations, is described in §6. The DESI sky footprint is defined in §3. The delivered seeing, transparency, sky brightness, and uptime over the first 1.1 years are described in §7. We detail simulations of the survey in §8 and compare them with the observed survey performance to date. Finally, we conclude in §9. The code and data used to produce the tables and figures in this paper are available at https://doi.org/10.5281/zenodo.8010818.

2. THE DARK ENERGY SPECTROSCOPIC INSTRUMENT

The Dark Energy Spectroscopic Instrument is a 5000-fiber multi-object spectrograph on the Mayall telescope at Kitt Peak. The instrument and survey were conceived, designed, and built over a roughly ten year period from 2010–2020 (Levi et al. 2013; DESI Collaboration et al. 2016b, 2022). DESI was designed to measure the expansion history of the universe using the three-dimensional clustering of galaxies and the Lyman-alpha forest over the course of a five-year survey (DESI Collaboration et al. 2016a). The instrument collects light from astronomical sources with the 4-m Mayall primary mirror and focuses it through the new corrector onto a 3.2° diameter focal plane (Miller et al. 2023). 5000 robotically actuated fibers fill this focal plane (Silber et al. 2023), piping light through fibers to an array of ten high throughput spectrographs with three channels each spanning the wavelength range 3600–9800 Å.

The focal plane is divided into ten "petals," nearly identical wedges of the focal plane. Each petal has 500 positioners, connects to one spectrograph, and contains a guide-focus array imaging camera (GFA). Four of the petals' GFAs are dedicated to determining the focus of the instrument and deliver out-of-focus images. The other six deliver in-focus images and are used for guiding, point spread function measurements, and throughput measurements. The petals are designed to function independently of one another, so that problems with one petal do not affect any other petals.

The main survey will observe millions of stars and galaxies over the course of five years. Initial results from the survey validation program are now available (DESI Collaboration et al. 2023a,b). The primary targets are quasars (Yèche et al. 2020; Chaussidon et al. 2023), emission line galaxies with $0.6 < z < 1.6$ (Raichoor et al. 2020, 2023), luminous red galaxies with $0.4 < z < 1$ (Zhou et al. 2020, 2023), bright galaxies with $z < 0.4$ (Ruiz-Macias et al. 2020; Hahn et al. 2022), and stars (Allende Prieto et al. 2020; Cooper et al. 2022). Targeting catalogs (Myers et al. 2023) for these images were drawn mainly from Data Release 9 of the DESI Legacy Imaging Surveys (Dey et al. 2019), which included imaging from the Dark Energy Camera on the Blanco telescope (Flaugher et al. 2015), the 90prime imager on the Bok telescope (Williams et al. 2004; Zou et al. 2017), and the Mosaic3 imager on the Mayall telescope (Dey et al. 2016). Targeting catalogs also incorporated flux and astrometric measurements from Gaia, the Wide-field Infrared Survey Explorer, and the Siena Galaxy Atlas (Gaia Collaboration et al. 2016; Cutri et al. 2013; Meisner et al. 2018; Schlafly et al. 2019; Moustakas et al. 2023).

Each night, DESI observes roughly twenty tiles containing ∼100,000 sources. By the following morning, the offline pipeline automatically calibrates the resulting exposures, extracts the sources' spectra, subtracts background light, and fits the redshifts of the targets (Guy et al. 2023; Bailey et al. 2023). The performance of the pipeline was confirmed via a collaboration-wide effort to visually inspect tens of thousands of spectra and their derived redshifts (Lan et al. 2023; Alexander et al. 2023).

The DESI ~~guide-focus array cameras~~ GFAs and sky monitor provide real-time information on the seeing, transparency, and sky brightness seen by the Mayall (DESI Collaboration et al. 2022; Tie et al. 2020). This allows the DESI system to tune the length of exposures to achieve target depths; DESI closes the shutter and reads out the exposure when we have achieved the target signal to noise ratio (Kirkby et al. 2023, §5.7). This process allows us to produce spectra of relatively homogeneous quality even in changing conditions.

3. SURVEY FIELDS

The Dark Energy Survey Instrument Final Design Report calls for a baseline survey of 14,000 sq. deg. (DESI Collaboration et al. 2016a), with a science fiber density of $\sim 3000/\text{deg}^2$ for the dark program and $\sim 700/\text{deg}^2$ for the bright program. Given the DESI fiber density of $\sim 600/\text{deg}^2$, this corresponds to each region of the sky being covered by five observations for the dark program and one observation for the bright program. The bright and dark programs nevertheless require more passes to target multiple galaxies within a fiber patrol radius and to obtain reasonable completeness on lower priority main survey programs. We describe here the specific implementation of these broad requirements for the dark and bright programs.

We define a set of 9929 dark tiles and 5676 bright tiles that cover 14,200 sq. deg.: 9800 sq. deg. in the North Galactic Cap and 4400 sq. deg. in the South Galactic Cap. Each tile is a location on the sky that DESI will observe. These tiles are distributed among several passes where each pass consists of 1,427 non-overlapping

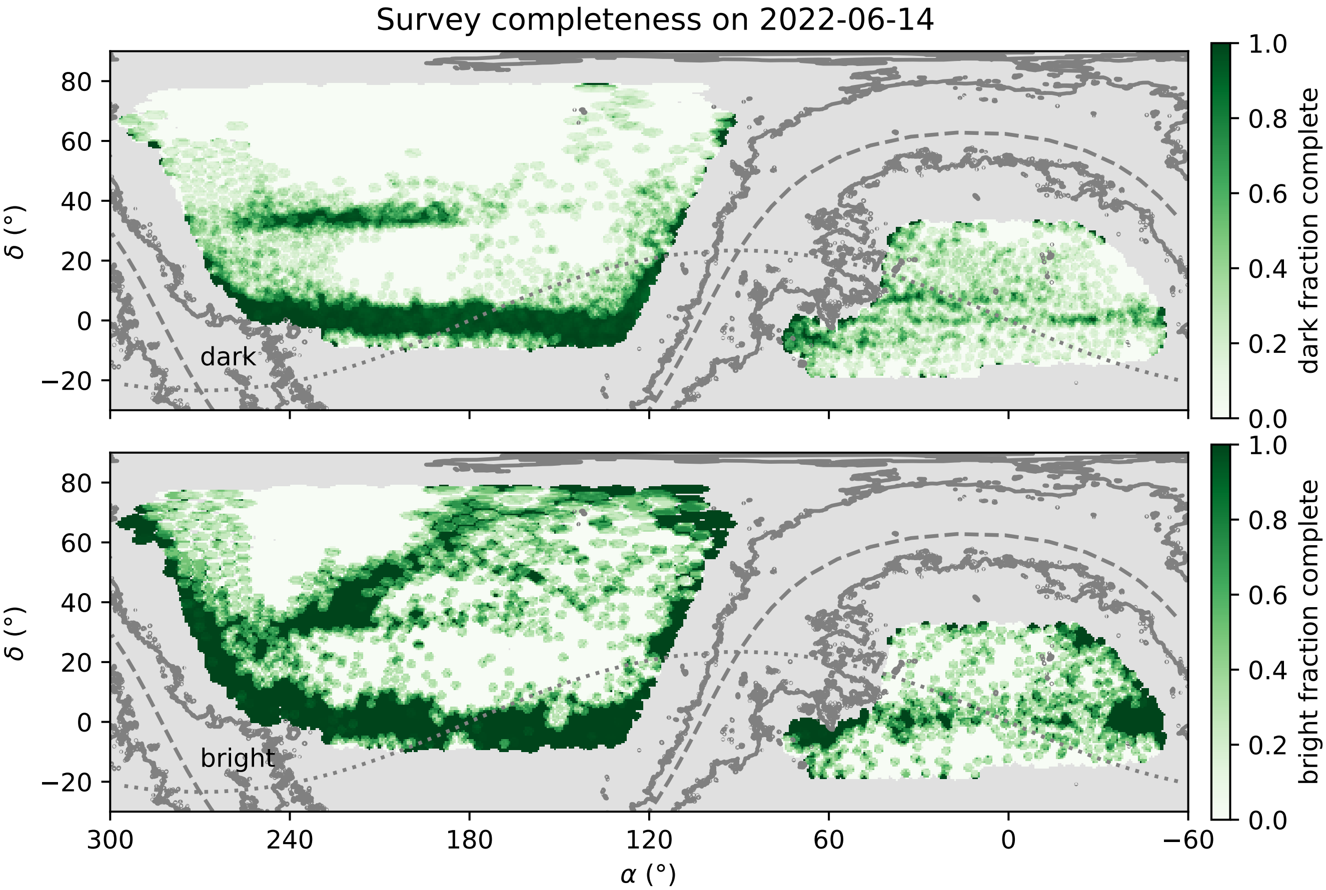

Figure 1. Survey completeness on 2022–06–14, in the dark (top) and bright (bottom) programs. Green areas are completely finished, while white areas are unfinished. Areas not included in the footprint are in gray. Regions with $E(B-V) > 0.3$ are outlined by the solid contours. The dotted and dashed lines show the ecliptic and Galactic planes. The survey aims to start observations near $\delta = 0°$ and build out. Notable deviations from that pattern are areas just above $\delta = 30°$, which are driven by needing to avoid strong winds from the south, and a region 50° from the ecliptic in the bright program in the north, driven by moon avoidance.

tiles. Approximately 75% of the footprint can be reached by a DESI fiber in a tile in a particular pass. The dark program consists of seven such passes, rotated with respect to one another to fill in gaps between the tiles, while the bright program consists of four such passes. This leads to an average coverage of 5.2 for the dark program and 3.2 for the bright program.

The pattern of tiles in a single pass is given by the Hardin et al. (2000) icosahedral tiling with 4112 tile centers distributed over the full sphere. This tiling matches the size of the DESI focal plane closely and provides a uniform distribution of tiles with the additional feature that no two tiles overlap one another within a single pass. The fraction of the sky accessible to a given number of tiles for the seven pass dark program and four pass bright program is shown in Figure 2. The geometry of the regions of relatively high and low coverage is complicated, and is shown for the seven-pass dark program in Figure 3.

The goal of the DESI tile selection was to select a large, contiguous region that could be efficiently observed for extragalactic targets as part of a year-round survey from Kitt Peak. These objectives imply limits on declination to avoid tiles that are only available at high airmass, and limits on extinction and Galactic latitude to avoid regions where extragalactic targets are both extinguished and more often blended with Milky Way stars.

We define the footprint as follows:

1. In the footprint of the DESI Legacy Imaging surveys Data Release 9
2. $-18° < \delta < 77.7°$
3. $b > 0°$ or $\delta < 32.2°$
4. $|b| > 22°$ for $-90° < l < 90°$, otherwise $|b| > 20°$

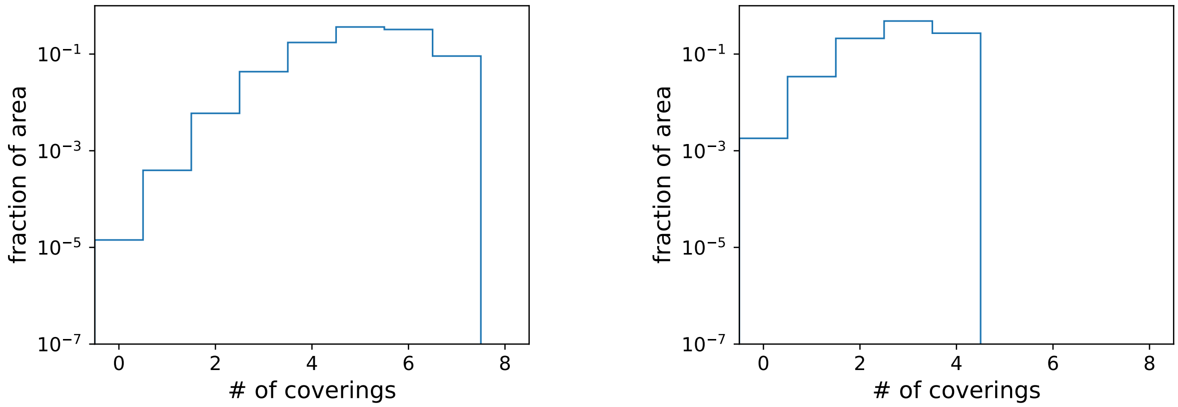

Figure 2. The fraction of the sky that is covered by a given number of tiles in the seven-pass dark tiling and the four-pass bright tiling. On average, a given part of the sky is covered by 5.2 dark tiles and 3.2 bright tiles.

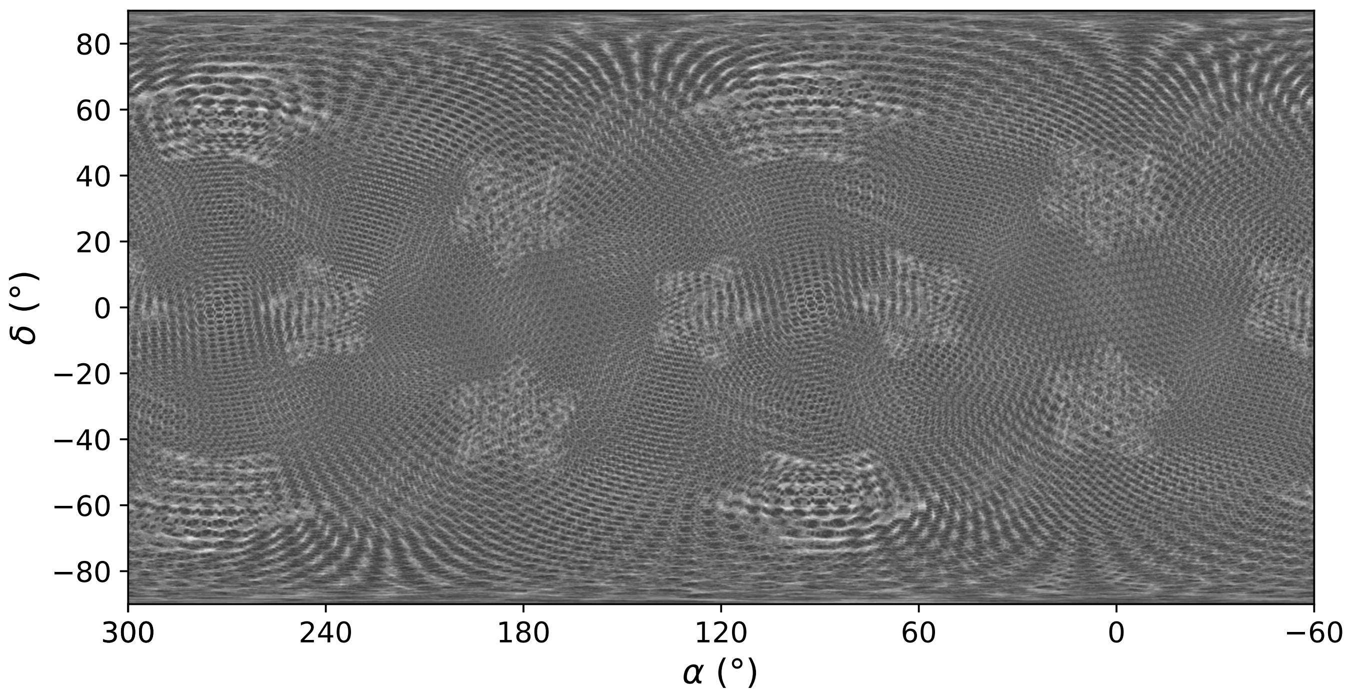

Figure 3. The number of exposures that can reach any particular point of the sky, for the seven-pass dark program, were no areas excluded (e.g., due to low Galactic latitude or low declination). The twelve star-like regions with with slightly lower coverage corresponds to the points of the underlying icosahedral tiling of Hardin et al. (2000).

These constraints produce the footprint shown in Figure 4.

Though we have imposed no explicit cuts on Galactic extinction, we only target regions of the sky with imaging from the DESI Legacy Imaging Survey. That survey explicitly avoided high $E(B-V)$ regions, so these regions are naturally avoided in the DESI footprint without need for further adjustment. Cuts on Galactic latitude do trim the edges of the imaging footprint slightly, however.

