The Astronomical Journal, 168:168 (13pp), 2024 October https://doi.org/10.3847/1538-3881/ad6cdb

SPRITEly: Time-domain Millimeter Interferometry at the Owens Valley Radio Observatory

Nitika Yadlapalli Yurk[1,2], Vikram Ravi[1,3], Mark W. Hodges[3], James W. Lamb[3], Richard Hobbs[3], David P. Woody[3], and Anthony C. S. Readhead[1,3]
[1] Cahill Center for Astronomy and Astrophysics, MC 249-17, California Institute of Technology, Pasadena, CA 91125, USA
[2] Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 91109, USA
[3] Owens Valley Radio Observatory, California Institute of Technology, Big Pine, CA, 93513, USA
Received 2023 November 14; revised 2024 August 2; accepted 2024 August 6; published 2024 September 20

Abstract

Though the time-domain millimeter sky is yet to be well characterized, the scarcity of millimeter observing resources in the world at present hampers progress toward it. In efforts to bolster the exploration of millimeter transients, we present the Stokes Polarization Radio Interferometer for Time-Domain Experiments (SPRITEly). Located at the Owens Valley Radio Observatory, SPRITEly is currently deployed as a two-element short-baseline 90 GHz interferometer uniquely focused on monitoring bright variable millimeter-continuum sources. We leverage two existing 10.4 m antennas and their existing receiver systems to begin, but we make significant upgrades to the back-end system during the commissioning process. With the ability to achieve rms noise of a few mJy, we plan to monitor known variable sources along with new nearby transients detected from optical surveys at high cadence, with the goal of producing well-sampled light curves. Interpreting these data in conjunction with multiwavelength observations stands to provide insight into the physical properties of the sources that produce transient millimeter emission. We present commissioning and early-science observations that demonstrate the performance of the instrument, including observations of the flaring BL Lac object S2 0109+22 and a periastron passage of the binary T Tauri system DQ Tau.

Unified Astronomy Thesaurus concepts: Radio interferometry (1346); Time domain astronomy (2109); Radio transient sources (2008); Astronomical instrumentation (799)

1. Introduction

The recently discovered prevalence and frequency of millimeter transients have generated a need for observing resources dedicated to monitoring and follow-up of interesting sources. For example, transient search pipelines applied to observations from the South Pole Telescope 3G instrument (S. Guns et al. 2021) and the Atacama Cosmology Telescope (S. Naess et al. 2021) yielded a sample of 18 events on timescales of days to weeks, surprisingly dominated by Galactic stars. These discoveries only scratch the surface of what can be accomplished with wide-field millimeter surveys, described further in N. Yadlapalli et al. (2022) and T. Eftekhari et al. (2022).

In addition to wide-field surveys, it is also important to conduct targeted studies that aim to constrain the rates and luminosities of millimeter transients from different classes of object. Common origins of extragalactic millimeter transient emission include synchrotron sources such as reverse shock emission from gamma-ray bursts (A. de Ugarte Postigo et al. 2012; T. Laskar et al. 2013), jet emission from some tidal disruption events (TDEs; B. A. Zauderer et al. 2011; Q. Yuan et al. 2016; I. Andreoni et al. 2022), fast blue optical transients (A. Y. Q. Ho et al. 2019, 2022), and interacting supernovae (Y. Cao et al. 2013; A. Horesh et al. 2013). Galactic accreting compact-object sources, such as black hole X-ray binaries (A. J. Tetarenko et al. 2017), also produce millimeter synchrotron emission. Recent detections of millimeter flares from nearby active stars have revealed yet another frequent source of transient millimeter emission needing deeper study (M. A. MacGregor et al. 2018; A. M. MacGregor et al. 2020; K. Burton et al. 2022). A few detections of giant millimeter flares from young stellar objects (e.g., G. C. Bower et al. 2003; J. Vargas-Gonzalez et al. 2023) also await further systematic characterization. Rapidly varying emission and short-duration flares at millimeter wavelengths are also seen in blazars. Here the axis of the relativistic jet is aligned close to the line of sight and the observed emission is strongly boosted by relativistic beaming, leading to fluctuations on timescales of days to years (see, e.g., L. Fuhrmann et al. 2016). In general, transient millimeter emission probes the physical extremes of the sources: the earliest stages of the evolution of jets and outflows, environmental conditions closest to the central engines, and the limits of magnetization and particle acceleration.

Even as millimeter observing resources become more abundant, through both wide-field surveys and highly sensitive facilities such as the Atacama Large Millimeter/submillimeter Array (ALMA), there still remains a need for observing resources dedicated to conducting studies of transients (e.g., K. D. Alexander et al. 2017). To fill this gap, we introduce the Stokes Polarization Radio Interferometer for Time-domain Experiments (SPRITEly), located at the Owens Valley Radio Observatory (OVRO), which will solely focus on observing transient and variable sources of millimeter emission. We describe an initial two-antenna deployment of SPRITEly and present commissioning and early-science observations that demonstrate the feasibility of the concept. We focus in particular on documenting several aspects of antenna and receiver bring-up for future users of SPRITEly deployments.

Figure 1. The two 10.4 m Leighton antennas, located at OVRO, that are used for the initial SPRITEly deployment. The receivers are housed within the sidecabs at the Nasmyth focus of each antenna, located nearer to the viewer. The smaller dishes are part of the 110-antenna Deep Synoptic Array. In the background is the Mayer control building, which houses the final steps of the signal processing chain and the data processing computers for the telescopes.

We describe the instrument in Section 2 and outline the observing and data-analysis strategy in Section 3. Commissioning and early-science observations are presented in Section 4. We discuss the future of the SPRITEly concept in Section 5 and conclude in Section 6.

2. Single Baseline Deployment Overview

2.1. Antennas

SPRITEly currently utilizes two Leighton 10.4 m antennas (R. B. Leighton 1977), shown in Figure 1. Built in the late 1970s, these antennas were designed to perform some of the earliest millimeter aperture-synthesis observations. Surface accuracies better than 50 μm delivered excellent response down to wavelengths as small as 1 mm; SPRITEly, however, will primarily utilize the 3 mm observing band. Each antenna was built with a sidecab, in which the receivers and supporting electronics are contained. Observations of compact sources were used to derive an antenna aperture efficiency of $\sim$0.5 (M. C. Wright et al. 2009).

Most recently, the Leighton antennas used here were a key component of the Combined Array for Research in Millimeter-wave Astronomy (CARMA; D. P. Woody et al. 2004). In 2015, CARMA was decommissioned and the dishes were moved back to OVRO. Five antennas were available to choose from to make up SPRITEly—the two we opted for comprised the shortest possible baseline of 24 m. This decision was motivated in part by minimizing atmospheric phase noise. As shown in C. Carilli & M. Holdaway (1997) for tests done on the Karl G. Jansky Very Large Array (VLA), baseline length correlates with the amount of rms phase noise present in observations. As the effect of atmospheric turbulence scales inversely with wavelength and SPRITEly will not make images that require high spatial resolution, the shortest baseline was the most prudent choice.

2.1.1. Monitor and Control

All communications for monitor and control of the Leighton antennas are done using a controller-area-network bus (CANbus) message system originally developed for CARMA (D. P. Woody et al. 2007). Each antenna sidecab houses its own computer that manages all incoming and outgoing CAN messages. Most monitor points are sampled at 0.5 s intervals, with the exception of encoder values, which are read out every 0.1 s. Communication with the CAN modules for SPRITEly occurs in a few different ways.

The antennas are driven using a graphical control system developed at OVRO for the Carbon Monoxide (CO) Mapping Array Project (COMAP; J. W. Lamb et al. 2022). This control system communicates with the modules controlling the antenna drives and other basic mechanics of the antenna. It also allows the user to check current azimuths and elevations of all sources in the source catalog, monitor current weather conditions, and run observing schedules. An additional feature of this control system is its storage of monitor data in an archive. Information for specific monitor points can be retrieved for arbitrary durations of time to be used for downstream analysis of antenna behavior, such as measuring coefficients of the pointing model as will be described in Section 2.1.2. The receiving system is largely controlled using a LabVIEW program, originally developed for use at CARMA, that interfaces with CAN modules communicating with receiver components. These monitor points include values such as a measure of the received power, attenuation levels at various stages, the receiver temperature, etc. Usage of this system for tuning and calibrating the receiver is described in Section 2.2.

2.1.2. Pointing

The procedure for determining pointing for the antennas is well established and consists of three steps: determining the tiltmeter offsets, measuring mount-dependent pointing coefficients with an optical camera, and determining the radio collimation offsets using radio measurements.

Tilts of the azimuth axis are tracked with a dual-axis tiltmeter on the azimuth axis and applied in the pointing model in real time. The two axes have small zero offsets that have to be updated on a monthly basis. These are determined by recording the tiltmeter outputs at eight azimuth points. The offset of sinusoidal curves fitted to these points gives the required tiltmeter zero offsets, which are stored in the control system.

Next, the parameters for the pointing model are determined. A seven-parameter model gives the corrections ΔAZ and ΔEL to convert the azimuth and elevation encoder readings, AZ and EL, to the true azimuth and elevation according to the following equations:

$$\begin{aligned}\Delta\mathrm{AZ} = {} & A_1 + M_1\cos(\mathrm{EL}) + M_3\sin(\mathrm{EL}) \\ & - M_4\sin(\mathrm{AZ})\sin(\mathrm{EL}) - M_5\cos(\mathrm{AZ})\sin(\mathrm{EL}) \quad (1)\end{aligned}$$

$$\begin{aligned}\Delta\mathrm{EL} = {} & -M_4\cos(\mathrm{AZ}) + M_5\sin(\mathrm{AZ}) \\ & + A_2 + A_3\cos(\mathrm{EL}). \quad (2)\end{aligned}$$

An optical telescope is permanently mounted on the primary backing structure, looking through a hole in the primary mirror. With the control system, the observer can image a star using a frame grabber and adjust offsets to center the star. These offsets are recorded in the archive for analysis. At least 10 (preferably 20 or more) stars are measured across the sky, and a least-

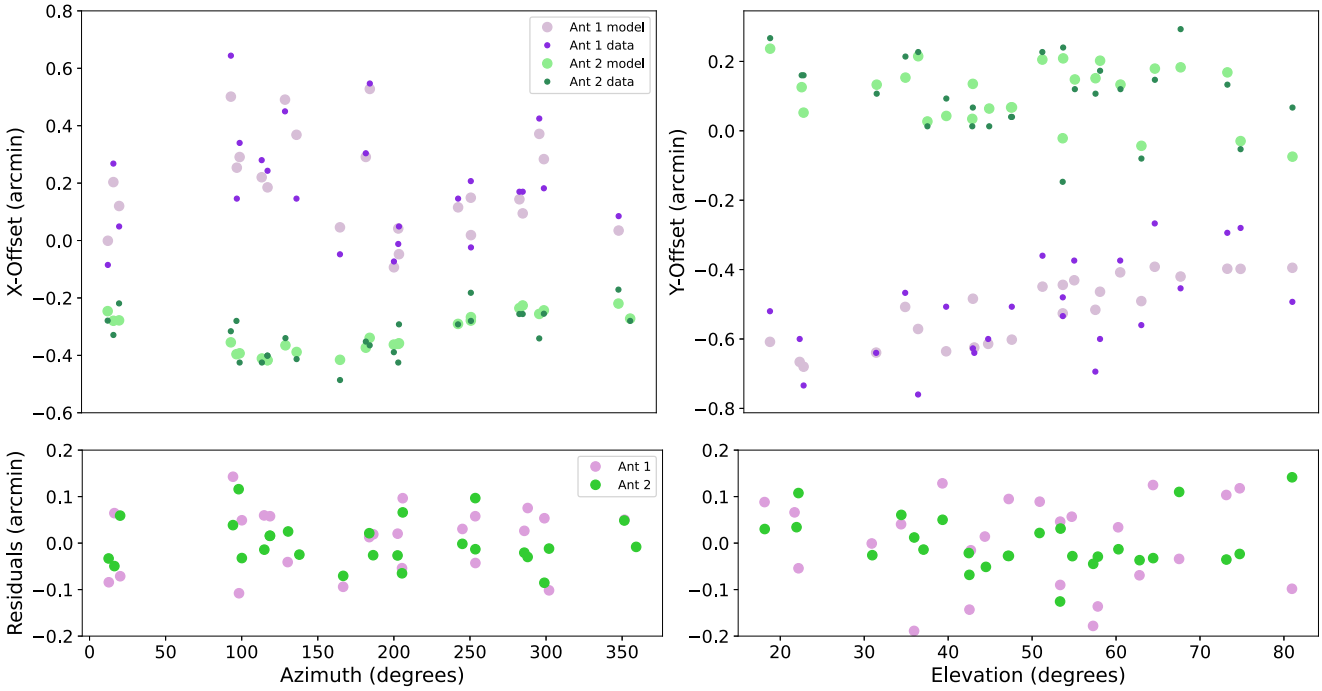

Figure 2. Example of data collected on 23 stars for calculation of the optical pointing model coefficients. Top: the smaller, darker markers shown are the true locations and offsets of the stars measured with each antenna. The larger, lighter markers show the results from applying pointing coefficients that best reproduce the observed offsets. Offsets are shown in the antenna-specific X and Y coordinates for each SPRITEly antenna. Bottom: residuals for the fit shown in the top panel, showing that there is no directional dependence to the fidelity of the model fitting.

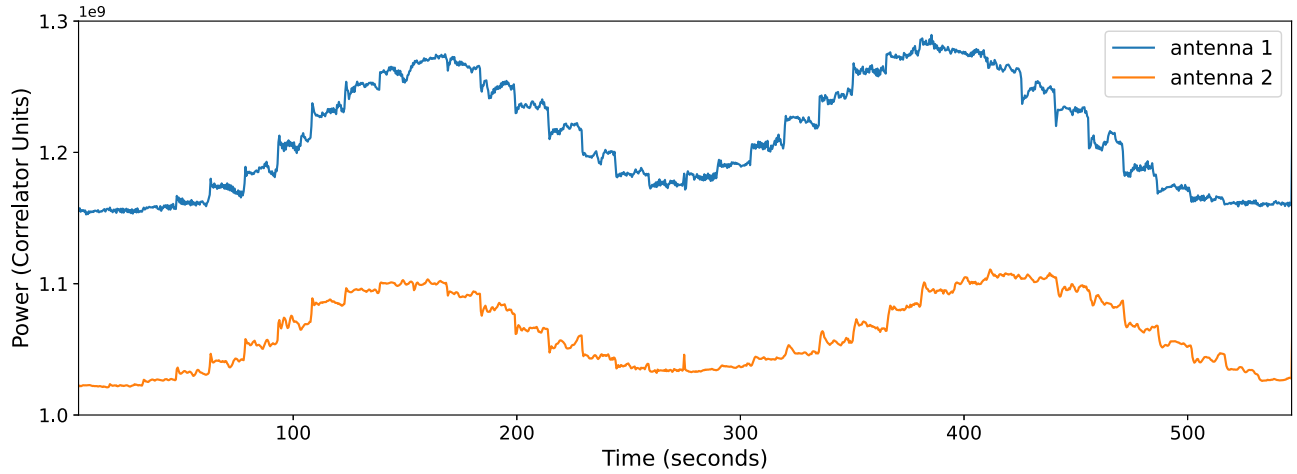

Figure 3. Example of data collected on Jupiter for single-dish radio pointing measurements. Light curves are shown for both SPRITEly antennas from scans in both the elevation and cross-elevation directions across the disk of the planet, which are represented by the first and second portions of the light curves, respectively. This scan measured total powers at offsets of up to $\pm 90''$ at $10''$ increments.

squares fit is used to determine all of the coefficients in Equations (1) and (2). We show an example of this in Figure 2, where we plot both x- and y-offset data taken on a sample of stars and the predicted offset values calculated using the best-fit coefficients.

The M terms are common to the optical and radio pointing models and represent pointing error contributions from the telescope mount. M_1 and M_2 represent azimuth and elevation encoder zero terms, while M_3, M_4, and M_5 represent tilt terms. Once they are found with the procedure described above and entered into the control system, only a few radio measurements are needed to determine the A aperture terms. A_1 and A_2 are collimation terms that capture the divergence between the radio and optical axes in two directions, while A_3 is an elevation-dependent sag term. We note that in Equation (2) we set $M_2 = 0$ and absorb the elevation encoder zero correction into the calculation of A_2.

Initial radio pointing offsets are made on bright radio sources (such as Saturn, Jupiter, or the Moon), even if they are resolved. Stepping in elevation or cross-elevation and recording the correlator total power (see Section 2.3 below) yields light curves such as those shown in Figure 3. The pointing recorded at peak brightness yields the respective offsets. Interferometric measurements with the two antennas yield more refined values. The elevation and cross-elevation collimation terms need only a few measurements, while the sag term needs a larger number over a wide elevation range. However, since the sag term is small and stable, it rarely needs to be revised.

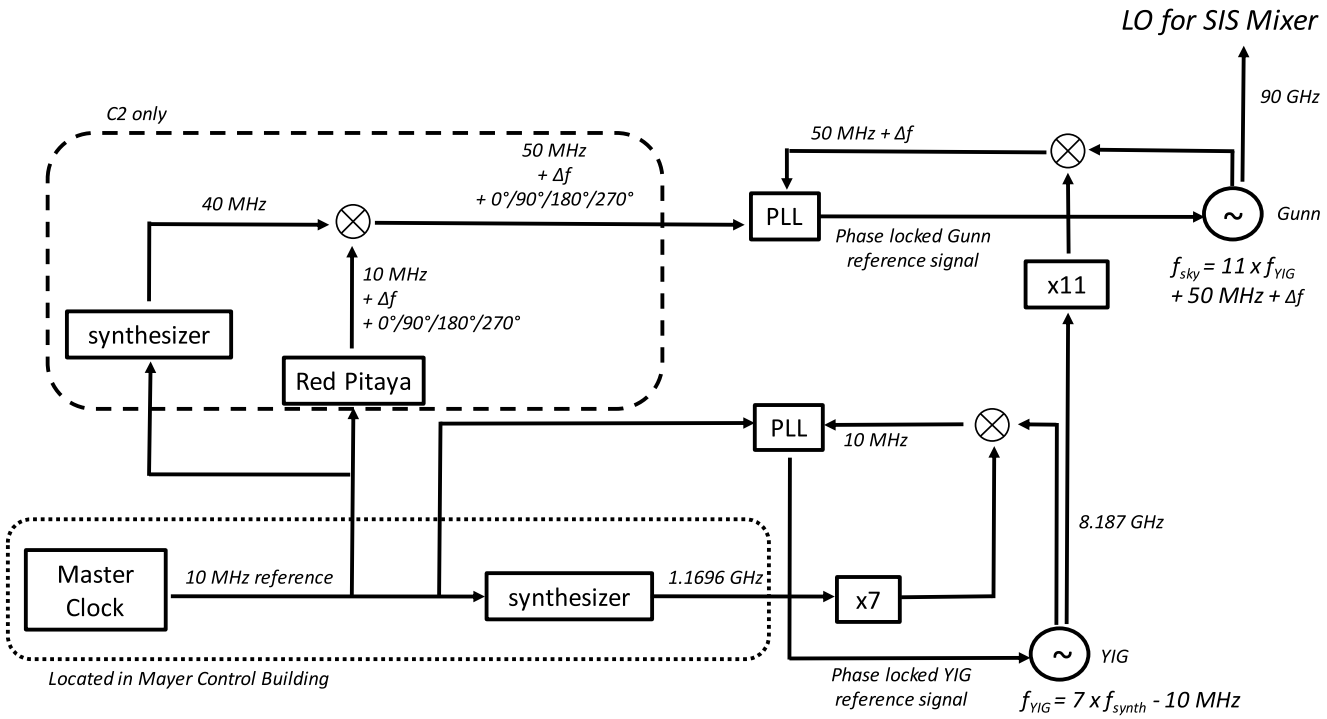

Figure 4. Diagram showing how the LO signal for the SIS mixer is generated. The components within the dotted box are located in the basement of the Mayer control building and are one set of components shared between both antennas. Signals leaving the Mayer control building are sent out over optical fiber and converted back to RF within each of the antennas. The components within the dashed line represent the system in the sidecab of SPRITEly Antenna 2 only. In Antenna 1, the dashed box is replaced by a single synthesizer that generates a 50 MHz signal.