The trend in exposure factor with declination in Figure 4 comes from the dependence of survey speed on airmass (§5.3). The SGC is significantly more expensive than the NGC due to a combination of extinction and airmass. No Legacy Survey imaging was available in the SGC north of $\delta = 32°$, though this region would otherwise be favorable for extragalactic studies. The

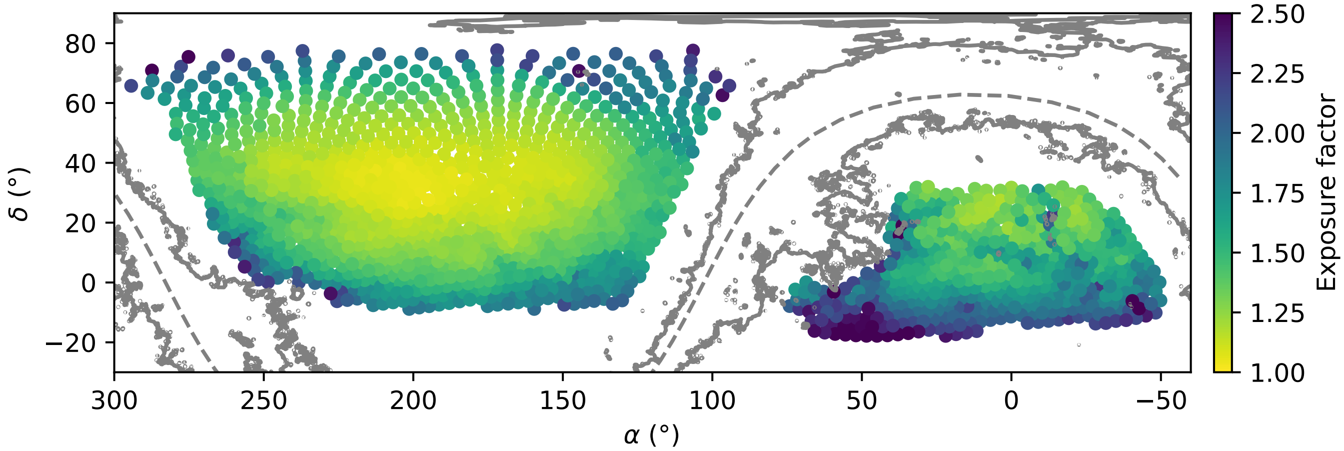

Figure 4. The footprint of the DESI survey resulting from the constraints of §3. Tiles are colored by the amount of time it would take to reach a fixed intrinsic galaxy depth, relative to observing at zenith in the absence of Galactic extinction. This is $f_{\rm dust} f_{\rm airmass}$, from Equations 1 and 2. Airmasses are computed using the design airmasses resulting from the optimization of §4.1. The Galactic plane is shown as a dotted gray line, and the gray contour shows $E(B-V) = 0.3$ mag. Tiles in extinguished regions and at the declination bounds of the survey are most expensive, owing to both atmospheric and Galactic extinction.

irregular small-scale variation comes from Galactic extinction.

The sky area within 1.6° of at least three tiles for the seven pass dark program is 14,246 sq. deg..

All main survey tile coordinates are rounded to the nearest 0.001° to improve legibility.

3.1. *Adjustments to tile centers*

The simple footprint definition of §3 describes our basic footprint selection strategy. Many tile centers are additionally adjusted to avoid bright stars.

The wide field of view (3.2°) of DESI means that bright stars cannot be completely avoided. However, bright stars are particularly damaging if they fall in a few special parts of the DESI focal plane.

First, it is problematic if a very bright star falls on a GFA. These can make it challenging to guide the telescope. Worse, the filter on the GFA reflects light falling outside of the GFA bandpass. Light from the bright star then ends up adding to a large out-of-focus ghost image covering a substantial portion of the DESI focal plane. This is avoided by shifting the tile centers to move bright stars off of the GFA filters. For tiles where a star with Gaia magnitude $G < 6$ lands nears a GFA, we searched for the smallest shift in RA or Dec, in steps of 10 arcseconds, that would put the star at least 25 arcseconds from a GFA.

Second, data from a petal can be rendered useless if a fiber is placed directly on a bright star, saturating large parts of the detector. This is mostly avoided by re-positioning such fibers (which will never have valid main survey targets) away from bright stars. But in rare cases a non-functional fiber happens to land on a very bright object. We adjust tile centers in these cases. After finding bright stars that land near the current set of non-functional positioners for each tile, we search for a small offset (up to 15 arcseconds) of the tile centers in order to minimize the total star light reaching non-functional positioners.

We periodically compute new offsets for tile centers to account for new or bumped non-functional positioners, but we do not do this on the fly when designing each tile.

4. SURVEY STRATEGY

The goal of DESI is to observe a large, homogeneous, ~~efficient,~~ reproducible, and cosmologically interesting set of targets over 14,000 sq. deg. of the sky (DESI Collaboration et al. 2016a). The survey further aims to operate in a "depth-first" fashion where all DESI observations in a particular region are completed before moving on to other parts of the sky.

A critical constraint on the DESI survey strategy is that each DESI observation of a field depends on all earlier, overlapping observations of that field. This is primarily motivated by the need to identify $z > 2.1$ quasars in fields from their initial observations, so that these Ly-α forest tracers can be targeted for repeat observations on subsequent overlapping fields (DESI Collaboration et al. 2016a). A secondary motivation is to obtain observations of targets where initial observations failed due to temporary glitches in fiber positioning or in the spectrographs. This dependence places important constraints on the survey strategy—an observation of a field cannot be made until earlier observations of all overlapping fields have been analyzed. In particular, no two overlapping fields of either the dark program or the bright program may be observed over the course of a single night.

The DESI survey definition (DESI Collaboration et al. 2016a) provides the basic information about each program, including the targets in each program, the amount of effective exposure time (in essence, signal-to-noise; §5.2) required to observe these targets, and the region of the sky where observations are needed. Three programs are defined. First, the dark program, which consists of 9,929 tiles observing luminous red galaxies, emission line galaxies, and quasars from $0.4 < z < 3.5$ (§3). Second, the bright program, which consists of 2,657 tiles observing bright galaxies and Milky Way stars (§3). Third, the backup program, which consists of brighter Milky Way stars. Each of these programs have independent target lists that are separately tracked. The bright and dark programs cover the same region of the high Galactic latitude sky, overlapping spatially; the backup program covers the same area as the bright and dark programs, as well as extending to lower Galactic latitudes.

The dark program is observed whenever conditions are good, and the survey speed for dark tiles is better than 0.4 (§5.3). When conditions are worse, due to bright skies or poor seeing or transparency, DESI observes the bright program, until the survey speed for bright tiles is worse than 0.08. In these poor conditions, DESI observes backup program tiles. This tiered approach is motivated by placing the brightest targets in the worst conditions, so that systematic uncertainties are limited. As an added benefit, this approach reduces overheads by placing the exposures needing the shortest effective exposure times in the worst conditions.

The next broad strategic element of the survey is to observe "depth-first", completing all DESI observations of a particular region of the sky as soon as possible. This allows these regions of the sky to be available early for cosmological investigations, and allows many scientific programs to proceed after the first year (albeit over a limited area). It also minimizes the negative impact of falling behind schedule; we would prefer to end the survey with a complete 13,000 sq. deg. survey than an inhomogeneous 14,000 sq. deg. survey. The depth-first goal is implemented in the nightly field selection (§5.5) by preferring low declination tiles tiles near the celestial equator[1], tiles for which neighboring observations have been made, and tiles which have already been started but for which observations are not yet complete.

[1] A preference for a particular sky region keeps the footprint spatially compact; equatorial fields also enable early science results combining DESI data with other equatorial surveys.

The remaining elements of survey planning focus on how we can observe the DESI footprint as efficiently as possible. This means optimizing the hour angles at which tiles are observed, attempting to observe all tiles as they transit the meridian while reconciling that with the actual distribution of tiles on the sky. It also means limiting the lengths of the slews between adjacent tiles.

4.1. *Airmass Optimization*

Survey planning assigns each tile an optimal hour angle. These optimal hour angles need to satisfy two requirements:

1. The distribution of local sidereal time (LST) needed to observe all the tiles should match the distribution of local sidereal time expected to be available to the survey.
2. The total time needed to finish the survey should be as short as possible.

Alternatively, for the dark program, these requirements could be rephrased as asking how to minimize the airmass of the observations subject to the time available to the survey. For the bright program, the computation of optimal hour angles is more complicated because moon avoidance becomes important.

The airmass optimization algorithm for DESI is simple. Initial hour angle assignments are made by sorting the tiles by right ascension and constructing the cumulative distribution function, weighting each tile by its expected observational cost. We compute the cumulative distribution function for the available local sidereal time in the same way, accounting for seasonal variations in the amount of time lost to weather and for monsoon season shutdowns. Right ascensions are mapped to sidereal times through their cumulative distribution functions so that the cumulative distribution functions match. The result is an assignment of local sidereal times to tiles, which are then used for deciding which tile to observe at any time during a nightAn initial guess of the assignment of hour angles to tiles is made by matching the LST distribution available to the survey to the right-ascension distribution of the survey's tiles, weighted by the tiles' expected observation times. The initial assignments of tiles to LSTs is then further optimized through a simulated annealing process to minimize the total amount of time needed to observe the tiles, while maintaining the match between the distribution of LST available to the survey and the distribution of LST needed to observe the tiles. See appendix §A for more details about the airmass optimization process used in DESI.

~~The~~ Ultimately, the optimization process aims to minimize the expected observation time of the DESI survey. This is simply the sum of the effective times needed for each tile multiplied by corrections for extinction and airmass. The extinction correction is given by

$$f_{\rm dust} = 10^{2\times 2.165\times E(B-V)/2.5} \qquad (1)$$

using reddening $E(B-V)$ from Schlegel et al. (1998) with the calibration of Schlafly & Finkbeiner (2011). This reddening is taken to be the median SFD reddening over the 3.21° diameter tile. Meanwhile the airmass correction is

$$f_{\rm airmass} = X^{1.75}\,, \qquad (2)$$

where X is the airmass of the observation. This airmass adjustment is an empirical adjustment accounting for lower atmospheric throughput, brighter sky background, and worse seeing at higher airmass.

The ~~initial assignment of tiles to right ascensions has only a single free parameter, essentially the local sidereal time that should be observed for tiles at right ascension zero. We try a number of starting points around the circle and select the one with the lowest total expected observation time, derived from summing the exposure times of all of the tiles times the exposure factors of Equations 1 and 2, using the airmasses implied by the assignment of tiles to sidereal times.~~

~~This initial assignment of local sidereal times (LST) to tiles is then improved by a simulated annealing process. At each iteration of the process, a δHA scale is selected. Tiles are perturbed by a random amount on this scale to reduce the total observational cost of the survey (i.e., to improve the airmass distribution) and to improve the match of the planned LST distribution to the available LST distribution. Then at each iteration the LST distribution is spatially smoothed and the δHA scale is reduced by 5%. This annealing process only changes the initial LST assigment slightly, however.~~

~~This process is~~ DESI airmass optimization scheme is close to optimal for situations when the moon is down. For the bright time survey when the moon is usually up, determining the optimal observing strategy is much more challenging. For DESI, this added challenge is ignored and we optimize both the dark and bright programs using the ~~simple airmass optimization described above—the~~ same airmass optimization algorithm—the moon is not included in the optimization process. The bright program efficiency could be improved by a more advanced optimization process.

The airmass optimization process should be performed periodically as the survey proceeds. We aim to do this about once a year, but did not update the design hour angles during the first 1.1 years of the survey.

The backup program is not optimized for airmass; we aim to observe all tiles at zero hour angle. This reflects the fact that completeness ~~&~~ and homogeneity are not as important to the backup program as they are to the cosmological programs.

4.2. *Slew Optimization*

Long slews reduce the amount of time each night during which DESI can be making science observations. A number of operations occur when ending one observation and starting a new one (DESI Collaboration et al. 2022):

1. Spectrograph readout
2. "Blind" positioner move
3. Slewing & settling
4. Field acquisition & guiding
5. "Correction" positioner move

The spectrograph readout and blind positioner move can occur simultaneously with slewing and settling, but the field acquisition and correction move must occur after slewing is complete. If the slew and settle time exceeds ten seconds, slews begin to increase the overhead between exposures. Settling time is 8 seconds, and it takes 16 seconds to slew between adjacent DESI fields. So slewing adds to DESI overheads regardless of slew length.

Nevertheless, even without any explicit slew optimization, slewing would only account for 3.1% of the open shutter time for the DESI survey, according to survey simulations. To try to reduce this, we do a simple greedy slew optimization where tiles nearby the current location of the telescope are preferentially observed. We penalize slews in the declination or negative right ascension directions, but not in the positive right ascension direction, since we do not want to penalize slews that are trying to keep up with the sky rotation. This simple prescription reduces the slew time to 2.9% in simulations, and inspection of the resulting slew patterns suggests limited potential for further improvement.

5. SURVEY OPERATIONS

Survey operations broadly refers to the process by which we complete the tiles composing the DESI dark, bright, and backup programs. Because ~~we use past observations to inform future observations of overlapping observations~~past exposures inform future exposures, we cannot observe tiles overlapping previously observed "pending" tiles ~~(in the same program)~~ until the analysis of those tiles has completed~~and, in the~~

~~dark program, $z > 2.1$ quasars have been identified~~[2]. So the basic operational scheme becomes:

1. Each night, observe tiles that do not overlap the footprint of pending tiles.

2. Each day, analyze observations and incorporate results into the targeting ledger (merged target list or MTL; see §6), clearing pending tiles.

If data reductions are delayed, we may skip step (2), in which case the footprint of pending tiles grows. We repeat this process until the survey is complete. The rest of this section details our implementation of this scheme.

The ability to reproduce the particular set of targets that DESI ultimately observes is a key requirement of this process. We need to be able to simulate the observational process on mock target catalogs in order to account for the effect of the DESI design on the final galaxy redshift catalogs. Accordingly, we must be capable of reproducing the assignment of every fiber to every target over the course of the survey. Since these choices depend on the current observational state of the targets and the current health of the instrument, we need to track these quantities through time (see §6, §5.13). We record the state of both the targets and the instrument in ledgers. In these ledgers, each row is time-stamped and changes are made by appending new rows to the ledger indicating the new state of a target or fiber. Thus, past decisions about the assignment of fibers to targets can be reproduced by reading the ledgers through to the time at which those decisions were made.

5.1. *Daily Observation Overview*

The broad operational model of DESI is specifically implemented in operations in a number of different steps, schematically illustrated in Figure 5. These steps are described in more detail later in this section, and include:

1. Afternoon planning identifies completed, pending, and unobserved tiles, and establishes priorities for the night's observations.

2. The Next Field Selector selects each program and tile to observe during the night.

3. Targets are assigned to each positioner on the fly immediately before the observation is made.

4. DESI positions fibers and the spectrograph shutter opens to observe the targets in the field (Silber et al. 2023).

5. The Exposure Time Calculator (ETC, Kirkby et al. 2023) computes the effective time obtained on each tile during an observation, determining when an observation is complete.

6. The spectroscopic pipeline reduces, classifies, and measures redshifts for all targets the following morning (Guy et al. 2023; Bailey et al. 2023).

7. The reproducibility of the on-the-fly tile design is confirmed by designing the tile a second time outside of operations on the mountain.

8. Humans perform quality assurance, visually inspecting summary figures and statistics on each tile, and declare tiles either finished or problematic.

9. Reduced data products for tiles passing quality assurance are archived.

10. The Merged Target List (MTL) is updated with the new data, updating the observation state and redshift of the observed targets.

11. The state of the robotic positioners is updated, should any have failed.

12. The results of the previous nights' observations are available for afternoon planning, and the process repeats for the next night's observations.

Some of these steps need not occur every day. Pipeline reductions, quality assurance, MTL updates, and focal plane state updates can all be delayed, as illustrated by the dashed box in Figure 5. When MTL updates are delayed, tiles will be left in a "pending" state and the survey will be forced to observe new parts of the sky rather than completing the survey in already observed regions. Delaying focal plane state updates causes only a slightly inefficient assignment of positioners to targets. In practice, we perform MTL updates roughly weekly in bright time when progress is slow and roughly every other day in good weather in dark time.