The procedures described above have resulted in a blind radio pointing accuracy that is better than $\pm 10''$ (peak to peak), as determined through interferometric pointing measurements. Given a Gaussian primary half-power beamwidth at 90 GHz of $81''$, this corresponds to a maximum intensity loss of $<4\%$. Occasional pointing checks during observing remove small residual effects due to temperature and weather. Note that tilts are automatically compensated for by the tilt meters. Both optical and radio aperture coefficients are stored and can be selected as needed.

2.2. Receiving System

SPRITEly mostly utilizes the original millimeter receiving system designed for use in CARMA. The receivers are located in the sidecabs of the antennas at a Nasmyth focus. They support both 3 mm and 1 mm observations and employ superconducting–insulating–superconducting (SIS) mixers providing double-sideband conversion. For the two-antenna deployment, only the single-linear-polarization 3 mm channel was commissioned. The receivers are cryogenically cooled with a first-stage radiation shield cooled to around 70 K and a second stage cooling down to around 4.5 K as required for the superconducting components. While much of the local oscillator (LO) signal generation procedure remains the same, major modifications were made to the lobe rotation and phase switching system.

2.2.1. LO Signal Generation

The LO signal is mixed with the incoming astronomical signals in the SIS mixer to perform the first stage of downconversion for SPRITEly. The frequency of the LO signal determines the center frequency of our observing band, and high phase stability of the LO signal is required for achieving good signal-to-noise ratio in the final cross-correlation products. Stable LO phase is achieved using multiple steps of multiplication and phase locking from a lower-frequency reference signal, as shown in Figure 4. A tone from a low-phase-noise synthesizer at 1.1696 GHz in the Mayer control building is sent on fiber to the antenna, where a diode multiplier generates the seventh harmonic at 8.1872 GHz. This is mixed with the output of a yittrium–iron–garnet (YIG) tuned oscillator at 8.1772 GHz, yielding a 10 MHz difference that is servoed to a 10 MHz reference using a phase-locked loop (PLL) controlling the YIG tuning coil. A harmonic mixer combines the YIG output and that of a bias-tuned Gunn oscillator operating at 90 GHz producing a 50 MHz difference. A second PLL keeps the Gunn oscillator locked to a 50 MHz reference. The Gunn oscillator provides the LO for the SIS mixer, and, as described below, the 50 MHz is used for phase switching and lobe rotation. The 10 MHz from the observatory master clock containing a rubidium oscillator locked to GPS is used as the reference for the Mayer building synthesizer and sent to the antenna on an optical link. As explained in the following section, it is also used for the 50 MHz reference, which was developed specifically for SPRITEly. It is crucial that both synthesizers responsible for the 50 MHz signal be referenced to the same 10 MHz signal to allow for precise control of the phase difference between the two LO signals.

No Doppler tracking is applied for Earth rotation or source velocity since it has negligible effect on the continuum data. The YIG and Gunn oscillator tuning is done by the respective CAN modules after a frequency is issued through the LabVIEW interface, which also monitors the lock state and signal levels.

2.2.2. Fringe Tracking

The fundamental output of an interferometer is the cross-product of the signals from a pair of antennas. This product, known as a visibility, is a complex number representing a Fourier component of the sky brightness over the antenna beam parallel to the baseline. The visibility can be imagined as a pair of sine and cosine fringes on the sky. Since the sky rotates through the fringes at the sidereal rate, any reasonable data recording rate will lead to the antennas averaging together many fringe cycles, integrating away the signal. However, given that the fringe frequency is a known function of the observing frequency and projected baseline length, the relative phases between two LOs on a single baseline can be modified with time to match the fringe frequency. This allows the fringes to be frozen relative to the baseline, enabling longer averaging times on source.

Referring to Equation (6.11) of A. R. Thompson et al. (2017), the phase difference across the upper sideband (USB) of a single baseline, comprising antennas m and n, after a single stage of downconversion is given as follows (assuming that both visibility phase and gain phase are zero). This would emulate the phase difference between signals exiting the SIS mixers in the two SPRITEly antennas

$$\phi_{mn} = 2\pi\nu_{\rm LO}\tau_g + (\theta_m - \theta_n). \tag{3}$$

Here ϕ_{mn} is known as the fringe phase, $\nu_{\rm LO}$ is the LO frequency, and τ_g is the geometric delay between the two antennas (the difference in times of the arrival of the same signal to antenna m compared to n), while θ_m and θ_n refer to the phase of the LOs in antennas m and n. To maintain a steady fringe phase with time (also known as "fringe stopping") in the final visibility, we need to adjust the phase difference between the LOs at the same rate that the geometric delay is changing, as follows:

$$\begin{aligned} 2\pi\nu_{\rm LO}\frac{d\tau_g}{dt} + \frac{d}{dt}(\theta_m - \theta_n) &= 0 \\ 2\pi\nu_{\rm LO}\frac{b\sin\theta}{c}\frac{d\theta}{dt} &= \frac{d}{dt}(\theta_m - \theta_n). \end{aligned} \tag{4}$$

In Equation (4), b is the baseline length and θ is the angle between the baseline vector and the source position vector. For observations at 90 GHz with a 24 m baseline, the maximum speed of the fringes (known as the natural fringe frequency) will be $\sim$0.5 Hz. Practically, this means that we would need to record data faster than once per second to minimize loss.

To stop the fringes, we control the phase difference between the LOs of the antennas by keeping the 50 MHz reference signal fixed in one antenna and ramping the phase of the 50 MHz of the other as shown in Figure 4. To generate the 50 MHz reference signal for the Gunn PLL, we mix a phase-adjusted 10 MHz signal with a 40 MHz signal locked to the reference 10 MHz. The phase adjustments to the 10 MHz signal include adjustments for both fringe tracking and quadrature phase switching, required for sideband separation of the correlations (discussed in Section 3.1).

We implement the phase switching of the 10 MHz signal on a small, inexpensive device produced by Red Pitaya (STEMLab 125-10; hereafter "Red Pitaya"). It features a 32-bit field-programmable gate array (FPGA) clocked at 125 MHz, as well as two 10-bit analog-to-digital converters (ADCs) and digital-to-analog converters (DACs), sufficient for a simple operation such as adjusting the phase of an input signal. The firmware described in this section was written using tools developed by the Collaboration for Astronomy Signal Processing and Electronics Research (CASPER) collaboration (J. Hickish et al. 2016). Both the Red Pitaya and the synthesizer generating the 40 MHz signal are located in an electronics rack within the antenna sidecab.

To start, we write the phase-shifted signal as a linear combination of the cosine and sine components of the original signal, as below, where ϕ is the phase shift required to track the fringe rate of the source and $\theta_{\rm switch}$ is the phase switching term:

$$\cos(x - \phi + \theta_{\rm switch}) = A\cos(x) + B\sin(x). \tag{5}$$

Let us assume that input 10 MHz signal at the ADC represents the cosine component, $\cos(x)$. To derive the sine component, we delay the incoming signal by one clock cycle of FPGA, which translates to a $\phi' = 0.5$ rad shift of 10 MHz signal. Then, the firmware linearly combines the original and delayed signal to retrieve the sine component

$$\begin{aligned} \sin(x) &= \frac{-\cos(x + \phi') + \cos(x)\cos(\phi')}{\sin(\phi')} \\ &= -2.086\cos(x + \phi') + 1.830\cos(x). \end{aligned} \tag{6}$$

The coefficients A and B are defined as

$$A = \cos(\phi + \theta_{\rm switch}) \tag{7}$$

$$B = \sin(\phi + \theta_{\rm switch}), \tag{8}$$

where the value of $\theta_{\rm switch}$ will cycle between $0°/90°/180°/270°$ at each data readout cycle. We define $d\phi/dt$ through a software register, and the firmware will increment ϕ at each clock cycle and calculate the new values of A and B accordingly. The firmware then multiplies with the cosine and sine components of the input signal with coefficients A and B, respectively, and sends the output on the DAC. A block diagram, as well as the original Simulink and compiled firmware files, can be accessed at https://github.com/nitikayad96/SPRITE/blob/main/firmware.

Given that we are incrementing the LO phase at each FPGA clock cycle, we can only adjust the phase by integer multiples of fringe cycles per clock cycle. As our FPGA resolution is 32 bits and 1 bit will be used for determining the sign of the fringe rate, the minimum value of fringe cycles per clock cycle we can use is $1/2^{31} = 4.66 \times 10^{-10}$. For a 125 MHz clock, this translates to a minimum incremental fringe rate of $4.66 \times 10^{-10} \times 125\ {\rm MHz} = 0.06$ Hz. Thus, the maximum residual fringe left over in the correlations after fringe tracking will be 0.06 Hz.

2.2.3. First Stage of Downconversion

The original CARMA intermediate-frequency (IF) signal chain supports 8 GHz of bandwidth from 1 to 9 GHz. As the SIS mixer is not sideband separating, both the USB and lower sideband (LSB) are mixed into the IF. The receiving system can be tuned to a specific frequency based on the signal generated by the Gunn oscillator along with the bias voltage applied to the SIS mixer. This bias point is at the center of the first photon step below the superconductor gap voltage. For more information, see the review by T. G. Phillips & D. P. Woody (1982). The tuning of the receiver is done autonomously by the CAN modules on receipt of a frequency message from the LabVIEW program.