The flow chart in Figure 5 is only intended to be schematic and ignores many details. For example, the exposure time calculator runs online and is simultaneous with the exposure. Some visits are split into multiple exposures and do not require new fiber assignment or full positioning & acquisition loops. The spectroscopic extractions and redshift determination begin during the night as the data are taken, and so do not strictly follow

[2] Note that the different programs are independent, so a pending bright tile does not block observation of an overlapping dark tile.

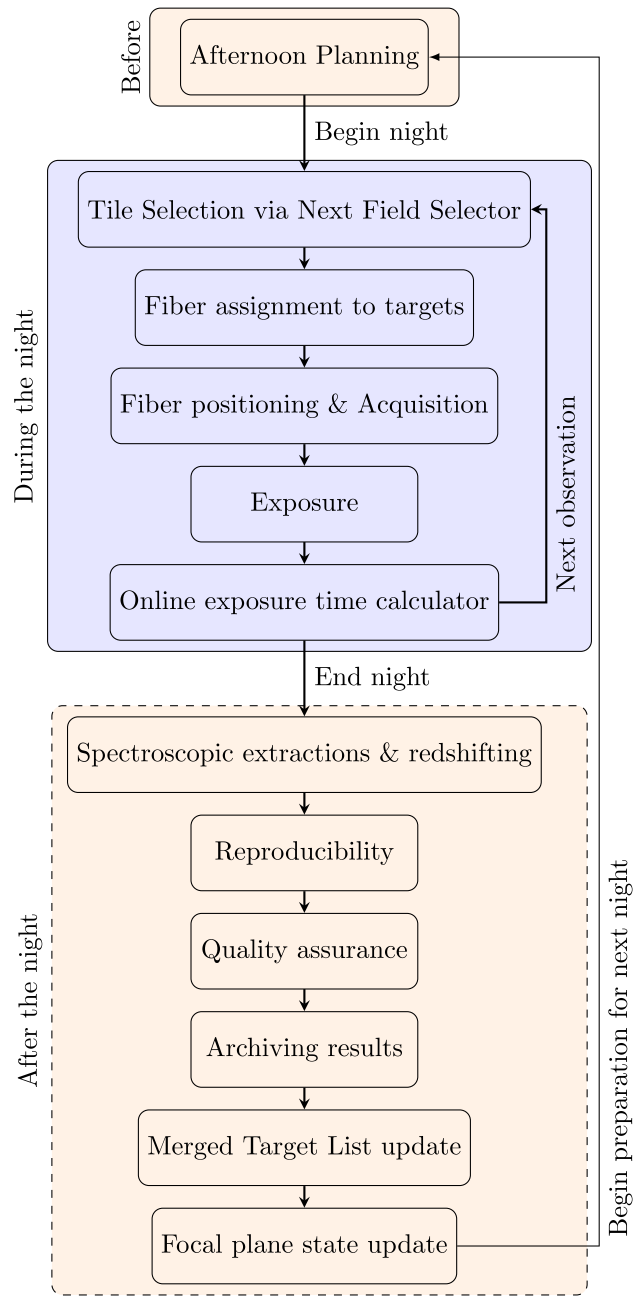

Figure 5. Schematic flow chart of DESI operations steps, running from planning for the night, through each night's observations, through their reduction and updates to the MTL. Steps in the dashed box are optional and may be skipped temporarily if systems are not available. See §5.1 for details.

the separation implied by the flow chart. Still, Figure 5 gives a good schematic overview of the DESI daily operation procedure.

5.2. *Effective Time*

The concept of "effective time" is important to DESI operations. We describe effective time briefly here; see Guy et al. (2023) for more details. Ultimately DESI seeks to measure the fluxes from distant galaxies to a specified accuracy. Rather than phrasing this accuracy in terms of the flux uncertainty at a particular wavelength, we parameterize it in terms of the amount of time it would take to reach a goal uncertainty in "nominal" conditions, defined to be 1.1″ seeing, a sky background of 21.07 mag per square arcsecond in the r band, photometric conditions, observations at zenith, through zero Galactic dust reddening. This "goal uncertainty" is weighted over wavelengths and spectral features in order to make it a good proxy for DESI's ability to find a redshift for a galaxy spectrum. Observations in the dark program aim for 1000 s of effective time, while bright program observations aim for 180 s.

The concept of effective time is made more complicated by the following effects:

1. Poisson noise from source flux,
2. different intrinsic source sizes (e.g., stars versus large galaxies), and
3. chromatic variation in the sky background and throughput.

The Poisson noise from source flux and the different intrinsic source sizes are challenging because they vary from source to source, making it hard to define the effective time for a tile. We adopt fiducial source fluxes and sizes for computing effective times for main survey tiles, which are given in Table 1.

Chromatic variation in the system throughput, sky brightness, and detector performance also complicates the notion of effective time. The goal is to have all tiles reach a nominal depth. However, for example, when comparing tiles observed through a red, moonless sky with tiles observed through a blue, moony sky, tiles with equal depth in the r band will have different depths in the g and z bands. A ~~specific~~ simple prescription for this nominal depth would be an average signal-to-noise ratio in a particular range of wavelengths for targets of a given magnitude. DESI instead adopts a detailed set of weights ~~derived from the spectra in Table 1 is used to average over this complication; see~~ over all wavelengths that is different for each program, reflecting the spectral lines in the different target classes and their redshift distribution. See Guy et al. (2023) for more details. These more detailed weights are intended to deliver something closer to a uniform redshift success rate for the different key target classes.

Finally, effective time accounts for Galactic extinction. The ETC (Kirkby et al. 2023, §5.7) aims to reach a fixed precision in the intrinsic r band flux of target galaxies. Accordingly, the real time needed to reach a given effec-

Table 1. Source properties used for effective tile effective time

program	profile	spectrum	source counts
dark	exponential, $r_{\rm half} = 0.45''$	LRG spectrum averaged over $0.68 < z < 0.97$	0.00 nMgy
bright	de Vaucouleurs, $r_{\rm half} = 1.5''$	BGS spectrum averaged over $0.13 < z < 0.37$	1.71 nMgy

See https://www.sdss4.org/dr17/help/glossary/#nanomaggie for the definition of nMgy.

tive time is increased by Equation 1 in the presence of Galactic extinction.

5.3. *Survey Speed*

The concept of survey speed is related to effective time, and is used for a variety of purposes, including the selection of program to observe during the night. The survey speed is computed using the current seeing, sky background, transparency, and airmass from the Exposure Time Calculator (ETC) (§5.7). The survey speed measures how many effective seconds DESI would be accumulating per second, were DESI observing a tile at zenith and zero dust extinction in the current conditions. Survey speeds range from zero in clouded-out conditions to ~ 2.5 in the best conditions, as shown in Figure 6. Dark tiles are never observed outside of 15° twilight or when survey speed measurements are unavailable, leading to a small number of bright observations in rather good conditions.

The relation between survey speed and seeing depends on the program, since programs observing point sources are more sensitive to seeing than programs observing large galaxies.

The survey speed is adjusted to airmass 1 when observations are made away from zenith following Equation 2. This adjustment is intended to account not only for atmospheric extinction, but also for worsened seeing and sky background at lower elevations. The ETC assesses the survey speed in real time; see §5.7 for more details.

5.4. *Afternoon Planning*

The role of afternoon planning is to determine the current status of survey progress in order to ~~determine~~ set the base priorities of tiles for the coming night's observations. Afternoon planning compiles a list of all observed exposures and their associated effective times (§5.2, §5.7), and combines these to determine the status of each tile: unobserved, pending, or completed. This status is used to determine the priority of each tile (§5.4.1), which determines which tiles are observed in the course of the night. Files describing the configuration of the survey strategy for each night and the state of the survey progress are created. The Next Field

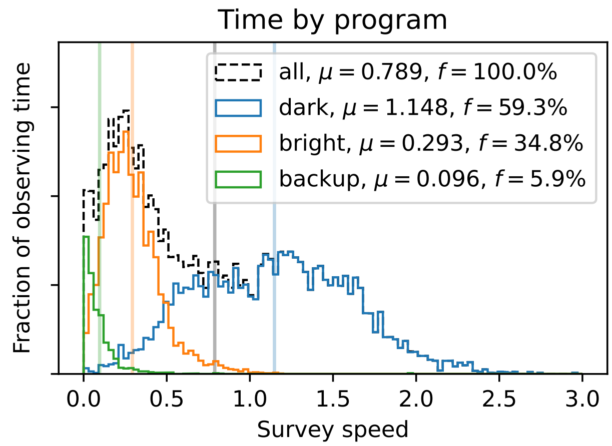

Figure 6. The survey speed delivered by the DESI main survey in different programs, as measured by the ETC. The survey speed describes the rate at which $(S/N)^2$ is accumulated relative to nominal dark conditions, and is highest when the seeing is good and the sky is clear and dark. The dark program is observed in the best conditions, while the bright and backup programs are observed in progressively worse conditions. ~~Dark tiles are never observed outside of 15° twilight or when survey speed measurements are unavailable, leading to a small number of bright observations in rather good conditions.~~ The legend gives the mean speed μ and the fraction of survey time spent in each program f.

Selector (§5.5) then uses these files in the course of the night's observing.

There are multiple sources for the effective time of each tile. The authoritative source of this information is the offline pipeline. Offline pipeline effective times become available in the morning after each night's observations, provided that no issues with the processing or computer systems prevent their computation. Absent information from the offline pipeline, afternoon planning uses effective times from the ETC (§5.7), which are computed on the mountain during each exposure and are always available.

5.4.1. *Tile Priorities*

A number of factors contribute to the priority assigned to a tile, which the Next Field Selector uses to select ~~a tile~~ tiles for observation (see §5.5). Note that these tile priorities are unrelated to the target priorities discussed in §6, which determine which targets get observed within a given tile. Afternoon planning sets a ~~fixed base priority~~ priority P of each tile for each night ~~. This priority is based on the status of the tile, its celestial coordinates, and~~ according to the following equations:

$$P = dsnB \tag{3}$$

$$d = \exp(-|\delta|/160^\circ) \tag{4}$$

$$s = 1 + 0.1 \times \texttt{is\_started} \tag{5}$$

$$n = 1 + 0.08 \times f_{\mathrm{neighbor}} \,. \tag{6}$$

Here δ is the ~~number of overlapping tiles~~ declination of a tile, `is_started` is one if a tile has been started and zero otherwise, and f_{neighbor} is the fraction of tiles overlapping this one that have been ~~observed. The Next Field Selector (§5.5) combines the base priority with additional factors depending on the current sidereal time and pointing of the telescope to select~~ finished. The factor B is a rarely used boost factor that can be set to manually change the priority of a tile. ~~The base priorities are close to one, and so primarily serve to "break the tie" among the many tiles at the appropriate hour angle for observation at a given time. The two most important contributions are the "finish-if-already-started" priority, which leads tiles to get another observation if they were started but not completed, and the low $|\delta|$ priority, which prefers equatorial tiles .~~

~~Tiles which~~

The broad goal of these priorities is to start the survey on the celestial equator and build out (d); to finish tiles that have already been started ~~are assigned the highest base priority, so that they can be finished. Unobserved tiles receive the default priority. Finished tiles are assigned a priority of zero and are not reobserved.~~

~~Spatially, tiles are assigned priorities according to their declination. Equatorial tiles are given the highest priority. This~~ (s); and to finish tiles where we already have a number of observations (n). The preference for equatorial tiles keeps the footprint spatially compact and leads to depth-first observations. ~~Beginning the survey~~ Starting on the equator also enables early science ~~involving~~ using cross-correlations with other equatorial surveys. Finally, it permits follow-up observations of interesting targets from telescopes in both hemispheres.

~~Tiles which overlap many other completed tiles are boosted in priority. This leads the survey to complete regions of the sky before moving to new areas.~~

5.5. *Next Field Selector*

The Next Field Selector (NFS) is responsible for selecting tiles to observe during each night. Roughly two minutes before each observation is expected to complete, the DESI Instrument Control System (ICS) requests a tile from the NFS. The NFS selects a program and computes ~~scores~~ a "score" for each tile in that program~~, selecting~~. It then chooses the tile with the highest score ~~, and begins designing this tile~~ and designs it on the fly (§5.6). The resulting tile is made available to the ICS and is observed.

Program selection is primarily driven by survey speed. When the survey speed is good, averaging > 0.4 for the past 20 minutes, dark program tiles are selected. When the survey speed is poor, $0.08 < \text{speed} < 0.4$, bright program tiles are selected. Otherwise, backup tiles are selected. In addition to this selection, dark tiles are never selected when the sun is within 15° of the horizon, and bright tiles are not selected when the sun is within 12° of the horizon.

The tile scores S used by the NFS are computed as the product of the base tile priority ~~(§5.4.1), a squared exponential penalty in the difference between the tile design hour angle and the current~~ P from afternoon planning (Equation 3), and two additional factors.

$$S = Pe^{-T_{\mathrm{slew}}/400\ \mathrm{s}} e^{-(H-H_D)^2/2\sigma^2} \tag{7}$$

$$\sigma = (d^2X/dH^2)^{-1/2}/4 \tag{8}$$

where T_{slew} is the estimated time needed to slew to the new tile from the current tile, H is the expected hour angle of the midpoint of the next observation, and H_D is the design hour angle of the tile~~, and a factor preferring short slews in the positive RA direction to other slews~~.

~~The squared exponential factor in hour angle has a variance that depends on the second derivative of the airmass at an hour angle of zero, so that tiles at low declination where the airmass changes rapidly with hour angle are observed close to their design hour angles, while tiles near the celestial pole are more flexible in their observation. Tiles above $\delta = 12^\circ$ are given a penalty factor with $\sigma = 15^\circ$, while below $\delta = 12^\circ$, σ decreases until reaching 10° at $\delta = -20^\circ$~~ $X(H)$ is the airmass of a tile as a function of its hour angle.

~~The slew time factor in the tile score is exponential in the amount of slew time, not counting any slew~~ The first factor prefers tiles near the current location of the telescope in order to reduce time spent slewing. The variable T_{slew} is based on the location of the new tile, the current location of the telescope, and the acceleration and cruise speed of the telescope on its hour angle and declination axes. For the computation of $T_{\rm slew}$ in the NFS, we do *not* count time spent slewing in ~~the positive right ascensiondirection~~direction of increasing right ascension. Slews in declination ~~may also be "free" if they are covered in the time needed to slew in the positive right ascension direction. Slewing in the positive RA direction is not penalized by the NFS in order to allow the NFS to~~ that occur while slewing toward increasing right ascension likewise do not contribute to $T_{\rm slew}$. This is to avoid penalizing the telescope for slewing to keep up with the sky. We do not, for example, ~~jump from~~ want the telescope to dawdle in one Galactic cap to ~~the other, without dawdling in the wrong cap until the hour angle penalty becomes overwhelming.~~ avoid slewing to the other to keep up with the sky.

The second factor penalizes tiles observed away from their design hour angles. When observing tiles away from their design hour angles, we prefer to observe high declination tiles to low declination tiles, because the airmass of a low declination tile varies more quickly with hour angle than a high declination tile. We implement this preference by letting σ depend on the second derivative of the airmass with hour angle, evaluated at hour angle zero. We clip σ to between 7.5° and 15° to avoid tiles with too-large or too-small observability windows. This ultimately leads to σ taking the value of 15° above $\delta = 12°$, and $\sigma \approx 10°$ at the southern boundary of the main footprint.

The NFS also places some constraints that may prevent a tile from being observed. For example, no tile may be observed within 50° of the moon, though this limit is occasionally relaxed when the location of the moon in the survey footprint would mean that no tiles were otherwise available. Similarly, no tiles may be observed within 2° of ~~any of the planets interior to Saturn's orbit~~a classical planet (one of the first six planets). Most importantly, no tile may be observed that overlaps a pending tile, as discussed in §5.1. Observers may impose additional constraints based on current conditions. These constraints are most often used to force observation in the north when strong southerly winds would otherwise shake the telescope and degrade the delivered image quality, though they can also be used to chase holes in the clouds.