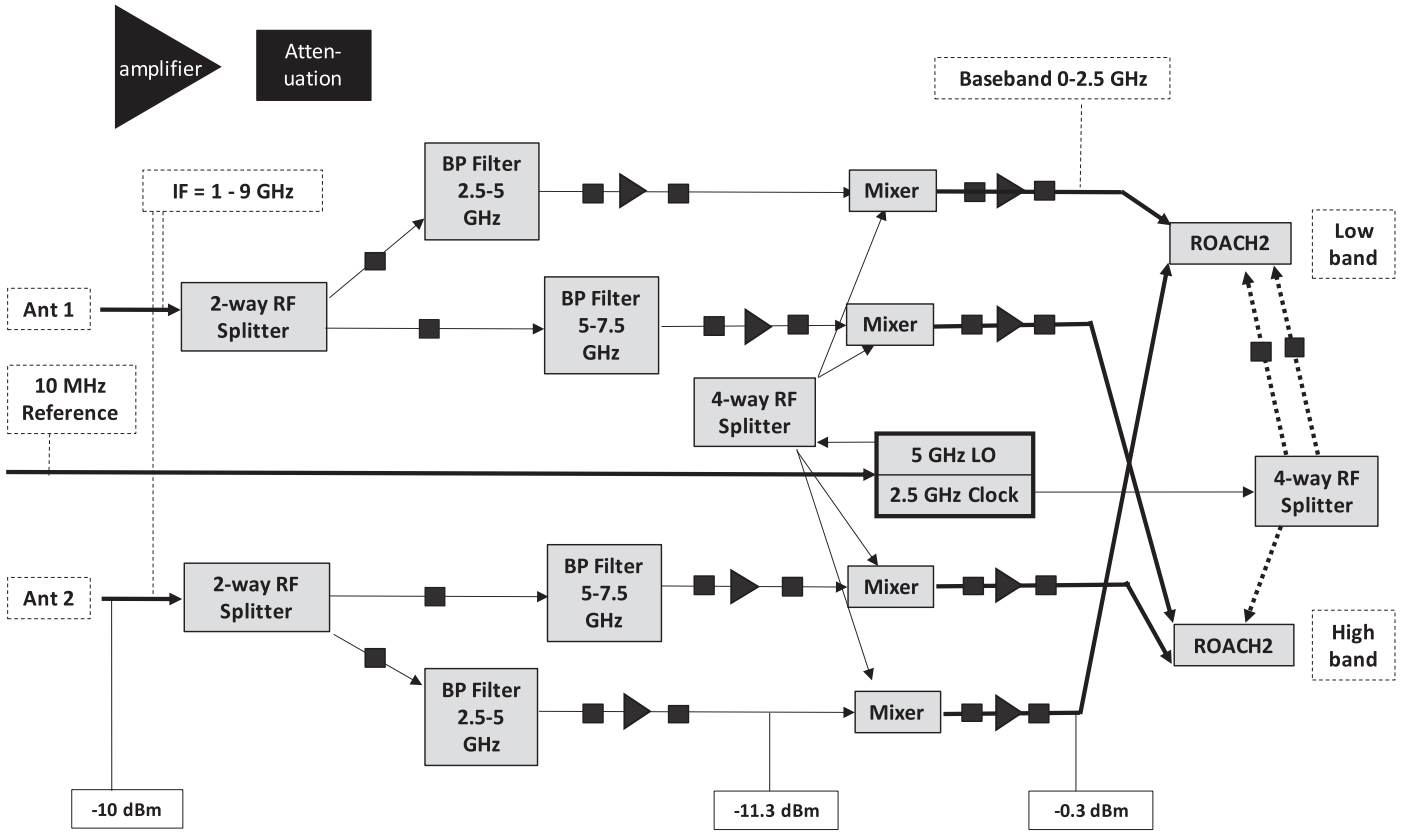

Figure 5. Diagram of the SPRITEly second-stage downconversion system showing how IF signals from both antennas are split into high- and low-band components to be fed to their respective ROACH2 boards for correlation. The thick arrows represent input signals to the downconverter and output signals to the ROACH2. The thick dotted lines indicate clocking signals. One of our ROACH2 boards contains two inputs for two clock signals, whereas the other ROACH2 board contains only one clock signal input and splits the signal internally.

For low-loss transfer of the signal, the IF signal is sent from the sidecab on a radio frequency (RF) over fiber link to the Mayer Control Building basement, where it is fed to the second-stage downconversion module.

2.2.4. Second Stage of Downconversion

Although 8 GHz of IF bandwidth is available, the current correlator is limited to sampling a maximum of 5 GHz of bandwidth (4 GHz of bandwidth in practice; see Section 2.3). Given this, the downconversion system designed for SPRITEly selects the central 5 GHz of the IF, 2.5–7.5 GHz, to convert to baseband, as shown in Figure 5. As will be discussed further in Section 2.3, our correlator handles the upper and lower halves of the selected IF band separately.

The input from each antenna is first split and fed into custom bandpass filters, one selecting 2.5–5 GHz, purchased from Reactel, and the other selecting 5–7.5 GHz, purchased from K&L Microwave. Using a Mini-Circuits ZX05-83-S+ mixer, these signals are each mixed with a 5 GHz LO signal generated by a Valon 5009 frequency synthesizer, referenced to the 10 MHz from the master clock. Baseband signals from the 2.5–5 GHz signal path of each antenna are directed toward two inputs of one half of the correlator, while the signals from the 5–7.5 GHz path are directed toward the other half. Attenuators are placed throughout the system to mitigate reflection from ports of successive components, and amplifiers are added to ensure that the instrument signal entering the ROACH2 boards is close to 0 dBm.

2.3. Correlator

The SPRITEly correlator employs two ROACH2 boards (A. Parsons et al. 2008), one for each half of the utilized IF bandwidth, for sampling and cross-correlating incoming signals from the antennas. The Reconfigurable Open Architecture Computing Hardware (ROACH) design was developed by CASPER to serve as a fast and flexible FPGA-based solution to computing correlation products from N-antenna interferometers. Each ROACH2 contains two interleaved 2.5 Gsps ADCs, and the clock signals for the ADCs are generated by a Valon 5009 synthesizer. For our FPGA firmware, we use the FX (Fourier transform followed by cross-correlation) digital correlator design implemented in the Arcminute Micro-Kelvin Imager (J. Hickish et al. 2018). At present, a design that processes 2 GHz of bandwidth is implemented and commissioned, which outputs complex auto- and cross-spectra with 2048 channels.

We use Python code to initialize the ROACH2 boards and handle data collection from them. In the initialization script, we reset the software registers, calibrate the ADCs, and set up an integration time. We currently integrate over times of 0.4096 s, the fastest time allowed by the overhead required to read from the board. The integration is sufficiently fast given that the maximum amount of residual fringe left from the fringe tracker is 0.06 Hz, requiring a sampling period of only $\sim$8 s or better. We then synchronize the two ROACH2 boards using a signal of 1 pulse per second. By resetting the control registers on each board to the same pulse, we guarantee that the data are being

collected at the same time on each board. This is required for accurately capturing data in sync with the phase switching.

The data collection script handles writing observation information and the output correlations to an output hdf5 data file. At the beginning of each observation, the name and coordinates of the source are written to the output file. Based on the coordinates of the source, the required fringe tracking rate is calculated and subsequently written to both the hdf5 file and the relevant Red Pitaya software register. When the observation begins, a counter in the ROACH2 firmware increments at the end of each integration, and the Python code then writes out the complex correlations and time stamps for each integration. At the end of each integration cycle, the phase switching software register is switched to the next phase in the cycle. Data are captured until the elapsed time equals the specified observation time.

3. Observing Strategy and Data Analysis

Observing schedules for SPRITEly are constructed in blocks of 24 hr sequences, and observed sources fall into one of four categories: flux calibration, bandpass calibration, gain calibration, and science. For all observations, we measure the system temperature, $T_{\rm sys}$, every 5 minutes by moving an ambient temperature load in front of the receiver. A Python script running in parallel with the correlator code communicates with the CAN module to control the movement of the load.

We use two observations of Uranus per sequence to derive flux calibration measurements. We calculate the brightness temperature of Uranus at 90 GHz using the frequency-dependent model presented in M. Hasselfield et al. (2013) and the updated model coefficients presented in T. Louis et al. (2014). From there, we calculate the flux of Uranus at a given observation time using the radius measurements presented in B. A. Archinal et al. (2011) and the planet ephemeris calculations built into the `astropy` software package. Similar formulations could be created for other solar system objects to be used as flux calibrators as well, so long as these objects are sufficiently bright and have an angular size much smaller than the half-power beamwidth. Uranus was the easiest object to use that fit these criteria, and we find that the flux calibration is stable enough over the course of $\sim$1 day to solely rely on it for now. For bandpass calibration, we use one scan per sequence of either 3C 84 or 3C 273.

For a specified science target, 1 minute of gain calibration scans is taken in between 10-minute intervals of science target scans. The set of gain calibrator sources is a combination of ALMA-monitored calibrator sources with flux densities greater than 0.5 Jy and the list of bright calibrator sources used at CARMA. The sources and their locations are shown in Figure 6. For each science target, we assign the nearest source from this set to serve as the calibrator and allow a generous 90 s of slewing time between science and gain calibrator sources.

We refrain from observing sources during times when their nearest calibrator source is on the opposite side of the 0° azimuth line to avoid cable-wrapping issues. Additionally, though the minimum elevation the Leighton antennas can observe is 15°, we refrain from observing sources lower than 30°, due to high atmospheric optical depths at low elevations.

To carry out an observing sequence, a schedule file detailing at what time and for what duration to observe each source is written. This initial schedule need only contain bandpass and flux calibrator scans and science targets. A Python-based scheduling package written for SPRITEly then automatically tags all observations of Uranus as flux calibration scans and all observations of 3C 84 and 3C 273 as bandpass calibrators. The nearest gain calibrator is identified for each science target, and the scans are automatically added. This full schedule is then passed to the control system, which schedules movement of the antenna drives, and to a wrapper script for the correlator, which parses the schedule and synchronizes times of data collection.

3.1. Data Processing

All of the post-processing of the visibilities is done using custom Python 3.x routines. The first step in reducing data from a sequence is to perform scan-by-scan preprocessing. For each scan, we separate the sidebands, remove geometric delay effects, and then remove cable length delay effects.

As the SIS mixers used for SPRITEly are not sideband separating, we rely on the fact that each set of four visibilities read out from the correlator cycles through phases 0°, 90°, 180°, and 270° to separate the sidebands downstream from the receiver. One only needs one set of 90° switched visibilities to perform sideband separation, but the combination with 180° switching helps cancel cross talk introduced after the first mixer. Referring to Section 6.1.12 of A. R. Thompson et al. (2017), the complex response of a double-sideband system is the sum of the USB and LSB responses. Referring to Equation (9), which assumes that the gain amplitude is unity, ψ_u is the visibility phase of the USB and ψ_l is that of the LSB:

$$V_{\theta=0^\circ} = |V|[(\cos\psi_u + \cos\psi_l) + i(\sin\psi_u - \sin\psi_l)]. \quad (9)$$

Delaying the output of antenna 2 by 90° using the Red Pitaya, $\psi \rightarrow \psi - \pi/2$, the output visibility becomes

$$V_{\theta=90^\circ} = |V|[(\sin\psi_u + \sin\psi_l) + i(-\cos\psi_u + \cos\psi_l)]. \quad (10)$$

In our data processing code, we use the following relations to retrieve sideband-separated signals:

$$\begin{aligned} V_{\rm LSB} &= |V|(\sin\psi_l + i\cos\psi_l) \\ &= \frac{1}{2}[(\mathrm{Re}\{V_{\theta=90^\circ}\} - \mathrm{Im}\{V_{\theta=0^\circ}\}) \\ &\quad + i(\mathrm{Im}\{V_{\theta=90^\circ}\} + \mathrm{Re}\{V_{\theta=0^\circ}\})] \end{aligned} \quad (11)$$

$$\begin{aligned} V_{\rm USB} &= |V|(\sin\psi_u + i\cos\psi_u) \\ &= \frac{1}{2}[(\mathrm{Re}\{V_{\theta=90^\circ}\} + \mathrm{Im}\{V_{\theta=0^\circ}\}) \\ &\quad - i(\mathrm{Im}\{V_{\theta=90^\circ}\} - \mathrm{Re}\{V_{\theta=0^\circ}\})]. \end{aligned} \quad (12)$$

We can switch $V_{\theta=0^\circ} \rightarrow V_{\theta=180^\circ}$ and $V_{\theta=90^\circ} \rightarrow V_{\theta=270^\circ}$ in Equations (11) and (12) to retrieve separated sideband responses, $V'_{\rm LSB}$ and $V'_{\rm USB}$. The correlated signal component in V is equal to $-V'$, so $0.5(V - V')$ gives the sideband-separated signals; however, cross talk or common-mode signals after the first mixer do not have their phase inverted, so the difference cancels them out if they are sufficiently stable.

As each ROACH2 board is handling one half of the IF, we essentially have four bands of interest. We refer to each half of the IF as either high band or low band and refer to the sidebands within those as upper bands and lower bands. The

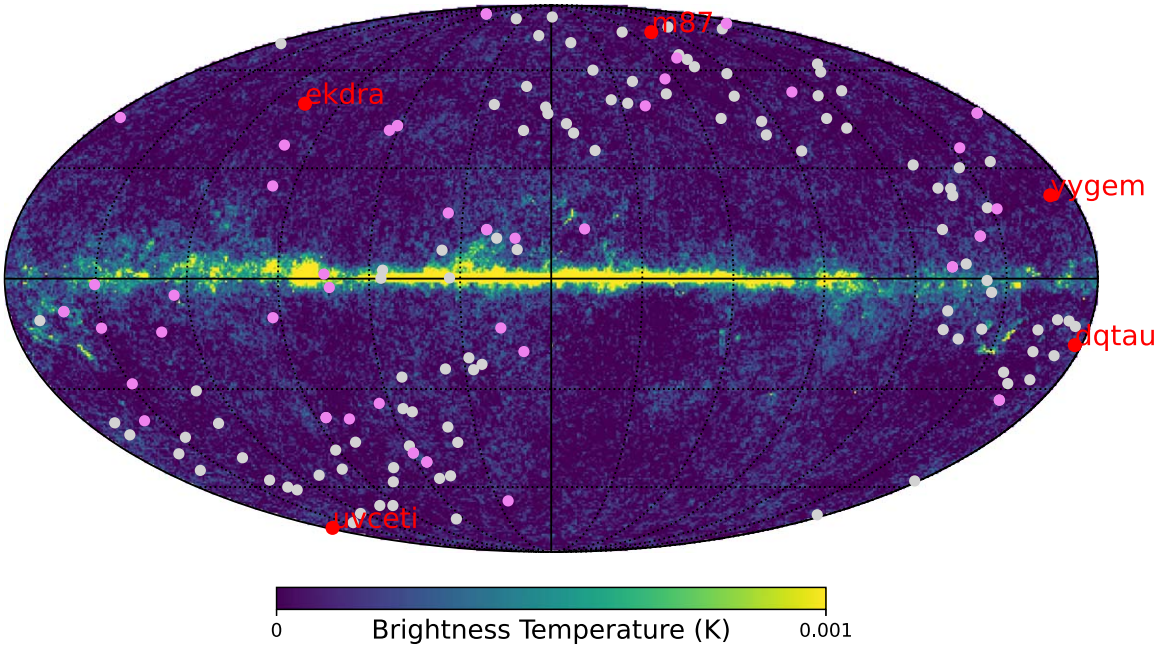

Figure 6. Locations of all sources allowed for use as gain calibrators for SPRITEly science observations overlaid on a Mollweide projection of a 100 GHz all-sky map (with the cosmic microwave background subtracted) created with observations from the Planck High Frequency Instrument (Planck Collaboration et al. 2020). The gray points represent sources taken from the ALMA calibrator list, while pink points represent sources used as calibrators for CARMA. The red points represent a few science targets of interest located at a variety of coordinates to demonstrate the broad availability of calibrator sources.

sky frequency ranges they span are as follows:

$$F_{\mathrm{HB},l} = 83 - 85\ \mathrm{GHz}$$
$$F_{\mathrm{LB},l} = 85 - 87\ \mathrm{GHz}$$
$$F_{\mathrm{LB},u} = 93 - 95\ \mathrm{GHz}$$
$$F_{\mathrm{HB},u} = 95 - 97\ \mathrm{GHz}.$$

For each of these four bands, we calculate model visibilities using the coordinates of the source and the time stamps recorded from the data collection code. We use CASAtoolkit (CASA Team et al. 2022) to calculate the (u, v, w) terms at each time, deriving phase variations due to the time-varying geometric delay. We then add a correction to these visibilities to simulate the effect of removing a fringe rate with the Red Pitaya. As the last step, we remove the effect of a fixed cable length delay from the visibilities. This value is calculated using the delay spectrum of a bright calibrator source and adjusted only on a monthly basis. Fixed delay changes also occur each time the ROACH2 boards are reinitialized, but these are absorbed into the bandpass solution, which stays fixed between scans and sequences. At this point, the visibility phase should be flat across all spectral channels and should be constant with time.

The final step of preprocessing the scans is applying antenna-dependent corrections to the bandpass amplitude based on measurements of the system temperature, T_{sys}. As mentioned above, we take measurements of the system temperature every 5 minutes using the ambient load. The ratio of the power from the receiver looking at the ambient temperature load to the power while observing the sky is used to derive T_{sys} referred to outside the atmosphere (M. L. Kutner & B. L. Ulich 1981; C. L. Carilli et al. 1999), which is likely accurate to ∼10%. Using archived monitoring measurements of the ambient temperature and assuming $T_{\mathrm{sky}} = T_{\mathrm{CMB}} = 2.73$ K, we calculate the system temperature referenced to the top of the atmosphere for each channel of each antenna using the following equations, where P denotes power measured in each state:

$$Y_\nu = \frac{P_{\nu,\mathrm{load}}}{P_{\nu,\mathrm{sky}}} \tag{13}$$

$$T_{\mathrm{sys},\nu} = \frac{T_{\mathrm{load}} - Y_\nu T_{\mathrm{sky}}}{Y_\nu - 1}. \tag{14}$$

The visibility cross-correlation spectra amplitudes are converted to noise temperature units using a factor $c(\nu) = \sqrt{T_1(\nu)T_2(\nu)}/\sqrt{A_1(\nu)A_2(\nu)}$ in kelvin per correlator unit. T_1 and T_2 are the system noise temperatures for the two antennas, and A_1 and A_2 are the time-averaged autocorrelations. We mask sections of data when the ambient temperature load was in the optical path, and we also mask frequency channels affected by correlator artifacts and aliasing. The effective remaining double-sideband bandwidth is 3 GHz. At this point, the time and elevation effects of the atmosphere are removed from the data, and the scan is fully preprocessed.

After all scans in a sequence are preprocessed, we move on to calibration of the whole sequence. To derive the bandpass solution for the sequence, we compute the time-average bandpass across the bandpass calibrator scan. We smooth the bandpass solution for the real and imaginary component separately using an implementation of a Savitzky–Golay filter within Python's `scipy.signal` package and then apply it to all scans in the sequence. Flux density calibration values are then derived from scans on Uranus using an average value of the visibility across the scan, after detrending a constant phase slope to account for any atmospheric or instrumental drifts.

Before applying the flux calibration solution, the complex gain variations within tracks of a single source must be calibrated. For each scan on a gain calibrator, a single value of complex gain is derived. The gains derived from the calibrator source are then normalized to unity, the science scans are then corrected for the measured gain-phase variations, and the flux

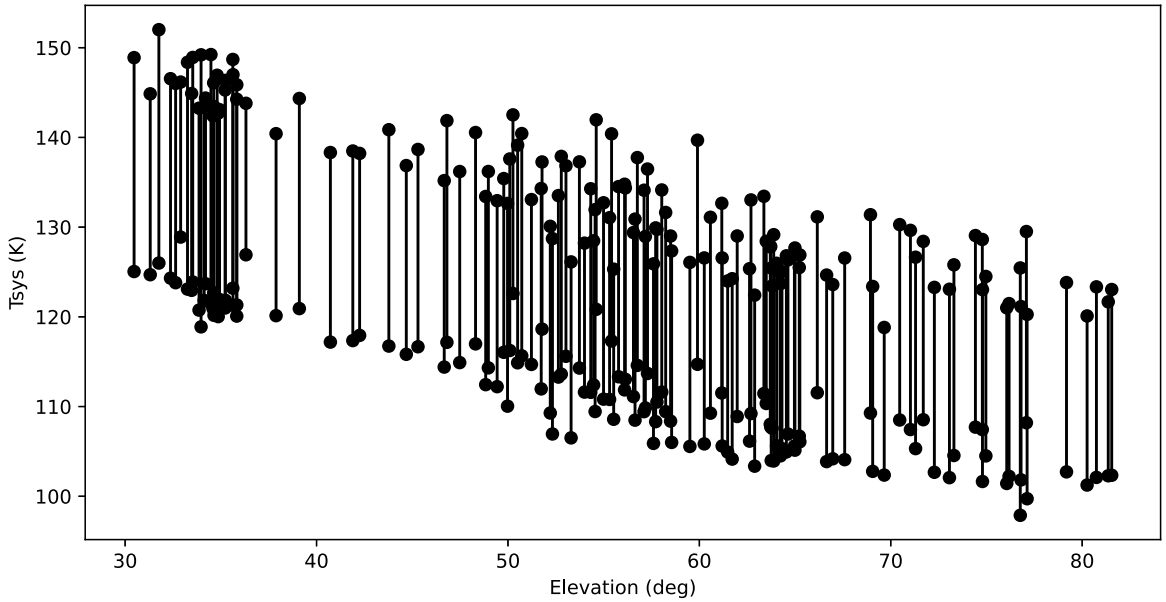

Figure 7. Measurements of elevation-dependent double-sideband system temperature of the SPRITEly baseline in reasonable weather conditions. Each measurement contains the $T_{\rm sys}$ of both the high band and low band, shown as circles connected by a line.

density solution is then applied. At this point, fully calibrated scans may be integrated down to achieve higher sensitivity.