5.6. *On-the-fly Fiber Assignment*

Tiles are designed on the fly when requested by the NFS. This means that we do not know which fibers will be assigned to which targets until minutes before observations begin. When requested, the `fiberassign` package (Raichoor et al. 2023) uses the MTL (§5.12, §6) and focal plane state (§5.13) to determine how best to allocate fibers to targets. Secondary targets and targets of opportunity are also optionally included.

Tile design takes roughly thirty seconds. Two minutes are allocated to cover rare cases in dense fields and when latency on the DESI computers is higher than typical.

Because reproducible assignments are critical to the large scale structure analysis of the final redshift catalog, `fiberassign` inputs are all in the form of ledgers recording the state of the system and targets at any given time. Moreover, the complete state of the software and input data to `fiberassign` is logged at run time. We also recreate each tile designed on-the-fly at the mountain at the National Energy Research Scientific Computing Center (NERSC) on the following day to verify that the same assignments are made (§5.9).

On-the-fly assignment is convenient because it allows decisions about which tile should be observed to be made in response to current observing conditions, while also allowing every tile to depend on all of its observed neighbors. A disadvantage of on-the-fly assignment is that it limits the optimization possibilities of `fiberassign`. In this mode, `fiberassign` does not know about future observations and cannot adjust its assignment of fibers to targets using that information.

Fiber assignment also needs access to the current state of the DESI focal plane. A substantial number of DESI positioners ($\sim$ 700) cannot be assigned to science targets and are usually left fixed in place. Small numbers of additional positioners occasionally become non-functional. In order to optimally assign targets to positioners, `fiberassign` must avoid assigning functional positioners to locations that would collide with non-functional positioners. Additionally, we assess whether each non-functional positioner lands on a location which can be used to measure the sky spectrum. If so, we reduce the number of functional positioners allocated to determining sky. This has a beneficial impact on survey efficiency, since the number of fibers allocated to sky is nearly 10% of the total fiber budget, and is similar to the number of non-functional fibers.

5.7. *Exposure Time Calculator*

The ETC (Kirkby et al. 2023) is responsible for deciding how long to observe each tile, and how much effective time (§5.2) each tile has accumulated during the night. It is also responsible for tracking survey speed and deciding when to split long observation sequences into multiple exposures[3].

[3] ~~Splitting an exposure reduces the impact of cosmic rays and enables fibers to be repositioned to account for changing airmass.~~

The ETC uses measurements of the sky background, seeing, and transparency to perform these tasks. Sky measurements come from the DESI sky camera, which uses 20 dedicated sky fibers to measure the sky brightness in the r band (two sky fibers on each petal) (DESI Collaboration et al. 2022). Seeing and transparency measurements come from the DESI ~~Guide Focus Array cameras (GFAs)~~GFAs, which are also used for guiding and focusing the telescope (DESI Collaboration et al. 2022). Measurements of the amount of flux entering a fiber relative to nominal—the combination of seeing, throughput, and fiber mis-centering most relevant to the effective time—are computed from GFA frames every eight seconds.

These measurements of the terms contributing to the signal and noise accumulated in the spectrograph are then used to estimate the $(S/N)^2$ obtained in the exposure in real time, which is calibrated to effective time by a single scale factor in each program. The ETC makes very good predictions for the completeness of dark tiles, leading DESI to have final tile spectroscopic effective times that very closely match their desired goal effective times, as shown in Figure 7. Bright program tiles show worse agreement due primarily to the varying ~~sky color~~ color of the sky background depending on the phase and location of the moon. The ETC has access only to the r band sky brightness, while the spectroscopic effective times use the observed brightness of the sky at all wavelengths (§5.2). Bright program tiles taken in conditions of ~~very~~ bright moon tend to be overexposed.

We cap the length of any single exposure to 1800 s for two reasons. First, long exposures suffer more cosmic ray hits, which wipe out all signal in affected pixels. By splitting long observations into multiple exposures, a cosmic ray wipes out only the signal in the exposure in which it occurs. Second, the airmass of a field changes slowly over the course of an exposure. Splitting long exposures allows us to adjust the atmospheric dispersion corrector for the new location of the field relative to zenith and to reposition the positioners accordingly. If the ETC determines that an observation is likely to exceed 1800 s, it aims to split it into a series of exposures of equal length. We cap the amount of time spent on a single tile per night to 90 minutes; if an observation does not reach depth in this time we return to it on a later night.

The required inputs for the ETC are the requested effective time for a tile, the program, and the ~~Galactic extinction averaged over each tile footprint~~median Galactic extinction over all targets on each tile. The requested remaining effective time is provided by the NFS, while the program and extinction are available in the tile files created by `fiberassign`.

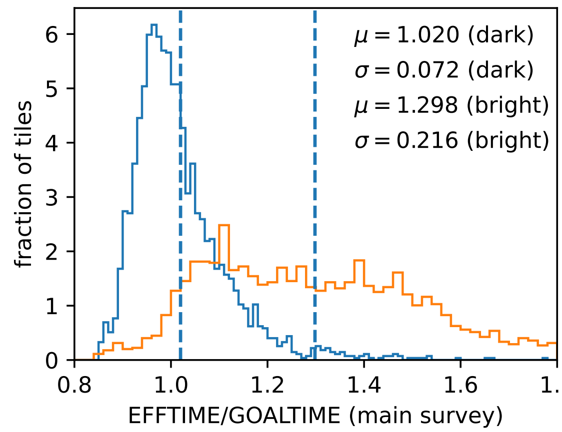

Figure 7. Completed dark time tiles have a narrow distribution in EFFTIME around the goal time of 1000 s, with tiles having on average 102% of their goal effective time, with a standard deviation of 7% (blue histogram). This demonstrates that the ETC is able to accurately predict the spectroscopic effective times from the real-time transparency, seeing, and sky brightness measurements. Bright time tiles have a much broader range of effective time fractions ~~, owing to the range of sky colors in which they are observed,~~ and tend to be observed 30% longer than necessary (orange histogram).

5.8. *Spectroscopic Pipeline*

The DESI spectroscopic pipeline ~~(Guy et al. 2023; Bailey et al. 20~~ runs each morning following observations, aiming to complete processing by 10:00 AM Pacific time. ~~This includes~~The pipeline carries out a large number of tasks, detailed in Guy et al. (2023) and Bailey et al. (2023). These include:

1. processing nightly calibration images (zero second, arc lamp, and flat field exposures),
2. finding wavelength and ~~point-spread-function~~ two-dimensional line-spread-function solutions for each exposure,
3. extracting the one-dimensional spectra from the ~~detrended~~ two-dimensional frames after correction for calibration images,
4. subtracting sky background light,
5. calibrating spectra to physical units (10^{-17} erg/s/cm^2/Å),
6. determining redshifts and classifications for each spectrum, and

7. evaluating the status of each tile and spectrum.

These tasks are all routinely completed within a few hours of the end of the night, for more than 10^5 fibers on a typical night.

The redshifts are used to update the MTL (5.12), promoting newly detected $z > 2.1$ Ly-α quasars to become the highest priority targets on future, overlapping tiles in the dark program. Other targets are marked with their new redshifts and with flags indicating whether the spectrum is valid or if for some reason the observation should be ignored (e.g., because the positioner did not reach its target location).

5.9. *Fiber Assignment Reproducibility*

Galaxy clustering measurements and cosmological analyses of the DESI redshifts depend on being able to reproduce the algorithm by which fibers were assigned to targets. The on-the-fly assignment of fibers to targets during the night raises concerns that a configuration problem may lead to different assignments when `fiberassign` (§5.6) is run on the mountain from when it is run at NERSC.

We reproduce every tile designed over the course of each night at NERSC the following morning to ensure that this does not occur.

5.10. *Quality Assurance*

The DESI survey uses the information on each tile to inform later observations of overlapping tiles, via incorporation into the MTL. The spectroscopic pipeline (Guy et al. 2023; Bailey et al. 2023) identifies Ly-α quasars in each observation, so that later tiles can be tasked with reobserving those high-priority targets. It also identifies which spectra are good, and which spectra are affected by issues with the hardware and should be ignored.

Accordingly, it is important to assess the quality of each observation so that problems with the data are identified before they are incorporated into the MTL. We make a number of quality assurance (QA) plots for each tile when pipeline reductions of that tile are completed. These plots include the redshift distribution of the objects on each tile, the redshifts as a function of fiber number, the effective time as a function of location in the focal plane, and the fiber positioning errors as a function of location in the focal plane[3]. The QA also indicates whether the pipeline identified any problems with the tile, like missing standard stars, large reduced-χ^2 values in the sky fibers after sky subtraction, or poor line-spread-function fits.

A member of the operations team reviews the QA for each tile looking for peculiarities. Most tiles are quickly marked good ($\sim$ 30 s per tile). The remaining more complicated and potentially problematic tiles are marked "unsure" and flagged for follow-up investigation. Examples of such rare cases include tiles with extremely bright stars leading to contamination and sky determination difficulties; cases where small amounts of air leak into the spectrograph, leading to increased glow from the ion pump inside the cryostat and associated enhanced backgrounds; cases where large turbulence in the volume of air between the primary and focal plane causes most positioners to be off target by more than 30 microns RMS; and cases where imperfect sky subtraction in very bright conditions lead to poor redshifts. Typically exposures affected by these kinds of problems are marked bad and reobserved.

Tiles passing QA are now ready for archiving before inclusion in the MTL (§5.12).

5.11. *Tile Archiving*

The daily offline spectroscopic reductions (§5.8) occasionally identify issues in the data or pipeline that need to be addressed before data can be incorporated into the MTL. In these cases, initial reductions are often deleted and replaced with improved reductions. For data that eventually enters the MTL, we want to more strictly archive the reductions that were the source of the MTL updates and therefore affect future observations. Accordingly, once redshift catalogs have been deemed acceptable for incorporation into the MTL, they are copied to a special "archive" directory and made read-only. Updates to the MTL are made only from archived tiles.

5.12. *Merged Target List*

The Merged Target List (MTL) records the current state of each potential DESI target. Before the survey began, it included entries for each potential target drawn from the imaging surveys, together with the class of that target and its priority. Following each tile's successful observation and quality assurance check, the archived results of the tile's spectroscopic analysis are used to update the MTL, adjusting the priorities of observed targets.

[3] Following fiber positioning, the fiber view camera images the focal plane with the fibers back-lit to identify the final location of the fibers. The fiber positioning errors shown in QA are the difference between the intended locations and the locations derived from this image. This is an imperfect proxy; it does not include any systematic errors in the map between true location and location in the fiber view image, and does include turbulent errors due to dome seeing in the fiber view camera image. However, it at least highlights any dramatic errors in fiber positioning.

The most important element of the MTL update is to mark successfully observed objects, so that they may be excluded from future tiles. The next most important element is to mark newly detected Ly-α quasars as high priority targets which should be observed whenever possible.

These updates are performed by adding new rows to the MTL corresponding to each observed target. All entries include a timestamp indicating when they were entered into the MTL. This ledger system enables `fiberassign` to be run in a reproducible fashion by specifying the latest timestamp in the ledger when `fiberassign` was run. Future `fiberassign` runs can read the ledger through that same timestamp in order to see the same survey state that the original assignment used. See §6 for much more detail about the MTL.

5.13. *Focal Plane State Update*

The DESI focal plane state describes which positioners are functional, which positioners are not functional, and which regions of the focal plane must be avoided to prevent collisions with non-functional fibers. The state of the focal plane changes occasionally as positioners malfunction or as positioners are brought back to life. Malfunctioning positioners are also occasionally moved; this changes the areas of the focal plane which must be avoided. The operations database at Kitt Peak is the authoritative source of information on the health of each positioner; information from this database must be synced into the state file used by `fiberassign` in order for fiber assignment to make use of this information.

Like the MTL, the current state of the DESI focal plane is stored in a ledger with timestamps included in every entry. The state of each positioner at a given point in the history of the instrument can then be obtained by reading the ledger through to that specific time. We update this ledger via synchronization with the operations database once each day.

Note that the ledger tracking the focal plane state that is used by `fiberassign` sees only a coarse, daily picture of the state of the positioners. The online system tracks every move of every positioner and its current state. When, for example, a positioner fails during a night, `fiberassign` and the ledger do not see it until the following night. This means that fiberassign will try to assign targets to non-functional fibers during the night following the failure of a positioner. The online system then rejects these assignments. Since at present only roughly one positioner fails per week, there is not much benefit to tracking the focal plane state with better granularity.

Following the focal plane state update, the daily operations loop is ready to repeat. The MTL and focal plane state have been updated, and afternoon planning (§5.4) can prepare for the coming night's observations using the results of the previous night's observations.

6. OVERVIEW OF THE MERGED TARGET LIST

The Merged Target List (MTL) tracks the observational state of all targets which the DESI survey may observe. These targets are drawn from a variety of different programs and classes, which may significantly overlap one another, and are denoted by a unique `TARGETID`, as described in Myers et al. (2023). Distinct target classes often need to be treated differently during DESI operations — for instance $z > 2.1$ quasars ideally need to be observed on 4 overlapping tiles to improve signal-to-noise in the Ly-α forest, whereas emission line galaxies require only a single observation. The main purpose of the DESI MTL code is to enforce a set of decisions for targets that span multiple target classes and so may have competing observational requirements (i.e. effectively "merging" those targets). In this section, we discuss the form of the various MTL ledgers and the logic used to update them during survey operations.

6.1. *The Initial MTL Ledgers*

The MTL software operates on a set of ledgers that contain the minimal information expected to be needed to conduct operational decisions. These ledgers begin with a list of possible targets, which are updated as the survey progresses. Each ledger entry represents a target in a given state at a given time. Additional entries are added to the end of the ledger when a target's state changes. Crucially, under normal operational procedures, no entries are ever *removed* or *changed*. This means that the entire observational history of a target can be recovered by reading a target's ledger entries in order, starting from the initial record.

There are five initial sets of MTL ledgers for the DESI Main Survey: primary dark-time and bright-time ledgers; secondary dark-time and bright-time ledgers; and a set of ledgers for the backup program. Details about how targets are selected for these different programs are available in Myers et al. (2023). Structurally, each of these sets of ledgers populates a separate directory and is organized as a set of files split by HEALPixel (Górski et al. 2005) in the nested scheme at `nside = 32`. This means that each individual ledger covers $\sim 3.36\,\mathrm{deg}^2$ of the DESI footprint described in §3. Guidelines for creating initial MTL ledgers are included as part of a tutorial on processing DESI target files that is available on the desitarget GitHub site[4]. Details about the data model for, and content of, the MTL ledgers is available as part of the DESI data model[5].

6.2. *The Initial Observational State*

Each distinct DESI target class has an associated priority and requisite number of observations, which are inherited from the desitarget bitmask "yaml" files described in Myers et al. (2023)[6]. These initial priorities and numbers of observations are stored in the MTL ledgers as PRIORITY_INIT and NUMOBS_INIT. For example, low-priority emission line galaxies (ELG_LOP targets in Table 2) have PRIORITY_INIT=3100 and NUMOBS_INIT=2[7].

A source may be flagged as belonging to multiple target classes. The PRIORITY_INIT and NUMOBS_INIT values are set separately for dark-time and bright-time MTL ledgers, using only target classes belonging to the appropriate program. For example, a source could be targeted as a quasar *and* a low-priority emission line galaxy *and* a white dwarf. When constructing the dark-time ledgers, only the quasar and emission line galaxy priorities will be considered; the quasar will "win" because PRIORITY_INIT = 3400 (for unobserved quasars) exceeds PRIORITY_INIT = 3100 (for unobserved low-priority ELGs). When constructing the bright-time ledgers, only the bright-time white dwarf targeting bit will be considered, because the quasar and emission line galaxy target classes belong to the dark-time program; the white dwarf values will drive the PRIORITY_INIT and NUMOBS_INIT settings in the bright-time ledgers. An important principle, here, is that the analysis of the bright-time and dark-time programs are independent.