Fully calibrated data enable the detection of time-variable emission on all timescales greater than the minimum integration time (0.4096 s at present). Several further analysis techniques can be applied to these light-curve data, ranging from the detection of intrascan transients (e.g., from active stars) to monitoring the evolution of sources over months to years. The simplicity of the flux density time-series data products means that a broad range of widely used techniques from multiwavelength time-domain astronomy can be applied. A key challenge, which is beyond the scope of this paper to fully constrain, will be the characterization of the time-variable noise properties of the data in different observing conditions. Detailed analysis and interpretation of the light-curve data will ultimately depend on the specifics of the objects they represent and the sensitivity required for them.

4. Commissioning Observations

4.1. T_{sys} Measurement and Sensitivity Prediction

We first present measurements of $T_{\rm sys}$ obtained during each scan of a 24 hr sequence and predict a limiting sensitivity of the current SPRITEly deployment. Figure 7 shows band-averaged $T_{\rm sys}$ measurements from a sequence obtained under reasonable weather conditions as a function of elevation. These double-sideband $T_{\rm sys}$ measurements are approximately 50% lower than the effective $T_{\rm sys}$ in each sideband after separation. A more accurate treatment requires an analysis of the gain ratio between the sidebands. Then, the predicted system-equivalent flux density can be derived as follows, where η is the aperture efficiency, A is the collecting area of a single dish, and k is Boltzmann's constant:

$$\mathrm{SEFD} = \frac{4kT_{\rm sys}}{\eta A\sqrt{2}} = 11{,}000\left(\frac{T_{\rm sys}}{120\ \mathrm{K}}\right)\mathrm{Jy}. \tag{15}$$

Thus, given our utilization of 3 GHz of IF bandwidth, SPRITEly should achieve $\sigma_S = \mathrm{SEFD}\sqrt{\Delta\nu} = 200\left(\frac{T_{\rm sys}}{120\ \mathrm{K}}\right)\mathrm{mJy\ s^{-1/2}}$ rms noise. In practice, complex gain calibration uncertainties will add tens of percent of noise for phase-referenced observations.

4.2. Verification of Sideband Separation

To verify proper timing accuracy and performance of the phase switching system, we conducted an observation of SiO maser emission from R Leo (M. Balister et al. 1977). The $J = 2-1$ rotational transition of SiO occurs at 86.243 GHz and thus should occur only in the LSB of our observations. In addition, we should also see a significant improvement in signal-to-noise ratio of the maser signal after incorporating 180° phase switching to remove the effects of bright cross talk. These results are shown in Figure 8. We observe that the maser emission is not present in the USB but is seen in the LSB at the correct frequency accounting for a slight redshift due to R Leo's motion relative to the observing reference frame. We also observe an ∼8 dB improvement in signal-to-noise ratio after subtracting (rather than adding) 180° phase-switched integrations. This improvement will be frequency dependent owing to frequency structure in the cross talk.

4.3. Baseline and Bandpass Solutions

Accurate bandpass and complex gain calibration of SPRITEly observing sequences requires a robust solution for the baseline geometry. The antenna positions are first measured to subcentimeter accuracy using measurements of antenna locations post-processed using the Canadian Spatial Reference System Precise Point Positioning program and a local global navigation satellite system base station. Refinement at the submillimeter level, required for 3 mm observations, was then done by fitting for errors in four 2 hr tracks on bright compact sources. In the presence of baseline errors, the final data after performing all of the calibration steps outlined in Section 3.1 will show a residual fringe, $e^{i\phi}$, with time and spectral dependence as shown in Equation (16). The terms Δb represent x, y, z components of the baseline error, while $s(t)$ represents the time-dependent projected components of the unit source vector. The term τ_i accounts for any fixed instrumental delay

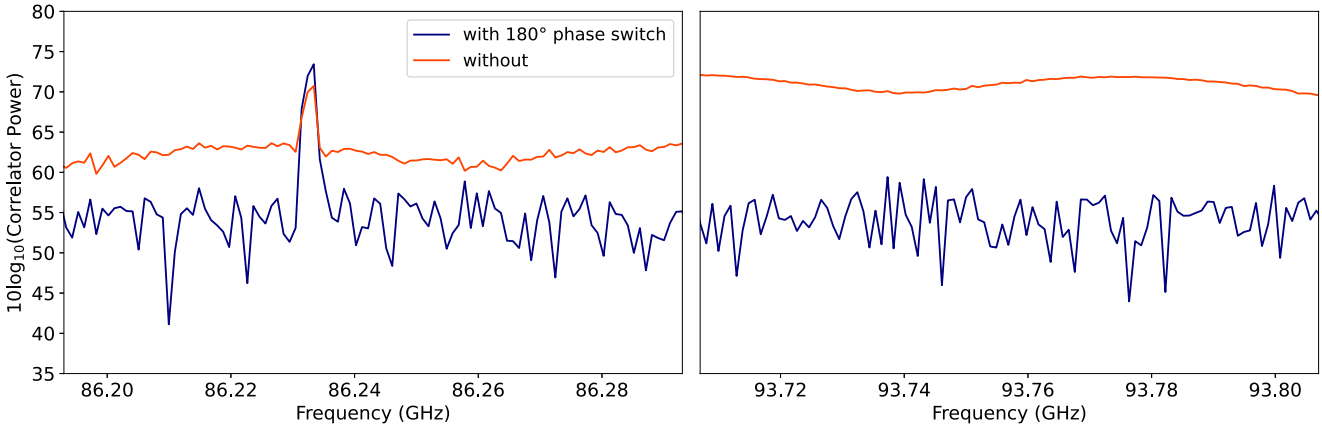

Figure 8. Time-averaged spectrum of SiO maser emission in R Leo, zoomed in to corresponding frequencies in both the USB (right) and LSB (left), after sideband separation, shown with and without cross-talk mitigation using 180° phase switching.

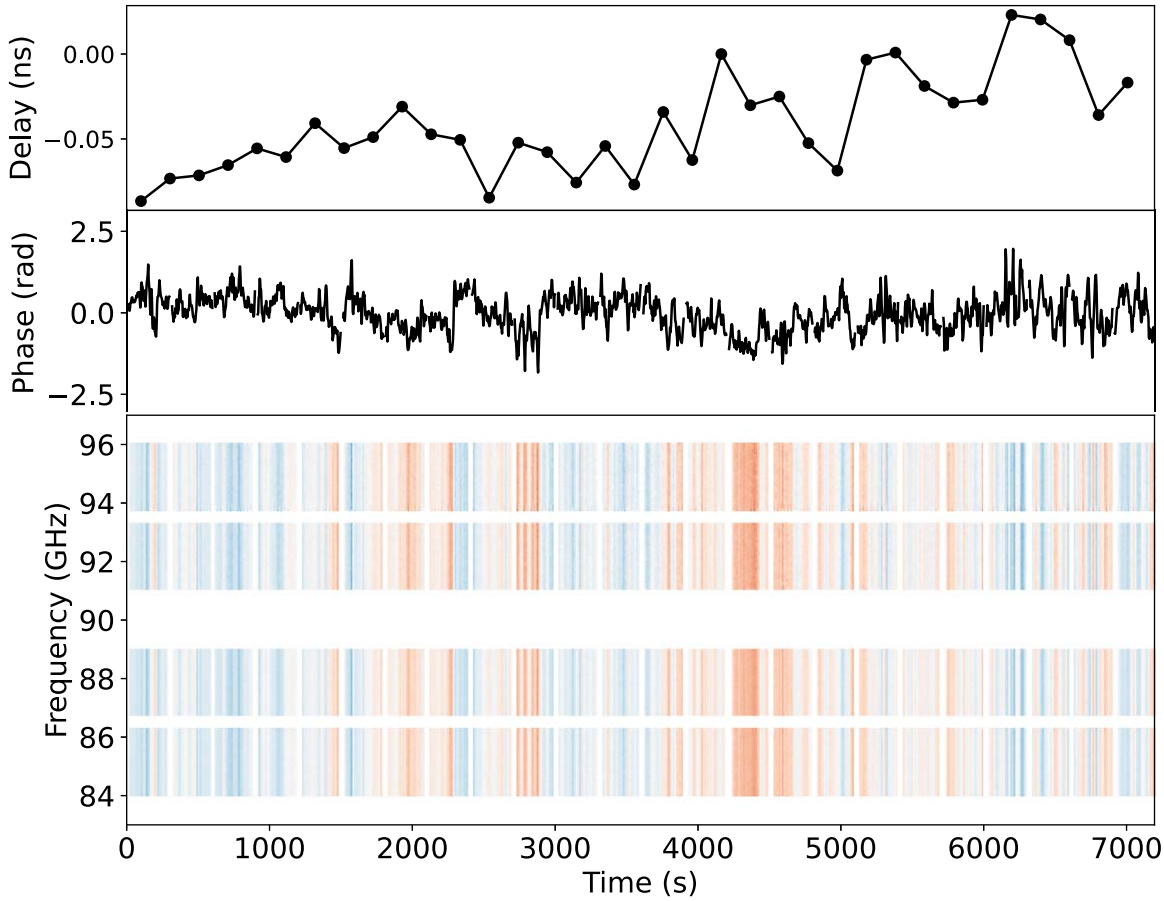

Figure 9. Behavior of phase as a function of time and frequency as seen on a 2 hr observation of bandpass calibrator source 3C 84 after refinements to the baseline solution. In the bottom panel, red indicates negative phase, blue indicates positive phase, and white indicates zero phase. The top panel of this figure shows the small amount of delay variations over the whole observation. The middle panel shows a band-averaged measurement of phase as a function of time.

errors:

$$\phi(t, \nu) = \frac{2\pi\nu}{c}[\Delta b_x s_x(t) + \Delta b_y s_y(t) + \Delta b_z s_z(t) + c\tau_i]. \quad (16)$$

We use the residual fringe data to fit for the terms Δb_x, Δb_y, Δb_z, τ_i. Figure 9 shows compelling evidence for a well-refined baseline, as we see no rapid changes in phase both across the bandpass and in the time-series data over a 2 hr observation.