Table 2. Initial priorities for some DESI target classes

Target name	Priority	Notes
Dark-time targets		
QSO	3400	Quasars
LRG	3200	Luminous red galaxies
ELG_HIP	3200	ELGs at highest priority
ELG_LOP	3100	ELGs at low priority
ELG_VLO	3000	ELGs at lowest priority
Bright-time targets		
MWS_WD	2988	White dwarfs
BGS_BRIGHT	2100	Bright-time galaxies
MWS_BROAD	1400	General stars
Rare secondary		
STRONG_LENS	4000	Gravitational lenses
"Filler" secondary		
PSF_OUT_DARK	90	Outlier point sources
Backup targets		
BACKUP_GIANT	35	Halo Giants
BACKUP_FAINT	20	General stars

Only a representative subset of target classes is displayed to illustrate the general prioritization schema.

6.2.1. *Relative Initial Target Priorities*

The relative initial priorities for targets[8] are broadly set by a simple underlying philosophy. Lower-density targets are more likely to be swamped by higher-density targets — so the rarest targets are typically assigned the highest priorities. For example, among dark-time targets, quasars have the highest initial priority, followed by luminous red galaxies and then emission line galaxies. Table 2 lists initial priorities for some representative target classes to help illustrate the general schema.

Bright-time targets are always assigned a lower initial priority than dark-time targets. Bright-time galaxies are prioritized over Milky Way targets, regardless of relative density. This ensures that the distribution of Galactic stars is not imprinted on patterns of large-scale structure traced by the bright galaxy program. The sole exception to this scheme is white dwarf targets, which are relatively rare and valuable. Potential white dwarfs are assigned a higher initial priority than all other bright-time targets (but still have a lower initial priority than dark-time targets).

Secondary targets have a range of initial priorities, driven by the intersecting needs of each specific campaign. Secondary targets are generally not allowed to have higher initial priorities than the DESI primary target classes, except for exceedingly rare, high-value targets. Broadly, secondary targets are prioritized by density with very large "filler" samples having very low initial priorities. The only targets that have an initial pri-

[4] https://github.com/desihub/desitarget/blob/master/doc/nb/how-to-run-target-selection-main-survey.ipynb

[5] See https://desidatamodel.readthedocs.io/en/latest/DESI_SURVEYOPS/mtl/index.html.

[6] See, e.g., https://github.com/desihub/desitarget/blob/1.1.1/py/desitarget/data/targetmask.yaml for the DESI Main Survey.

[7] As a hedge against potentially needing additional signal-to-noise, NUMOBS_INIT for DESI primary galaxy targets was set to 2 total observations. But, in the DESI Main Survey, the second observation is scheduled at very low priority (see, also §6.3.2).

[8] Listed in full in the "yaml" files discussed in §6.2.

Table 3. MTL observational states for DESI targets

State	Description
UNOBS	Unobserved (the PRIORITY_INIT state)
MORE_ZWARN	Ambigous redshift — observe more
MORE_ZGOOD	Good redshift, but observe more
MORE_MIDZQSO	$z < 2.1$ QSO; observe more at low priority
DONE	Enough observations have been obtained

ority lower than "filler" secondary classes are targets observed as part of the DESI backup program.

6.3. *Updating the Observational State*

As the DESI survey progresses and redshifts are obtained that reveal the nature of a source, the priority and observational state of a target are updated in the relevant MTL ledger[9]. Possible observational states for targets are listed in Table 3, and each observational state corresponds to a specific numerical priority for a given target class. For example, an unobserved quasar target has a priority of UNOBS=3400; a quasar target for which a good redshift is obtained — $z \geq 2.1$ for a quasar, corresponding to the Ly-α redshift boundary — has a priority of MORE_ZGOOD=3350; and a quasar target for which observations have been exhausted drops to a priority of DONE=2. Setting MORE_ZGOOD $<$ UNOBS for quasars ensures that pairs that are closer on the sky than the DESI fiber patrol radius are *both* typically observed, because an unobserved quasar has higher priority than one requiring additional observations. The numbers of observations conducted and required for a target are also updated with each acquired redshift, as detailed in §6.3.2 and §6.3.3.

6.3.1. *Redshift Information*

The standard DESI pipeline applies a template-fitting code called Redrock (Bailey et al. 2023) to derive classifications and redshifts for each target. The MTL code considers redshifts and redshift warnings from Redrock when updating the state of a target. These quantities are denoted by Z and ZWARN in the MTL ledgers, and we adopt this notation below.

The ZWARN information from Redrock is crucial for the MTL code to determine whether a sufficiently good observation was obtained to update the state of a target. If an observation has a ZWARN bit-value of BAD_SPECQA, BAD_PETALQA or NODATA[10] set then the observation is considered to not be "good." Such an observation is treated as if it had never been acquired, and the state of the corresponding target is never updated, regardless of the target type. The NODATA bit is set by Redrock (see Bailey et al. 2023, for more details), whereas the BAD_SPECQA and BAD_PETALQA bits — which we will describe here — are set as part of the DESI spectroscopic pipeline (Guy et al. 2023). Note that a good observation may still correspond to a poor redshift fit, where the most such common redshift failures set the SMALL_DELTA_CHI2 bit for low signal-to-noise spectra.

BAD_PETALQA, which denotes low-quality observations across an entire petal, is flagged when any bit in Table 4 is set. Quantitatively, the BADPETALSTDSTAR flag listed in Table 4, which denotes a petal that may have insufficient standard stars to extract high-quality spectra, is set when:

$$\begin{aligned} & N_{\rm good} < 2 \\ \text{OR} \quad & N_{\rm good} = 2 \;\; \& \;\; {\rm rms}(R_{\rm flux}) > 0.05 \\ \text{OR} \quad & {\rm rms}(R_{\rm flux}) > 0.2 \end{aligned} \tag{9}$$

where $N_{\rm good}$ is the number of good standard stars that the spectroscopic pipeline was able to fit and $R_{\rm flux}$ is the ~~ratio of the amount of flux~~ fraction of the expected flux (based on the photometric magnitude) entering the spectrograph~~to the corresponding imaging flux~~. A standard star is defined as a good fit if

$$\chi^2/{\rm dof} < 2 \;\; \& \;\; {\rm SNR(blue)} > 4 \;\; \& \tag{10}$$

$$|\Delta(g-r)| < 0.1 + 0.2E(B-V) \; . \tag{11}$$

Here, the "blue" region of the spectrum and the g- and r-camera magnitudes are detailed in Guy et al. (2023), and the $E(B-V)$ term allows for some flexibility in the assumed reddening correction.

BAD_SPECQA, which denotes a low-quality spectrum for a single DESI observation, is set when any bit in Table 4 *or* Table 5 is flagged. Effective time for a fiber is considered "too" low (i.e. the LOWEFFTIME bit is set) when:

$$t_{\rm eff} \; 10^{2\times 2.165 \;\, \Delta_{E(B-V)}/2.5} < 0.85 \times 0.85 \times {\rm GOALTIME} \; . \tag{12}$$

Here, $t_{\rm eff}$ is the effective integration time through the fiber and

$$\Delta_{E(B-V)} = E(B-V)_{\rm fiber} - {\rm median}(E(B-V)_{\rm tile})$$

accounts for different extinction by Galactic dust through the fiber, as compared to the extinction across

[9] These quantities are recorded in the PRIORITY and TARGET_STATE columns described in the DESI data model.

[10] See, e.g., the **zwarn_mask** bitmask at https://github.com/desihub/desitarget/blob/2.2.0/py/desitarget/data/targetmask.yaml#L230-L248.

Table 4. Flags used to construct the BAD_PETALQA mask

Flag	Description
BADPETALPOS	Fraction of fibers with bad positioning ($> 100\,\mu$m) is > 0.6 (corresponding to > 300 fibers on a petal)
BADPETALSTDSTAR	Too few standard stars or the rms between stars is too large in the petal (see §6.3.1 for more details)
BADREADNOISE	Bad readnoise (> 10 electrons/pixel)

The BAD_PETALQA flag is set if *any* bit in this table is set.

the entire tile. The factors of 0.85, which represent the per-tile and per-fiber minimum amount of integration time needed to complete an observation were set by trial-and-error during DESI Survey Validation (e.g. DESI collaboration et al. 2023). The quantity on the right-hand side of this inequality ends up being 722 seconds in dark time (GOALTIME = 1000 s) and 130 seconds in bright time (GOALTIME = 180 s), reflecting the effective exposure times listed in §5.2.

6.3.2. *General Updates*

The MTL uses a "good" spectroscopic observation to update the state of most targets via a relatively simple algorithm. The number of required observations (called NUMOBS_MORE in the MTL ledgers) is decremented by one and the number of obtained observations (NUMOBS) is incremented by one[11]. In addition, the PRIORITY of a target will be changed to the MORE_ZGOOD or MORE_ZWARN priority if ZWARN is zero or non-zero, respectively, for the acquired redshift. As soon as NUMOBS_MORE drops to zero, a target's priority is set to the DONE priority discussed in §6.2 (which is a very low value of 2 for all target classes). Similarly, if a target has reached a value equal to the DONE priority, then its NUMOBS_MORE value is reduced to zero[12]. Targets for which the MORE_ZGOOD priority is equal to the DONE priority will have NUMOBS_MORE drop to zero after their first ZWARN = 0 spectrum is obtained. Similarly, targets for which MORE_ZWARN is equal to DONE will no longer be observed after their first observation with ZWARN > 0. The MORE_ZGOOD, MORE_ZWARN and DONE priority values are typically identical for both bright-time and dark-time galaxy targets, meaning that such targets are usually only observed once.

6.3.3. *Updates for Quasars*

The logic for updating the MTL state is more complex for DESI primary quasar targets and any secondary targets that have flavor set to QSO in the scnd_mask bitmask[13] discussed in Myers et al. (2023). In particular, to improve information for Ly-α quasars (e.g. Farr et al. 2020), the MTL logic incorporates quasar classifications (denoted IS_QSO_QN) and redshifts (denoted Z_QN) from a line-fitting code called QuasarNET (Busca & Balland 2018; Green et al. 2023).

DESI quasar targets have an initial, unobserved priority of 3400 and are scheduled for 4 total observations. Then, such targets are treated in one of three ways, regardless of whether ZWARN indicates the Redrock redshift is confident or not:

- Quasar targets for which the Redrock redshift is Z ≥ 2.1 *or* which QuasarNET classifies as a definitive high-redshift quasar (IS_QSO_QN==1 *and* Z_QN ≥ 2.1) are denoted "Ly-α" quasars.
- Quasar targets for which the Redrock redshift is $1.6 \leq$ Z < 2.1 *and* which QuasarNET classifies as a definitive mid-redshift quasar (IS_QSO_QN==1 *and* $1.6 \leq$ Z_QN < 2.1) are denoted "mid-z" quasars.
- Otherwise, quasar targets are denoted "low-z."

Quasars in the "Ly-α" category have their priority set to MORE_ZGOOD and their NUMOBS_MORE decremented by one. Quasars in the "mid-z" category have their priority set to MORE_MIDZQSO and their NUMOBS_MORE decremented by one. Quasars in the "low-z" category have their priority set to MORE_MIDZQSO and their NUMOBS_MORE decremented by *three*. As with other targets, quasars are observed until their NUMOBS_MORE drops to 0, or below, at which point they are assigned the DONE priority and NUMOBS_MORE=0.

Note that this schema implies that a quasar target can *never* reach the MORE_ZWARN state during the DESI Main

[11] Note that NUMOBS_MORE will equal NUMOBS_INIT for an unobserved target (just as PRIORITY will equal PRIORITY_INIT).

[12] A target can, technically, be observed again once it has reached the NUMOBS_MORE=0 state — such an outcome is simply rendered unlikely because the DONE priority is very low.

[13] https://github.com/desihub/desitarget/blob/2.5.0/py/desitarget/data/targetmask.yaml#L131.

Table 5. Flags used to construct the BAD_SPECQA mask

Flag	Description
UNASSIGNED	Fiber is not assigned to a known target or sky location
BROKENFIBER	Fiber is broken
MISSINGPOSITION	Location information is missing for this fiber
BADPOSITION	Fiber was placed $> 100\,\mu$m from the target location
POORPOSITION	Fiber was placed $> 30\,\mu$m from the target location
LOWEFFTIME	Effective time for this fiber is too low (see §6.3.1 for more details)
BADFIBER	Fiber is unusable
BADTRACE	Bad trace solution
BADFLAT	Bad fiber flat
BADARC	Bad arc solution
MANYBADCOL	$>$ 10% of the pixels covered by this fiber have bad columns
MANYREJECTED	$>$ 10% of the pixels covered by this fiber were rejected during extraction
BADAMPB	Issues with the amplifier readouts of camera B render this fiber unusable
BADAMPR	Issues with the amplifier readouts of camera R render this fiber unusable
BADAMPZ	Issues with the amplifier readouts of camera Z render this fiber unusable

The BAD_SPECQA flag is set if any bit in this table is set *or* if any bit in Table 4 is set, although, strictly, LOWEFFTIME was not used to inform BAD_SPECQA until April 19, 2022 (see, e.g., https://github.com/desihub/desispec/pull/1722).

Survey. Note, also, that "low-z" quasars may eventually receive two observations as their NUMOBS_MORE will only drop to *one* after their first acquisition. The second observation, however, will be scheduled at a priority (MORE_MIDZQSO) that exceeds only the lowest-priority, highest-density DESI "filler" targets. This choice reflects the low density and relatively high scientific value of even $z < 1.6$ and ambiguously classified quasars.

6.3.4. *Special Cases*

There are two special cases that inform how the MTL ledgers are updated. First, any target that becomes a quasar in the "Ly-α" category is *locked into* that state until it reaches NUMOBS_MORE of 0 and the DONE priority. This provides some insurance in the case of genuine $z \geq 2.1$ quasars having a flawed observation or fluctuating in redshift around $z = 2.1$ due to noise. Second, only primary programs are allowed to determine the state in the primary ledgers *except* in the case of primary targets that are either for calibration or are only in the Milky Way Survey (MWS) program. Such primary targets *are* allowed to be updated by secondary target classes that have updatemws set to True in the scnd_mask bitmask discussed in Myers et al. (2023). This allows the MWS (see Cooper et al. 2022) to better prioritize highly desirable secondary target classes for Galactic science without impacting primary analyses of extragalactic large-scale structure.

6.3.5. *Reprocessing the MTL Ledgers*

Beyond the routine MTL updates discussed in §6.3.2, §6.3.3 and §6.3.4 the MTL ledgers can be fully *reprocessed* when redshift information from the DESI spectroscopic pipeline needs to be altered. This can occur when a DESI hardware glitch is identified after the MTL ledgers have already been updated for certain tiles, or due to improvements in the DESI spectroscopic pipeline software. Reprocessing of the ledgers is achieved by adding new entries to the ledger with the original state of each affected target, and then reprising the MTL updates, in the original tile-order, using the new redshift information.

The root directory for the MTL ledgers includes two "done" files (named `mtl-done-tiles.ecsv` and `scnd-mtl-done-tiles.ecsv`) that list each tile that has been processed through the MTL logic. These files communicate to afternoon planning that a tile's analysis is complete and overlapping tiles may be observed. The files include a column (named ARCHIVEDATE) that records when the redshift information used to update the MTL ledgers that touch a given tile was archived (§5.11). As is the case for the other MTL ledgers, new entries are only ever appended to the "done" files (i.e. no information is ever overwritten). If a tile appears in a "done" file multiple times, then that tile was reprocessed, using information from redshifts on the recorded `ARCHIVEDATE`. The corresponding ledgers will contain entries, in order, for both the original MTL state changes and any updates based on reprocessed redshift information.

6.4. *Other Ledgers*

Two bespoke types of MTL ledgers exist in addition to the five initial sets detailed in §6.1; a single, monolithic ledger listing targets of opportunity (henceforth ToO), and sets of ledgers used to override the MTL logic.

The ToO ledger is read by `fiberassign` to design special tiles to follow up gravitational wave detections, neutrino bursts, or other time-critical events (e.g Palmese et al. 2021). Entries in the ToO ledger can also be used to requisition fibers on existing tiles (see §5.6), although this mode is yet to be used in the DESI Main Survey. The ToO ledgers differ from other MTL ledgers as they contain just the minimal information needed by `fiberassign`, plus columns that are only relevant to time-critical observations.

Override ledgers are used to *force* an observational state *into* an MTL ledger. This is particularly beneficial when rare, high-value targets have been studied using newly available data and found to have a different redshift or classification to that assigned by the DESI pipeline. For example, the override mechanism currently ensures some quasars from a $z \sim 5$ secondary program (Yang et al. 2023) — which have been definitively classified through visual inspection of their DESI spectra — are always available to receive a DESI fiber. Override ledgers closely resemble other MTL files, as they essentially contain the state that will be forced into an MTL ledger.

7. ~~SKY FOOTPRINT~~

~~The Dark Energy Survey Instrument Final Design Report calls for a baseline survey of 14,000 sq. deg. (DESI Collaboration et al. 2016a), with a science fiber density of $\sim 3000/\mathrm{deg}^2$ for the dark program and $\sim 700/\mathrm{deg}^2$ for the bright program. Given the instrumented fiber density of $\sim 600/\mathrm{deg}^2$, this corresponds to each region of the sky being covered by five tiles for the dark program and one tile for the bright program. The bright and dark programs nevertheless requires more passes to target multiple galaxies within a fiber patrol radius and to obtain reasonable completeness on lower priority main survey programs. We describe here the specific implementation of these broad requirements for the dark and bright programs.~~

~~We define a set of 9929 dark tiles and 5676 bright tiles that cover 14,200 sq. deg.: 9800 sq. deg. in the North Galactic Cap and 4400 sq. deg. in the South Galactic Cap. These tiles are distributed among several passes where each pass consists of 1,427 non-overlapping tiles. Approximately 75% of the footprint can be reached by a DESI fiber in a tile in a particular pass. The dark program consists of seven such passes, rotated with respect to one another to fill in gaps between the tiles, while the bright program consists of four such passes. This leads to an average coverage of 5.2 for the dark program and 3.2 for the bright program.~~

~~The pattern of tiles in a single pass is given by the Hardin et al. (2000) icosahedral tiling with 4112 tile centers distributed over the full sphere. This tiling matches the size of the DESI focal plane closely and provides a uniform distribution of tiles with the additional feature that no two tiles overlap one another within a single pass. The fraction of the sky accessible to a given number of tiles for the seven pass dark program and four pass bright program is shown in Figure 2. The geometry of the regions of relatively high and low coverage is complicated, and is shown for the seven-pass dark program in Figure 3.~~

~~The fraction of the sky that is covered by a given number of tiles in the seven-pass dark tiling and the four-pass bright tiling. On average, a given part of the sky is covered by 5.2 dark tiles and 3.2 bright tiles.~~

~~The number of exposures that can reach any particular point of the sky, for the seven-pass dark program, were no areas excluded (e.g., due to low Galactic latitude or low declination). The twelve star-like regions with with slightly lower coverage corresponds to the points of the underlying icosahedral tiling of Hardin et al. (2000).~~

~~The goal of the DESI tile selection was to select a large, contiguous region that could be efficiently observed for extragalactic targets as part of a year-round survey from Kitt Peak. These objectives imply limits on declination to avoid tiles that are only available at high airmass, and limits on extinction and Galactic latitude to avoid regions where extragalactic targets are both extinguished and more often blended with Milky Way stars.~~

~~We define the footprint as follows:~~

1. ~~In the footprint of the DESI Legacy Imaging surveys Data Release 9~~
2. ~~$-18° < \delta < 77.7°$~~
3. ~~$b > 0°$ or $\delta < 32.2°$~~
4. ~~$|b| > 22°$ for $-90° < l < 90°$, otherwise $|b| > 20°$~~

These constraints produce the footprint shown in Figure 4.

The footprint of the DESI survey resulting from the constraints of §3. Tiles are colored by the amount of time it would take to reach a fixed intrinsic galaxy depth, relative to observing at zenith in the absence of Galactic extinction. This is $f_{\mathrm{dust}} f_{\mathrm{airmass}}$, from Equations 1 and 2. Airmasses are computed using the design airmasses resulting from the optimization of §4.1. The Galactic plane is shown as a dotted gray line, and the gray contour shows $E(B-V) = 0.3$ mag. Tiles in extinguished regions and at the declination bounds of the survey are most expensive, owing to both atmospheric and Galactic extinction.

Though we have imposed no explicit cuts on Galactic extinction, we only target regions of the sky with imaging from the DESI Legacy Imaging Survey. That survey explicitly avoided high $E(B-V)$ regions, so these regions are naturally avoided in the DESI footprint without need for further adjustment. Cuts on Galactic latitude do trim the edges of the imaging footprint slightly, however.

The trend in exposure factor with declination in Figure 4 comes from the dependence of survey speed on airmass (§5.3). The SGC is significantly more expensive than the NGC due to a combination of extinction and airmass. No Legacy Survey imaging was available in the SGC north of $\delta = 32°$, though this region would otherwise be favorable for extragalactic studies. The irregular small-scale variation comes from Galactic extinction.

The sky area within 1.6 of at least three tiles for the seven pass dark program is 14,246 sq. deg..

All main survey tile coordinates are rounded to the nearest 0.001 to improve legibility.

6.1. *Adjustments to tile centers*

The simple footprint definition of §3 describes our basic footprint selection strategy. Many tile centers are additionally adjusted to avoid bright stars.

The wide field of view (3.2) of DESI means that bright stars cannot be completely avoided. However, bright stars are particularly damaging if they fall in a few special parts of the DESI focal plane.

First, it is problematic if a very bright star falls on a GFA. These can make it challenging to guide the telescope. Worse, the filter on the GFA reflects light falling outside of the GFA bandpass. Light from the bright star then ends up adding to a large out-of-focus ghost image covering a substantial portion of the DESI focal plane. This is avoided by shifting the tile centers to move bright stars off of the GFA filters. For tiles where a star with Gaia magnitude $G < 6$ lands nears a GFA, we searched for the smallest shift in RA or Dec, in steps of 10 arcseconds, that would put the star at least 25 arcseconds from a GFA.

Second, data from a petal can be rendered useless if a fiber is placed directly on a bright star, saturating large parts of the detector. This is mostly avoided by re-positioning such fibers (which will never have valid main survey targets) away from bright stars. But in rare cases a non-functional fiber happens to land on a very bright object. We adjust tile centers in these cases. After finding bright stars that land near the current set of non-functional positioners for each tile, we search for a small offset (up to 15 arcseconds) of the tile centers in order to minimize the total star light reaching non-functional positioners.

We periodically compute new offsets for tile centers to account for new or bumped non-functional positioners, but we do not do this on the fly when designing each tile.

7. SURVEY PERFORMANCE

Planning the DESI survey requires predicting the amount of effective time the survey can deliver over the year. The amount of effective time delivered depends on the point spread function delivered to the focal plane (§7.1), the transparency of the night sky (§7.2), the sky brightness (§7.3), the overall survey speed (§7.4), and the time off sky due to weather and technical downtime (§7.5).

In this and subsequent sections, we study the performance of the DESI survey from 2021–05–14 to 2022–06–14. The start date corresponds to the start of the DESI main survey; after this point we limited engineering observations and observed almost exclusively main survey tiles. The stop date corresponds to the beginning of a long shutdown due to damage to Kitt Peak infrastructure from the Contreras wildfire. The DESI survey restarted operations on 2022–09–11; we do not include this more recent data here.

We compare DESI's performance with expectations from the Mayall Telescope's long history. The Mayall has been observing the sky since 1973, providing a historical record of seeing, transparency, sky brightness, and downtime, based on the tireless, careful effort of the Mayall's observers. We focus here particularly on the record from 2007–2017, where records were most readily available. We compare DESI's observed performance with simulations based on on this historical record (§8).

An important concept in ~~the~~ this section is the survey "margin": the amount of time available to the survey divided by the time needed to finish the survey, minus

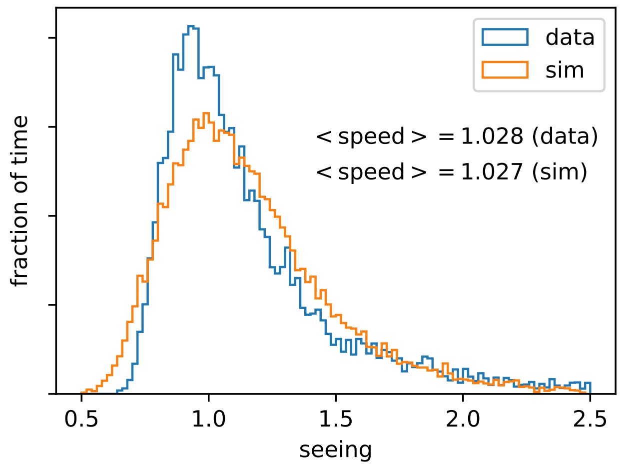

Figure 8. DESI delivered point spread function FWHM. The blue curve shows the measurements from the ~~guide arrays~~ GFAs during the DESI survey, while the orange curve shows data from simulations based on the MzLS. The inferred average survey speeds for both the real data and the simulated data (proportional to the square of the fraction of flux entering a fiber) is given for each case, and agree closely.

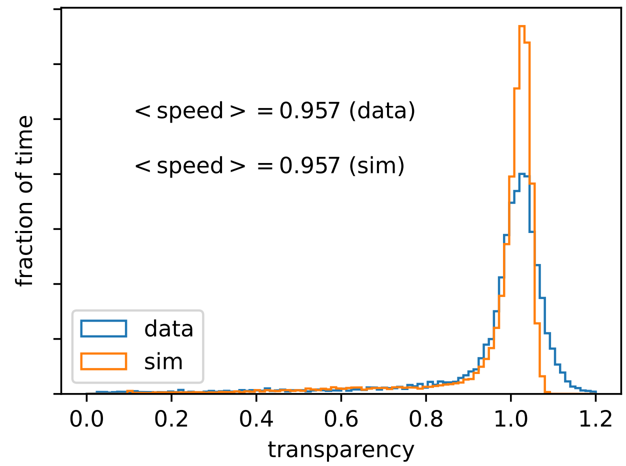

Figure 9. DESI observed transparency. The simulations show a narrower distribution of transparencies than observed, due to the simulations having deconvolved the observed distribution slightly to reduce the effect of measurement errors. The inferred average survey speeds are proportional to the square of the transparency, and are identical between observations and simulations by construction.

one. DESI aims to operate with a healthy margin to enable finishing the survey in the allotted five year survey window. Factors which speed the survey by a certain percentage increase the margin by the same percentage, in the limit that the margin is close to zero.

7.1. *Point spread function*

The DESI corrector was designed to contribute negligibly to the PSF delivered to the focal plane. This means that historical records from for example, the Mayall z-band Legacy Survey (Dey et al. 2019, MzLS), can be used directly to predict DESI's seeing. Comparison of predictions from simulations (§8) and the actual seeing in the first year of the survey show good agreement, as shown in Figure 8. The observed distribution is somewhat tighter than the simulated data based on the MzLS, plausibly due to DESI's improved control of focus using the ~~guide focus array cameras~~GFAs. However the overall inferred average speed (the square of the fraction of source flux entering a fiber, the critical element to survey planning) agrees closely with expectations from MzLS and the survey simulations.

7.2. *Transparency*

Similarly, survey planning and simulations assume that the transparency of the night sky as seen by DESI will closely match the historical performance obtained by MzLS. Again, predictions from simulations and DESI's observations in the first year show good agreement, as shown in Figure 9. The average survey speed, proportional to the square of the transparency, shows excellent agreement between the data and the simulations, though this is by construction.

An unexpected challenge in matching the observations to the simulations stems from the definition of "photometric." The distribution of transparencies seen by DESI (Figure 9) is strongly peaked near unity, but the peak has a width of about 3.5%. This width partially reflects measurement uncertainties, but also appears to reflect true variations in the transparency of the night sky, as confirmed by comparison with the amount of light delivered to the spectrographs and seen by the ~~guide cameras~~GFAs. The nights that were used to define a transparency of 1 for DESI were $\sim 3\%$ less transparent than the peak of the transparency distribution. The results shown in Figure 9 have been updated to account for this discrepancy.

7.3. *Sky Brightness*

Survey planning focused on the main dark program, with less emphasis on the bright program, which accounts for only roughly 10% of the survey effective time. The sky brightness when the moon is up is a relatively complex function of the moon phase, location, and the line of sight. However, when the moon is down, our model of the sky brightness is a simple function of airmass. Survey planning then chose an extremely simple description of the sky brightness: equal to a nominal dark sky brightness when the moon is down; equal to

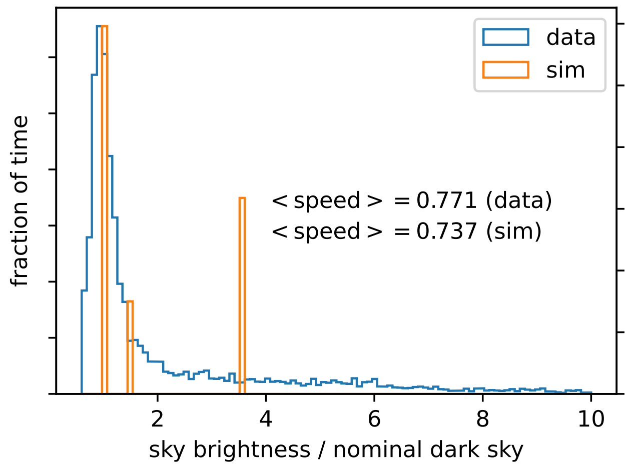

Figure 10. DESI observed sky brightness, relative to a nominal dark sky brightness of 21.07 mag. The observed sky brightness peaks about 7% darker than this. The sky brightness models in the simulation are very simple, assigning a sky brightness of 1, 1.5, or 3.6 depending on the phase and location of the moon. The overall average survey speed, proportional to one over the sky flux, are reasonably well-matched, though the simulations are 8% slower largely due to the slightly darker peak of the observed sky distribution than the simulated sky distribution.

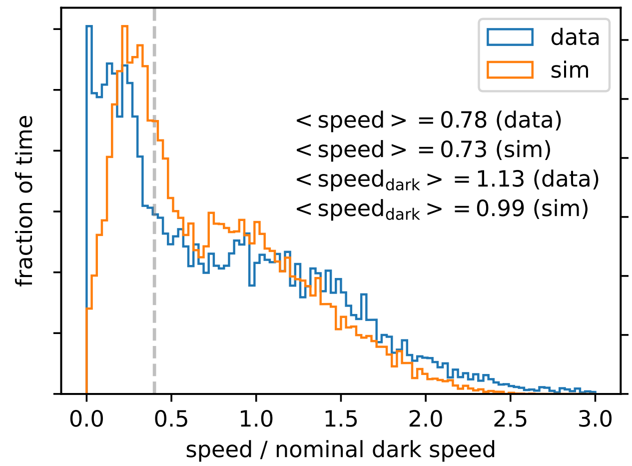

Figure 11. DESI delivered survey speed, compared with speeds delivered in the simulations. This is the product of factors relating to seeing, transparency, and sky brightness. The average delivered speed is 7% higher in the actual observations, but 14% higher when limiting to exposures taken in dark conditions. Note that the small difference between the average speeds here and in Figure 6 comes from the fact that here the speeds are computed from the measured seeing, transparency, and sky brightness, and in Figure 6 they are computed from the effective time delivered on each tile.

1.5× nominal when the moon is up but less than 60% illuminated and the product of the moon phase and distance from the horizon was smaller than 30 degrees; and equal to 3.6× nominal otherwise. Figure 10 compares this simple model in the simulations with DESI's observations. The work of Hahn et al. (2022) includes an improved sky model important for accurate modeling of the bright program.

This model is clearly limited, but because dark, moon-down time is the source of most of the survey's effective time, it is largely adequate. The average survey speed, proportional to one over the sky flux, is about 8% faster in the actual data than in the simulations. This is largely because the dark sky brightness peaks 7% darker than the nominal 21.07 mag forecast in survey planning.

7.4. *Overall speed*

The total delivered survey speed is a combination of the seeing, transparency, and sky brightness. Breaking these terms out separately, one expects the simulations to run 8% slower than the actual observations due to the different sky brightness modeling. Figure 11 compares the actual total delivered speeds in the simulations and observations.