4.4. Flux Calibration Solution

To verify our flux calibration procedure, as well as all other intrasequence calibrations described above, we routinely observe ALMA calibrator sources.[4] Here we present a 90-minute observation of J1342+2709, obtained on 2023 June 28, and last observed by ALMA on 2023 May 23 to be 105 ± 5 mJy at 90.4 GHz. We reduced all scans on this source

[4] https://almascience.nrao.edu/sc/

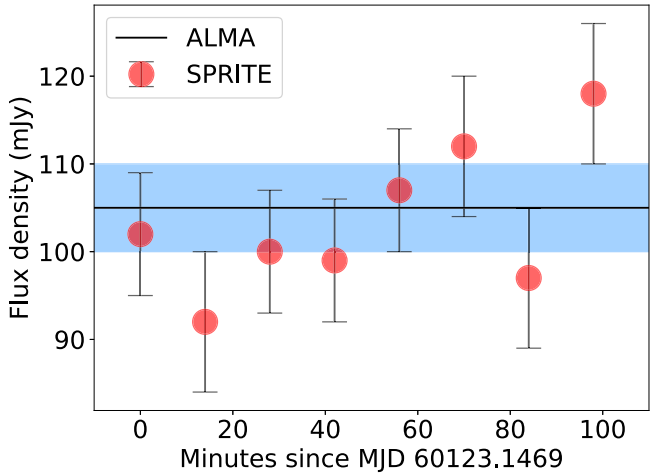

Figure 10. SPRITEly measurements of the flux density of ALMA calibrator source J1342+2709 (red points). An ALMA measurement obtained 36 days earlier is indicated as a black line (1σ uncertainty shown as a blue shaded region).

using the procedure outlined in Section 3.1. Bandpass calibration was performed using a 10-minute observation of 3C 273. Flux calibration was performed by averaging two 10-minute observations of Uranus, obtained 4 hr apart. Individual 10-minute scans on J1342+2709 were interspersed with 1-minute scans on J1333+2725 for gain calibration. We show the results in Figure 10. Measurements in each scan are consistent with the ALMA result, and we measure a mean flux density over this observation of 103 ± 7 mJy. The uncertainty includes statistical errors in the measurement and in all calibrations; the dominant source of error is the flux calibration.

4.5. End-to-end Demonstration: S2 0109+22 and DQ Tau

We describe two examples of early-science observations here that demonstrate the utility of an instrument like SPRITEly. Results are shown in Figure 11.

SPRITEly responded to an Astronomer's Telegram report (F. Roesch et al. 2023) of unprecedentedly high centimeter-wavelength flux densities of the BL Lac object S2 0109+22 (e.g., S. Ciprini et al. 2004; MAGIC Collaboration et al. 2018). S2 0109+22 is detected from radio to very high energy γ-ray wavelengths and is a classical intermediate spectrum peak source at a redshift of $\sim$0.36. The source is unlike other TeV blazars in its high degree of optical polarization variability (7%–30%) and its lack of correlation between the radio and optical light curves. The latter suggests that the radio and optical emission may not originate from the same zone of the jet. F. Roesch et al. (2023) measured flux densities of $\sim$4 Jy at 22 and 40 GHz on 2023 June 15, in excess of the 0.3–3 Jy range observed over the past four decades at these frequencies. In agreement with previous estimates of the radio–millimeter spectral energy distribution (S. Ciprini et al. 2004), we observed brighter emission at 90 GHz on 2023 June 28–29, with detectable variability over 24 hr. The measured flux densities on these 2 days were 5.1 ± 0.3 Jy and 4.2 ± 0.3 Jy.

DQ Tau is an eccentric 15.8-day binary system of $\sim$0.65 $M_{\odot}$ classical T Tauri stars (R. D. Mathieu et al. 1997). At periastron, the stars approach to within $8R_{*}$, suggesting that their magnetospheres briefly overlap. This appears to result in coincident millimeter (D. M. Salter et al. 2008, 2010) and X-ray (K. V. Getman et al. 2011) flares. The millimeter flares are explained as synchrotron emission from electrons accelerated in magnetic reconnection during periastron passage, and the observation of the Neupert effect in the relative timing of the millimeter and X-ray flares suggests that the same electrons are responsible for the nonthermal X-ray bremsstrahlung and the synchrotron emission. Despite the uniqueness of this system, only four millimeter observations of DQ Tau during periastron have been presented, with incomplete orbital-phase coverage. These observations indicate significant variations in the nature and timing of the periastron flares. We observed DQ Tau with SPRITEly during a periastron passage on 2023 June 26 (ephemeris from I. Czekala et al. 2016) and detected significantly variable emission above the typical quiescent level of 13–17 mJy observed on the days surrounding periastron passage (D. M. Salter et al. 2010). The variability timescale and magnitude are consistent with previous observations of DQ Tau, with the exception of the initial extreme ($\sim$0.6 Jy) flare observed by D. M. Salter et al. (2008). Continued monitoring of periastron passages of DQ Tau by SPRITEly will reveal the full phenomenology of these flares.

5. Future Observations and SPRITEly Deployments

The current implementation of SPRITEly will continue to observe every day with agreeable conditions. Our observations can currently be scheduled for blocks of up to 24 hr at once and include a mix of several classes of sources. We will spend several hours per day monitoring selections of nearby active stars in an experiment to measure the rate and luminosity function of millimeter flares (such as those reported in S. Guns et al. 2021). A few stars of interest include UV Ceti and YY Geminorum, which historically have been observed to have bright radio flares (P. D. Jackson et al. 1989). In light of new discoveries of millimeter emission from young stellar objects (e.g., J. Vargas-Gonzalez et al. 2023), we also plan to monitor DQ Tau and similar systems. We also plan to continue monitoring nearby radio active galactic nuclei (AGN) of interest to millimeter very long baseline interferometry observations, including M87 and Sgr A*, along with samples of AGN selected for their interesting variability. In addition to routine monitoring, we also plan to schedule observations of any newly discovered and nearby transients sources, including supernovae (in particular those of Type IIn; N. Yadlapalli et al. 2022), active X-ray binaries, and jetted TDEs.

The instrument details described in this paper represent only a starting point for the full potential of SPRITEly. With the two-antenna deployment described herein, we have made a significant effort to document existing hardware and procedures and implement scalable new hardware and software. For example, phase switching and sideband separation for multiple baselines can be accomplished by the installation of a Red Pitaya device in each new antenna. Our favored upgrade path is to make the following additions with corresponding sensitivity increases:

1. Add a third antenna, providing a factor of $\sim\sqrt{3}$ improvement in sensitivity.
2. Replace the existing 3 mm receivers with dual-polarization sideband-separating MMIC systems based on an existing design (J. W. Lamb et al. 2016), which will add a factor of $\sim 2\sqrt{2}$ improvement in sensitivity.

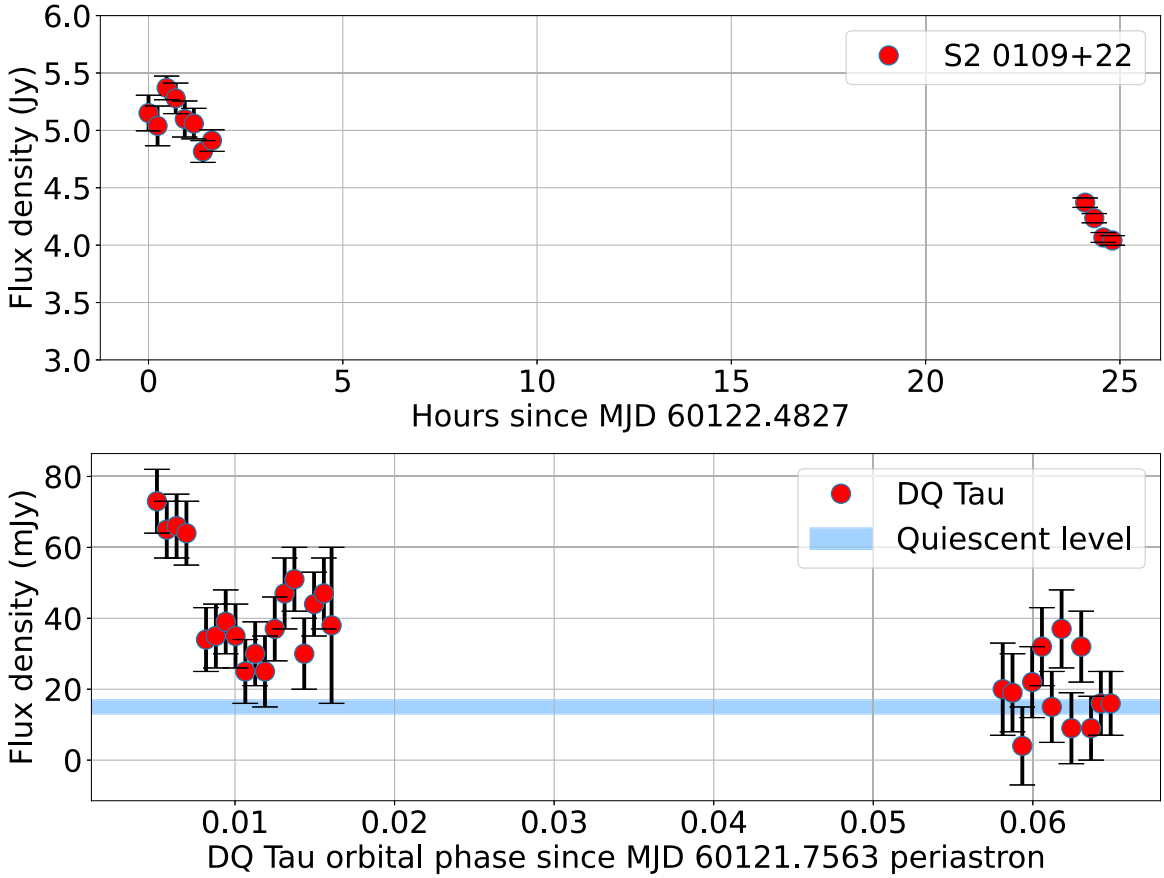

Figure 11. SPRITEly measurements of the flux density of the flaring BL Lac object S2 0109+22 (top) and the binary classical T Tauri system DQ Tau (bottom). The blue shaded region in the bottom panel indicates an approximate quiescent level for the 90 GHz emission from DQ Tau (see text for details).