The observed average survey speed has been 7% faster than expected in the simulations, consistent with the difference in sky brightness. Additionally, the variance in the observed speeds is larger than predicted by our simple simulations, leading the average speed in the dark program—observed when conditions are good—to be 14% larger than in the simulations. This is the largest factor in leading to discrepancies between the observed and expected survey progress (see Table 10). This is largely driven by times when the skies are especially dark.

7.4.1. *Solar Cycles*

The DESI survey started survey validation near the start of solar cycle 25. The next solar maximum will occur in 2025, near the end of the DESI main survey. It is therefore likely that sky brightness distribution observed so far is darker than what we will have for the remainder of the survey (Walker 1988; Patat 2008; Noll et al. 2012). The impact of the solar cycle on DESI's overall performance will depend on the amplitude of the solar cycle. Investigations using past data from the Sloan Digital Sky Survey and its extensions, as well as the DECam Legacy Survey (Dey et al. 2019), suggested potential impacts on survey speed of between 5% and 20%. For comparison, Patat (2008) measure an approximately 30% difference in dark sky brightness between solar minimum and solar maximum.

7.5. *Downtime*

Another key element of survey planning is the amount of time the system is down, due to bad weather or technical problems with the system. DESI's downtime has been very close to expectations, with the exception of two significant shutdowns in the summer of 2021 and 2022. Table 6 lists the time lost to various causes, and the total time remaining. We exclude the second shutdown from the time range considered in this work, but we describe it briefly here for completeness.

The first shutdown of the DESI main survey was from 2021–07–10 to 2021–09–20, when the focal plane electronics were upgraded. The second shutdown was from 2022–06–14 to 2022–09–11, when a wildfire swept through Kitt Peak, requiring repair to the site's infrastructure. Such large events are not directly incorporated into planning, and instead come out of the overall survey margin. However, survey planning does include a nominal three week shutdown during Arizona's summer monsoon season, when nights are shortest and frequent clouds and rain slow observing. Both shutdowns occurred during monsoon season, leading them to have a much smaller impact on survey progress than suggested by their duration.

Outside of these two shutdowns, DESI's downtime has been very modest and consistent with expectations. The DESI performance database tracks the state of the system every second, recording a wealth of information, including whether the spectrograph shutter is open, whether the telescope is guiding, and whether the system is in a weather, instrument, telescope, or other hold. Defining "on sky" time as time when the spectrograph shutter has been open while guiding within the last 2.5 minutes (to cover overheads between exposures and long slews), DESI has spent 76.6% of its time on sky during "dark time". Here we define "dark time" as time on nights more than two days from full moon, with the sun more than 15° below the horizon, with the moon down, and outside of one of the two major shutdowns. The majority of the downtime (22.2% of the dark time) is due to the weather, with another 1.4% due to instrument downtime and less than 1% to other sources.

The instrument has has met the goal of $< 2\%$ downtime, and other sources of downtime are negligible for planning purposes. The weather loss of 22.2% is typical for the Mayall outside of the major shutdowns DESI has experienced. Specifically, replaying the years 2007–2017 as if they were 2021–05–14 to 2022–06–15, excluding time during major shutdowns, and weighting nights by the length of the night between 15° twilight, the Mayall would have been closed due to weather 23.7% of the time on average, with a standard deviation of 3.6%; DESI's observed weather loss so far of 22.2% is typical.

Table 6. Dark Time Spent on Sky or Down

Category	% of moon-down time	% of all time
On sky[a]	76.6%	69.1%
Open shutter	66.2%	58.4%
Any recorded loss	23.7%	30.3%
Weather loss[b]	22.2%	27.9%
Instrument loss	1.4%	1.9%
Telescope loss	0.2%	0.4%
Other loss	0.4%	1.2%

Fraction of time spent either on sky or down, according to the DESI performance database. We tabulate values for both "moon-down" time and "all" time. "All" time includes all time outside of monsoon shutdowns with the sun more than 12° below the horizon. "Moon-down" time is the subset of "all" time where the moon is below the horizon and excluding four nights around full moon. Engineering activities take priority around full moon, and are concentrated in moon-up time in general, leading to better on-sky fractions in moon-down time. The various sources of loss need not sum to the "any recorded loss", since the system can be down for multiple reasons simultaneously. Small differences between 100% and the sum of the on sky time and the any recorded loss time can stem from the definition of "on sky" time.

[a] On sky time is defined as time within 2.5 minutes of a moment when the spectrograph shutters were open and the telescope was guiding.

[b] The weather loss here tabulates both time the observers mark as being lost due to weather as well as time when the instrument control system was not in "observing" mode. The latter case usually corresponds to nights that cloud out but where the observers do not mark the time as a weather loss. However, other, more rare cases will be incorrectly grouped with weather loss here.

The amount of time available for observation with the Mayall per month is given in Table 7, based on the years 2007–2017. This table uses the time between 15° twilight, adjusted for seasonal variability in the weather. We have not removed planned engineering time around full moon or during the annual monsoon season, however, because the alignment of these shutdowns with month boundaries can artificially increase variability.

As noted earlier, survey planning includes a three week shutdown around full moon during the Arizona monsoon season. So far, our monsoon season shutdowns have been significantly longer than forecast there, owing to electronics upgrades and the Contreras wildfire. On the other hand, we would plan to run DESI through the monsoon season if weather and engineering requirements allowed. Table 7 gives a sense for how much that ad-

Table 7. Weather-adjusted hours available per month

Month	Hours	Month	Hours
January	240 ± 47	July	104 ± 21
February	211 ± 25	August	148 ± 26
March	240 ± 21	September	191 ± 26
April	216 ± 16	October	258 ± 36
May	201 ± 14	November	254 ± 24
June	185 ± 22	December	222 ± 27
Annual	2468 ± 80		

The number of hours available for observation with the Mayall per month, accounting for varying weather and the changing length of the night, but excluding engineering and monsoon shutdowns. Uncertainties reflect year-to-year standard deviations due to weather.

Table 8. Amount of Effective Time per Year

Parameter	Value	Notes
Time per year[a]	3481 hr	Planning
Open shutter fraction	58.4%	observed
Fraction of dark time[c]	59.3%	observed
Fraction of bright time[c]	34.8%	observed
Fraction of backup time[b,c]	5.9%	observed
Average dark speed	1.148	observed
Average bright speed	0.293	observed
Average backup speed[b]	0.096	observed
Average overall speed	0.789	observed
Dark effective time per year	1383 hr	computed
Bright effective time per year	207 hr	computed
Backup effective time per year[b]	12 hr	computed
Number of dark tiles	9929	design
Number of bright tiles	5676	design
Effective time for dark tiles	1000 s	design
Effective time for bright tiles	180 s	design
Effective time for backup tiles	60 s	design
Mean airmass & dust adjustment	1.51	design
Dark time needed per year	833 hr	computed
Bright time needed per year	86 hr	computed
Outside major unplanned shutdowns[d]	87%	observed
Time on tiles not counted[e]	2%	observed
Average dark tile overexposure[f]	2%	observed
Dark margin, computed	39%	computed
Bright margin, computed	105%	computed
Dark margin, observed	36%	observed
Bright margin, observed	93%	observed

Parameters controlling the amount of effective time available to the survey (top of table), compared with parameters controlling the time needed to complete the survey (middle of table).

[a] The number of hours derived from ephemerides; see § 7.6 for details.

[b] Backup program parameters are especially uncertain because backup tiles were not regularly observed until December 2021.

[c] We are defining the time available to the program according to the amount of time selected for that program based on the NTS program selection. See §5.3 for more details.

[d] This fraction is the expected time available to the survey given the long summer 2021 shutdown divided by what the survey would have had with the planned shutdown.

[e] Tiles "not counted" as main survey tiles were either observed for other programs (1%) or discarded (1%).

[f] The average completed dark tile has 1.02× the required effective time.

justment would speed the survey—recovering the bright part of July would be roughly a quarter as valuable as a January.

7.6. *Effective hours delivered per year*

When planning programs for DESI, it can be valuable to have a sense for the total number of effective hours DESI can deliver in a year. Table 8 tabulates some key numbers for making this calculation.

We were able to get a good match between the observed dark margin and the margin expected from a relatively simple calculation based on the number of hours available to the survey and the survey's average speed in different programs. The calculations count every hour with the sun more than 12° below the horizon, excluding an 18 night shutdown around full moon each monsoon season for engineering purposes.

Matching the computed margin to the actual margin requires accounting for the longer-than-expected DESI shutdown in the summer of 2021 (§7.5). Other small adjustments are needed to account for time DESI has spent on tiles for programs other than the main survey (1%) and on exposures that needed to be discarded (e.g., due to wind shake, or temporary instrument problems; 1%).

Note that this calculation folds in true values of critical parameters DESI achieved during the 2021–05–14 to 2022–06–15 time window under consideration—it uses the observed open shutter fraction and the observed average speeds and fractions of time in different programs. This effectively folds in the real weather and conditions that DESI has experienced and all technical downtime.

These values are useful for the planning of future DESI-like surveys, but the match between the observed DESI margin and the computed value from this computation is somewhat artificial.

We can check the consistency of this table by comparing the number of hours accumulated on dark tiles between 2021–05–14 and 2022–06–15 with the expectations from this table. On the basis of the ephemerides, there are 3248 total hours excluding the long shutdown in the summer of 2021. Using the open shutter fraction, fraction of time in the dark program, and average dark program speed from Table 8, we obtain 1291 effective hours at zenith through no extinction. Counting all time accumulated on dark exposures in that window, and adjusting by Equation 1 and Equation 2 to account for extinction and airmass, we obtain 1247 observed effective dark hours. These are different by 3.5%. Much of the difference is "time on tiles not counted", e.g., time we spent observing tiles for special programs or tiles that we eventually deemed bad. Another issue surrounds the accounting for engineering time; engineering time spent on guided observations with the spectrographs open counts as open shutter time in Table 8, though this kind of open shutter time needs to be separately accounted when computing the amount of time DESI can deliver on science tiles. Still, these are small effects, and Table 8 provides a useful description of the number of effective hours the DESI system can deliver.

8. SURVEY SIMULATIONS

We perform survey simulations to verify that the DESI survey will complete in its allotted five-year mission. The survey simulations step through the survey at ten second intervals each night of observations. The simulation generates a realistic realization of the observing conditions (seeing, transparency, sky brightness) based on modeling of past observing conditions from the Mosaic z-band Legacy Survey (Dey et al. 2019, MzLS). Downtime due to weather is also included, following patterns from observations at the Mayall from 2007–2017.

At each time step, if the system is not already observing, a new tile is selected, and the telescope begins tracking a new field overhead (Table 9). Otherwise, when the system is observing, effective time is accumulated according to the current seeing, sky brightness, and transparency. Observing continues until the tile is complete or the tile needs to be split or abandoned due to overly long exposures or too-high airmass. When splitting, a separate tile split overhead is incurred (Table 9). Weather-related downtime may also close the dome at any point, stopping the current observation and advancing the simulation to the next time the dome opens.

Table 9. Selected Survey Simulation Parameters

Parameter	Value
Nightly beginning & end of observations	15° twilight
New field overhead	139 s
Split exposure overhead	70 s
Engineering nights per lunation	4
Monsoon shutdown nights per year	18

A selection of important parameters in the simulations, and their values.

The survey simulations use the same airmass optimization and next-tile selection algorithms as the real survey. Accordingly the simulations follow the same moon & planet avoidance algorithms as the real survey. They use a simplified model of the ETC and a simple model of the instrument. They model only per-tile quantities and ignore any details relating to individual fibers and target selection; the survey simulations seek only to accumulate the required effective time on each tile.

The survey simulations include realistic models of the weather based on historical data from the Mayall. Comparisons of modeled seeing, transparency, sky brightness, and delivered speed are shown in Figures 8, 9, 10 and 11. The sky modeling in the simulations is rudimentary, but the seeing and transparency distributions match the observations closely. Moreover, the time correlation of the variations in the seeing and transparency is modeled with a Gaussian process, with power spectral densities chosen to closely match observations from the MzLS. That said, the accuracy of the time correlations of variations in the weather makes only a minor impact on survey planning.

Overheads due to stopping and splitting exposures are modest. For the dark program as of 2022–10–04, the mean exposure time is 1093 s, over 3725 observations of 2913 tiles. This implies an overhead of about 9%, which is well captured by the simulations. Slew time is ignored in the simulations, and would account for an additional overhead of about 3%, using the slews from a simulated survey and a realistic model for the telescope slew time as a function of the change in hour angle and declination.

The survey simulations can incorporate past data and use them to make forecasts for the future given different scenarios. This is valuable to, for example, understand the impact of different planned maintenance activities requiring shutting down the telescope to the final survey margin.

8.1. *Comparing survey simulations with the observed survey progress*

Figure 12 shows an example survey simulation run. For this run, we chose to exactly duplicate DESI long summer 2021 shutdown, as described in §7.5. No additional sources of downtime were included except for normal weather losses, which were chosen to replicate randomly-sampled years of the Mayall's historical weather record.

The survey simulation matches the dark program reasonably well. In the survey simulation, ~~26.53~~26.96% of the dark program is completed before 2022-09-21, while in the real survey, 28.97% of the survey was completed. The DESI survey is proceeding 7% faster than forecast in the simulations, our top line result. However, the comparison is complicated by the different average speed in the dark program in the simulations than in reality; see §7.4. Accounting for this makes the dark program 14% faster while being the active program on the telescope for 3% less time than expected. Additional minor differences between the simulations and real observations are that the simulations neglect slew overheads and technical downtime (3% and 2% effects). More importantly, the simulation year one weather realization is particularly poor, with 11% more lost time than DESI observed from 2022–05–14 to 2023–06–15, outside the summer 2021 shutdown. Finally, 2% of the time in the real survey was spent either on tiles we end up discarding or on tiles that were not for the main survey, and another 2% of time was spent overexposing dark tiles. Table 10 summarizes the different contributions to discrepancies between the simulation completeness and the observed completeness. We conclude that the main survey is running 4% slower than we would expect from the simulations after accounting for all of these effects, which we consider good agreement.

Table 10. Contributors to differences in dark margin

Cause	Fraction
Observed progress through 2022-06-14	29.0%
Simulated progress through 2022-06-14	27.0%
Expected effective time through 2022-06-14	21.4%
Dark speed	+14%
Fraction of time in dark program	−3%
Neglected slew time	−3%
Neglected technical downtime	−2%
Actual weather versus simulated	+11%
Time on tiles not counted	−2%
Dark tiles are overexposed	−2%
Adjusted simulated completeness	30.2%
Ratio of observed and simulated completeness	+7%
Ratio of completeness after adjustments	−4%

Important contributions to the difference between the observed completeness in the simulations and the actual observed completeness of the survey. The signs are chosen so that improving the simulations would change the simulated completeness in the indicated direction. A number of minor effects are present, which together would lead the simulations to run 12% faster, exceeding the 7% difference between the observed and simulated completeness. A large number of effects come into play.

We have focused on the dark program, which accounts for most of DESI's effective time, and for which the survey simulations are best suited. The bright program is running much faster than expected from the simulations, due primarily to the following:

- The simulations include no observations when the sun is within 15° of the horizon; in fact we aim to start observing the backup program at 10° twilight and the bright program at 12° twilight.
- The simulations include no observations within 4 days of full moon; in practice, this time is often used for observing when no engineering work is planned.
- The simulation sky modeling in bright conditions is rudimentary. (§7.3).

The bright program was more than 40% complete prior to the summer 2022 shutdown, after little more than a year of main survey observations! This program will need to be expanded in order to accommodate the available time.