3. Use the full 16 GHz of available sideband-separated IF bandwidth, adding a factor of ∼2 improvement in sensitivity.

The cumulative sensitivity increase of a factor of 6.9 over the existing two-antenna deployment would see the system achieve a noise of $\sigma_S \approx 20$ mJy s$^{-1/2}$. This is comparable to the 3 mm sensitivity of CARMA.

6. Conclusion

SPRITEly represents a new resource to explore the growing field of millimeter transients. Taking a unique lower-sensitivity and nonimaging approach allows SPRITEly to have a narrower focus purely on time-domain studies. Additionally, it demonstrates the capabilities of cost-effective and sustainable instrumentation to push forward ambitious science goals. This work outlines how we utilize a preexisting antenna and receiver system with easily implemented back-end and control system upgrades to demonstrate a two-antenna deployment of SPRITEly. Low-cost and scalable solutions, such as using the Red Pitaya device to handle lobe rotation and phase switching, may become more prevalent for future smaller-scale experiments. Another modern feature of SPRITEly is its use of open-source software codebases, such as Astropy (Astropy Collaboration et al. 2022), to rapidly develop an integrated control and data-analysis software system.

Equipped with the flexibility to schedule observations without proposal-driven constraints, we aim to use SPRITEly to monitor and follow up bright and nearby transient sources. As demonstrated by comparable instruments across the electromagnetic spectrum, including, for example, the Arcminute Micro-Kelvin Imager, the Las Cumbres Observatory, and the Neil Gehrels Swift Observatory, SPRITEly will fill an important niche even in the era of the next-generation VLA and the upgraded ALMA. Systematic studies will fill in the rate and luminosity phase space of millimeter transient sources. High temporal resolution, especially at millimeter wavelengths, holds the power to unveil rapidly variable behavior and key physical properties of energetic synchrotron sources and their environments. As the era of wide-field synoptic surveys broadens in this coming decade, from new millimeter experiments such as CMB-S4 (K. Abazajian et al. 2019) to the nearly completed optical Vera Rubin Observatory, telescopes like SPRITEly will become increasingly more important in the effort to characterize the plethora of interesting transients to be discovered.

Acknowledgments

The Owens Valley Radio Observatory is located on the ancestral homelands of the Big Pine Paiute Tribe of the Owens Valley. We recognize and acknowledge the historical and cultural significance of these lands to members of the Tribe. The authors thank staff members of the Owens Valley Radio Observatory and the Caltech radio group, including Kristen Bernasconi, Stephanie Cha-Ramos, Charlie Harnach, Sarah Harnach, Lori McGraw, Paul Rasmussen, Andres Rizo, and Michael Virgin. Their tireless efforts were instrumental to the success of SPRITEly.

This material is based on work supported by the National Science Foundation under grant No. 1935980. We acknowledge funding from the Mt. Cuba Astronomical Foundation under a grant titled "SPRITEly: A New Window into the Explosive, Dark Universe." We would like to thank the director of OVRO, G. Hallinan, for access to the Alan Moffet Funds.

ORCID iDs

Nitika Yadlapalli Yurk https://orcid.org/0000-0003-3255-4617
Vikram Ravi https://orcid.org/0000-0002-7252-5485
James W. Lamb https://orcid.org/0000-0002-5959-1285
Anthony C. S. Readhead https://orcid.org/0000-0001-9152-961X

References

Abazajian, K., Addison, G., Adshead, P., et al. 2019, arXiv:1907.04473
Alexander, K. D., Berger, E., Bower, G., et al. 2017, arXiv:1703.04692
Andreoni, I., Coughlin, M. W., Perley, D. A., et al. 2022, Natur, 612, 430
Archinal, B. A., A'Hearn, M. F., Bowell, E., et al. 2011, CeMDA, 109, 101
Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, ApJ, 935, 167
Balister, M., Batchelor, R. A., Haynes, R. F., et al. 1977, MNRAS, 180, 415
Bower, G. C., Plambeck, R. L., Bolatto, A., et al. 2003, ApJ, 598, 1140
Burton, K., MacGregor, M. A., & Osten, R. A. 2022, ApJL, 939, L6
Cao, Y., Kasliwal, M. M., Arcavi, I., et al. 2013, ApJL, 775, L7
Carilli, C., & Holdaway, M. 1997, VLA Scientific Memos
Carilli, C. L., Carlstrom, J. E., & Holdaway, M. A. 1999, in ASP Conf. Ser. 180, Synthesis Imaging in Radio Astronomy II, ed. G. B. Taylor, C. L. Carilli, & R. A. Perley (San Francisco, CA: ASP), 565
CASA Team, Bean, B., Bhatnagar, S., et al. 2022, PASP, 134, 114501
Ciprini, S., Tosti, G., Terasranta, H., & Aller, H. D. 2004, MNRAS, 348, 1379
Czekala, I., Andrews, S. M., Torres, G., et al. 2016, ApJ, 818, 156
de Ugarte Postigo, A., Lundgren, A., Martin, S., et al. 2012, A&A, 538, 44
Eftekhari, T., Berger, E., Metzger, B. D., et al. 2022, ApJ, 935, 16
Fuhrmann, L., Angelakis, E., Zensus, J. A., et al. 2016, A&A, 596, A45
Getman, K. V., Broos, P. S., Salter, D. M., Garmire, G. P., & Hogerheijde, M. R. 2011, ApJ, 730, 6
Guns, S., Foster, A., Daley, C., et al. 2021, ApJ, 916, 98
Hasselfield, M., Moodley, K., Bond, J. R., et al. 2013, ApJS, 209, 17
Hickish, J., Abdurashidova, Z., Ali, Z., et al. 2016, JAI, 5, 1641001
Hickish, J., Razavi-Ghods, N., Perrott, Y. C., et al. 2018, MNRAS, 475, 5677
Ho, A. Y. Q., Margalit, B., Bremer, M., et al. 2022, ApJ, 932, 116
Ho, A. Y. Q., Phinney, E. S., Ravi, V., et al. 2019, ApJ, 871, 73
Horesh, A., Stockdale, C., Fox, D. B., et al. 2013, MNRAS, 436, 1258
Jackson, P. D., Kundu, M. R., & White, S. M. 1989, A&A, 210, 284
Kutner, M. L., & Ulich, B. L. 1981, ApJ, 250, 341
Lamb, J. W., Cleary, K. A., Gawande, R. S., et al. 2016, Proc. SPIE, 9914, 99140H
Lamb, J. W., Cleary, K. A., Woody, D. P., et al. 2022, ApJ, 933, 183
Laskar, T., Berger, E., Zauderer, B. A., et al. 2013, ApJ, 776, 119
Leighton, R. B. 1977, Tech. Rep. AST 73-04908, https://authors.library.caltech.edu/35046/
Louis, T., Addison, G. E., Hasselfield, M., et al. 2014, JCAP, 2014, 016
MacGregor, A. M., Osten, R. A., & Hughes, A. M. 2020, ApJ, 891, 80
MacGregor, M. A., Weinberger, A. J., Wilner, D. J., Kowalski, A. F., & Cranmer, S. R. 2018, ApJL, 855, L2
MAGIC Collaboration, Ansoldi, S., Antonelli, L. A., et al. 2018, MNRAS, 480, 879
Mathieu, R. D., Stassun, K., Basri, G., et al. 1997, AJ, 113, 1841
Naess, S., Battaglia, N., Richard Bond, J., et al. 2021, ApJ, 915, 14
Parsons, A., Backer, D., Siemion, A., et al. 2008, PASP, 120, 1207
Phillips, T. G., & Woody, D. P. 1982, ARA&A, 20, 285
Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020, A&A, 641, A3
Roesch, F., Kadler, M., Benke, P., et al. 2023, ATel, 16093, 1
Salter, D. M., Hogerheijde, M. R., & Blake, G. A. 2008, A&A, 492, L21
Salter, D. M., Kóspál, Á., Getman, K. V., et al. 2010, A&A, 521, A32
Tetarenko, A. J., Sivakoff, G. R., Miller-Jones, J. C. A., et al. 2017, MNRAS, 469, 3141
Thompson, A. R., Moran, J. M., & Swenson, G. W., Jr. 2017, Interferometry and Synthesis in Radio Astronomy (3rd ed.; Berlin: Springer)
Vargas-Gonzalez, J., Forbrich, J., Rivilla, V. M., et al. 2023, MNRAS, 522, 56
Woody, D. P., Beasley, A. J., Bolatto, A. D., et al. 2004, Proc. SPIE, 5498, 30
Woody, D. P., Wiitala, B., Scott, S. L., et al. 2007, RScI, 78, 094501
Wright, M. C., MacMahon, D., Corder, S., & Sandell, G. 2009, CARMA Memo 52
Yadlapalli, N., Ravi, V., & Ho, A. Y. Q. 2022, ApJ, 934, 5
Yuan, Q., Wang, Q. D., Lei, W.-H., Gao, H., & Zhang, B. 2016, MNRAS, 461, 3375
Zauderer, B. A., Berger, E., Soderberg, A. M., et al. 2011, Natur, 476, 425