9. CONCLUSION

The Dark Energy Spectroscopic Instrument's main survey began on 2021–05–14, and has observed more than 14 million galaxies and 4 million stars through 2022–06–14. The success of the survey has relied on the efforts and dedication of a large science collaboration, instrument, and operations team. The DESI instrument's performance largely exceeds expectations; the data management, processing, and analysis routinely delivers high quality redshifts within hours of observation, even while accommodating last-minute changes in instrument configuration & calibrations; and the operations team has put together a robust system to feed back past obser-

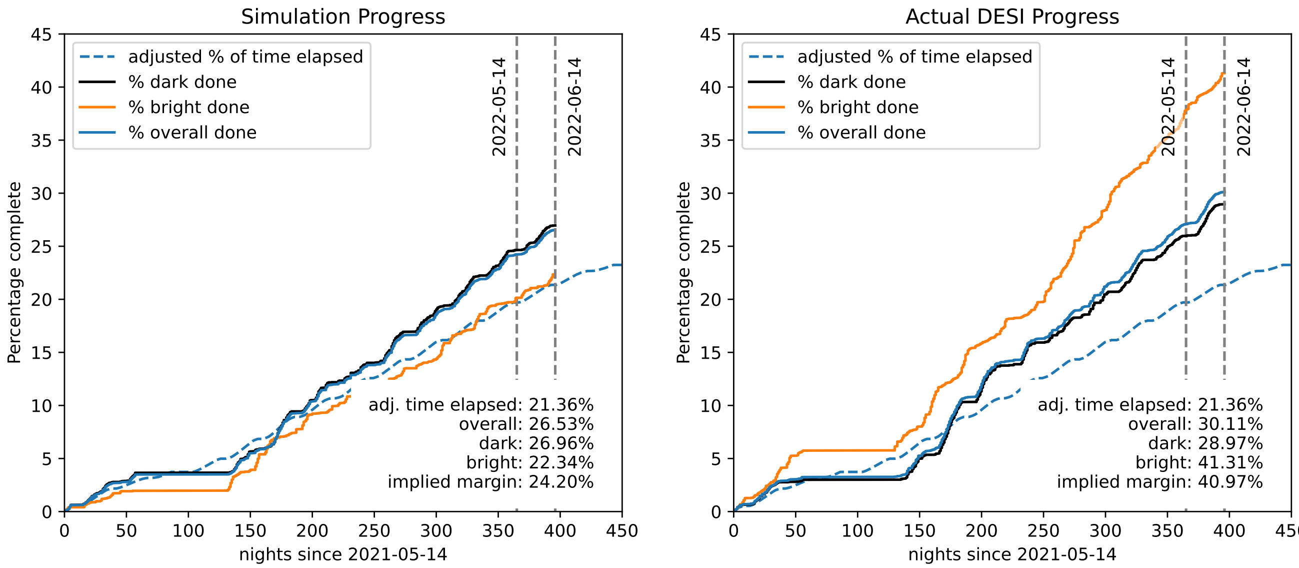

Figure 12. DESI observed progress compared with a nominal simulation using the same major shutdowns. The dark time progress of the simulation is a good match for the observed dark time progress; in the simulation, 26.53% of the dark program was completed before 2022-06-15, while in the real survey, 28.97% of the survey was completed. The fraction of time elapsed is shown with a dashed line, weighting nights by the length of the night, historical weather loss, and removing nights near full moon and planned monsoon shutdowns; see §8 for details. These survey simulations match the progress of the bright program poorly, however, with the actual bright survey progress running ahead of the simulations by almost a factor of two. This is due to limitations of the sky brightness modeling in the simulations, as well as the use of more time in twilight and near full moon for bright observations than expected.

vations into the design of future observations on a daily basis, while identifying and removing problematic observations. The collaboration's realization of the scientific potential of these observations is now underway.

We have laid out the choices made in the survey strategy—the survey footprint, the amount of observing time needed on each tile, the hour angles at which the tiles should be observed, and the tiles' priorities. The decision to require that all observations be fully processed before making subsequent overlapping observations allows the survey to reobserve any $z > 2.1$ quasar discoveries, and places strict requirements on the daily operations design and plan. We detailed the steps of the daily operations loop largely implied by this decision, from afternoon planning to nightly observations to data reduction to updating DESI Merged Target Lists. These Merged Target Lists play a central role in tracking DESI observations in operations, and we described the details of their construction and updates following targets' observation.

We also described the survey performance, which has somewhat exceeded projections made on the basis of historical data from the MzLS—the sky has been slightly darker than we expected. Instrument downtime has been kept low (excepting a major shutdown during the summer monsoon season for upgrading the focal plane electronics), leaving the survey with a healthy 36% margin on 2022–06–14. We compared the observed survey performance with detailed simulations and found good agreement, increasing our confidence in the simulations' value for predicting survey performance.

The first 1.1 years of DESI's operations have been an exciting success, and we look forward to a long, productive future for the instrument.

ADM was supported by the U.S. Department of Energy, Office of Science, Office of High Energy Physics, under Award Number DE-SC0019022.

This material is based upon work supported by the U.S. Department of Energy (DOE), Office of Science, Office of High-Energy Physics, under Contract No. DE–AC02–05CH11231, and by the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility under the same contract. Additional support for DESI was provided by the U.S. National Science Foundation (NSF), Division of Astronomical Sciences under Contract No. AST-0950945 to the NSF's National Optical-Infrared Astronomy Research Laboratory; the Science and Technologies Facilities Council of the United Kingdom; the Gordon and Betty Moore Foundation; the Heising-Simons Foundation; the French Alternative Energies and Atomic Energy Commission (CEA); the National Council of Science and Technology of Mexico (CONACYT); the Ministry of Science and Innovation of Spain (MICINN), and by the DESI Member Institutions: https://www.desi.lbl.gov/collaborating-institutions. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U. S. National Science Foundation, the U. S. Department of Energy, or any of the listed funding agencies.

The authors are honored to be permitted to conduct scientific research on Iolkam Du'ag (Kitt Peak), a mountain with particular significance to the Tohono O'odham Nation.

Facilities: DESI

Software: astropy (Astropy Collaboration et al. 2013, 2018, 2022)

APPENDIX

A. AIRMASS OPTIMIZATION

The DESI airmass optimization scheme works by assigning local sidereal times to tiles and computing the total time necessary to observe the tiles given that assignment. It aims to minimize a cost C:

$$C = T + R \tag{A1}$$

$$T = \frac{T_p - T_0}{T_0} \tag{A2}$$

$$R = \frac{1}{\sum_i^n P_i}\left(n \sum (sA_i - P_i)^2\right)^{1/2} \tag{A3}$$

$$s = \sum P_i / \sum A_i \,, \tag{A4}$$

where T_p is the total time needed to observe the survey given the planned local sidereal times and the implied airmasses, and T_0 is the time that would be needed to observe the survey were all tiles observed at an hour angle of 0. P_i and A_i are the number of planned and available hours in a particular bin i of LST, and n is the total number of bins of LST used. Note that hour angles HA and LSTs are related by $HA = LST - \alpha$, and that assigning an LST to a tile is equivalent to assigning an hour angle to a tile, since each tile has a defined right ascension α.

More explicitly, the total times T_P and T_0 are given by

$$T_0 = \sum T_{0,i} \tag{A5}$$

$$T_P = \sum T_{H,i} \tag{A6}$$

$$T_{H,i} = G_i \, 10^{2\times 2.165\times E(B-V)/2.5} X_{i,H}^{1.75} \tag{A7}$$

where $T_{H,i}$ is the estimated time needed to observe tile i at an hour angle of H, $X_{i,H}$ is the airmass of tile i at hour angle H, and G_i is the goal time for a tile (1000 s for a dark tile or 180 s for a bright tile). Note that sky brightness variations due to the moon are not accounted for here, and that one obtains the same solution for any G as long as it is constant in a program, as for DESI.

The term T (Equation A2) is proportional to the total observing time (up to an additive constant); we want to minimize it. The term R (Equation A3) is the root mean square difference between the binned, planned LST distribution and the available LST distribution. It is zero if the distribution of LST available to the survey exactly matches the planned distribution of LST. An alternative optimization algorithm would force these two quantities to match; the approach taken here allows these to diverge but includes the divergence in the cost function C. For DESI we choose bins 1.875° in size when binning the available and planned LST distributions A_i and P_i.

Our approach to assigning LSTs to tiles starts with an initial guess. This initial guess is then optimized by a simulated annealing algorithm, which perturbs the assignment to try to reduce the cost C.

To create the initial hour angle assignments, we first construct the cumulative distribution function of the tiles' observational costs as a function of right ascension, $CDF_O(\alpha)$. To construct this, we need to know what the observational cost of a tile is, and for that we need the tile's airmass—but we do not know the tile's airmass because we have not yet assigned it an hour angle. For this initial guess we presume that all tiles will be observed with an hour angle of zero. We also construct the cumulative distribution function of the available LST, $CDF_L(L)$, choosing $CDF_L(L_{\mathrm{start}}) = 0$ and integrating around the circle. We then find for each right ascension α_i the corresponding LST L_i such that $CDF_O(\alpha_i) = CDF_L(L_i)$. Conceptually, this corresponds to matching the first 10% of the tiles in right ascension to the first 10% of the LSTs (starting from L_{start}), and so on, until all tiles have been mapped to LSTs. This gives a mapping of tiles to LST that provides the initial guess for the simulated annealing. The only free parameter in this initial guess is L_{start}, the LST at which to start the cumulative distribution function; this corresponds to the LST to which to map tiles with $\alpha = 0°$. We choose a number of L_{start} values around the unit circle and use the L_{start} with the best score to produce the initial guess.

The simulated annealing process consists of a number of steps. In each step, we start by identifying LSTs where changing the assignment of LSTs to tiles by one bin in LST would most significantly improve R, the component of the cost coming from the difference between the planned and available times. These bins are identified by finding the locations where $|\Delta(sA_i - P_i)|$ is largest, where Δ represents taking the difference between bin i and bin $i-1$. One of the top five such bins is selected at random. A scale factor is chosen from a Rayleigh distribution. The LST of each tile j in the selected bin is adjusted by the scale factor and the new survey cost C_j is computed. The new plan with the minimum C_j is chosen (if any is better than the original C), and the process repeats. If instead no improvement was found, instead 20% of tiles are selected at random. Then again the LST assignment of each of these tiles is perturbed, the new cost C is computed, and the assignment with the best C is kept.

The simulated annealing steps are grouped into rounds. Each round consists of one simulated annealing step per tile in the program being optimized (i.e., 9929 steps for the dark program, and 5676 steps for the bright program). When a round is complete, the LST assignment to tiles is mildly smoothed. Each tile's hour angle is replaced by $H_i' = (1-\alpha)H_i + \alpha\bar{H}_i$, where $\bar{H}_i$ is the hour angle map convolved with a Gaussian with a length of 10°, and α is a parameter between 0 and 1 reflecting how aggressively to replace the hour angles with the smoothed version. This smoothing is expected to improve the cost, because the optimal solution should assign LSTs to tiles in a spatially smooth manner. Next, the perturbation scale is reduced to 95% of its previous value, from an initial values of 1°. Finally, α is reduced to 95% of its previous value, from an initial value of 5%. Then another round of simulated annealing is performed with the updated parameters. Rounds continue until both $R < 0.02$ and the fractional improvement in C is less than 1%, $C_i/C_{i-1} - 1 > -0.01$, where i indexes rounds.

In practice, the simulated annealing scheme does not shift the solution far from the initial guess. The primary limitation of the initial guess is that it gives all of the tiles at the same right ascension the same LST. An optimal solution, however, keeps tiles at low declination close to hour angles of zero and preferentially uses tiles at high declination to fill in the LST distribution. Experiments with alternative optimization schemes only improved the cost by roughly half of one percent.

REFERENCES

Alexander, D. M., Davis, T. M., Chaussidon, E., et al. 2023, AJ, 165, 124

Allende Prieto, C., Cooper, A. P., Dey, A., et al. 2020, Research Notes of the American Astronomical Society, 4, 188

Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33

Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ, 156, 123

Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, ApJ, 935, 167

Bailey et al. 2023, in preparation

Busca, N., & Balland, C. 2018, arXiv e-prints, arXiv:1808.09955

Chaussidon, E., Yèche, C., Palanque-Delabrouille, N., et al. 2023, ApJ, 944, 107

Cooper, A. P., Koposov, S. E., Allende Prieto, C., et al. 2022, arXiv e-prints, arXiv:2208.08514

Cutri, R. M., Wright, E. L., Conrow, T., et al. 2013, Explanatory Supplement to the AllWISE Data Release Products, Tech. rep.

DESI Collaboration, Aghamousa, A., Aguilar, J., et al. 2016a, arXiv e-prints, arXiv:1611.00036

—. 2016b, arXiv e-prints, arXiv:1611.00037

DESI Collaboration, Abareshi, B., Aguilar, J., et al. 2022, AJ, 164, 207
DESI Collaboration, Adame, A. G., Aguilar, J., et al. 2023a, arXiv e-prints, arXiv:2306.06307
—. 2023b, arXiv e-prints, arXiv:2306.06308
DESI collaboration et al. 2023, in preparation
Dey, A., Rabinowitz, D., Karcher, A., et al. 2016, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9908, Ground-based and Airborne Instrumentation for Astronomy VI, ed. C. J. Evans, L. Simard, & H. Takami, 99082C
Dey, A., Schlegel, D. J., Lang, D., et al. 2019, AJ, 157, 168
Farr, J., Font-Ribera, A., & Pontzen, A. 2020, JCAP, 2020, 015
Flaugher, B., Diehl, H. T., Honscheid, K., et al. 2015, AJ, 150, 150
Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, A&A, 595, A1
Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759
Green et al. 2023, in preparation
Guy, J., Bailey, S., Kremin, A., et al. 2023, AJ, 165, 144
Hahn, C., Wilson, M. J., Ruiz-Macias, O., et al. 2022, arXiv e-prints, arXiv:2208.08512
Hardin, R., Sloane, N., & Smith, W. 2000, Published electronically at http://www. research. att. com/njas/icosahedral. codes
Kirkby et al. 2023, in preparation
Lan, T.-W., Tojeiro, R., Armengaud, E., et al. 2023, ApJ, 943, 68
Levi, M., Bebek, C., Beers, T., et al. 2013, arXiv e-prints, arXiv:1308.0847
Meisner, A. M., Lang, D., & Schlegel, D. J. 2018, Research Notes of the American Astronomical Society, 2, 1
Miller, T. N., Doel, P., Gutierrez, G., et al. 2023, arXiv e-prints, arXiv:2306.06310
Moustakas et al. 2023, in preparation
Myers, A. D., Moustakas, J., Bailey, S., et al. 2023, AJ, 165, 50
Noll, S., Kausch, W., Barden, M., et al. 2012, A&A, 543, A92
Palmese, A., BenZvi, S., Bailey, S., et al. 2021, GRB Coordinates Network, 30923, 1
Patat, F. 2008, A&A, 481, 575
Raichoor, A., Eisenstein, D. J., Karim, T., et al. 2020, Research Notes of the American Astronomical Society, 4, 180
Raichoor, A., Moustakas, J., Newman, J. A., et al. 2023, AJ, 165, 126
Raichoor et al. 2023, in preparation
Ruiz-Macias, O., Zarrouk, P., Cole, S., et al. 2020, Research Notes of the American Astronomical Society, 4, 187
Schlafly, E. F., & Finkbeiner, D. P. 2011, ApJ, 737, 103
Schlafly, E. F., Meisner, A. M., & Green, G. M. 2019, ApJS, 240, 30
Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, ApJ, 500, 525
Silber, J. H., Fagrelius, P., Fanning, K., et al. 2023, AJ, 165, 9
Tie, S. S., Kirkby, D., Martini, P., et al. 2020, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 11447, Ground-based and Airborne Instrumentation for Astronomy VIII, ed. C. J. Evans, J. J. Bryant, & K. Motohara, 1144785
Walker, M. F. 1988, PASP, 100, 496
Williams, G. G., Olszewski, E., Lesser, M. P., & Burge, J. H. 2004, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 5492, Ground-based Instrumentation for Astronomy, ed. A. F. M. Moorwood & M. Iye, 787–798
Yang, J., Fan, X., Gupta, A., et al. 2023, arXiv e-prints, arXiv:2302.01777
Yèche, C., Palanque-Delabrouille, N., Claveau, C.-A., et al. 2020, Research Notes of the American Astronomical Society, 4, 179
Zhou, R., Newman, J. A., Dawson, K. S., et al. 2020, Research Notes of the American Astronomical Society, 4, 181
Zhou, R., Dey, B., Newman, J. A., et al. 2023, AJ, 165, 58
Zou, H., Zhou, X., Fan, X., et al. 2017, PASP, 129, 064101