

UC San Diego

UC San Diego Previously Published Works

Title

The Simons Observatory: deployment and current configuration of the observatory control system for SAT-MF1 and data access software systems

Permalink

<https://escholarship.org/uc/item/64x6p4dr>

ISBN

978-1-5106-7527-8

Authors

Bhimani, Sanah

Lashner, Jack

Aiola, Simone

et al.

Publication Date
















2024-08-16

DOI

10.1117/12.3020237

Peer reviewed

The Simons Observatory: Deployment and current configuration of the Observatory Control System for SAT-MF1 and data access software systems

Sanah Bhimani ¹, Jack Lashner ¹, Simone Aiola ^{2, 5}, Kevin T. Crowley ¹⁰, Nicholas Galitzki ^{3,4}, Remington G. Gerras ¹², Kathleen Harrington ^{7,8}, Matthew Hasselfield ², Alyssa Johnson ⁹, Brian J. Koopman ¹, Hironobu Nakata ¹¹, Laura Newburgh ¹, David V. Nguyen ¹, Michael J. Randall ¹⁰, and Max Silva-Feaver ¹

¹Wright Laboratory, Department of Physics, Yale University, New Haven, CT 06511, USA

²Center for Computational Astrophysics, Flatiron Institute, New York, NY 10010, USA

³Department of Physics, University of Texas at Austin, Austin, TX, 78712, USA

⁴Weinberg Institute for Theoretical Physics, Texas Center for Cosmology and Astroparticle Physics, Austin, TX 78712, USA

⁵Joseph Henry Laboratories of Physics, Jadwin Hall, Princeton University, Princeton, NJ 08544, USA

⁶Department of Physics, The University of Tokyo, Tokyo 113-0033, Japan

⁷Argonne National Laboratory, High Energy Physics Division, Lemont, IL 60439, USA

⁸University of Chicago, Department of Astronomy and Astrophysics, Chicago, IL 60637, USA

⁹Department of Astronomy & Astrophysics, University of California, San Diego, La Jolla, CA 92093, USA

¹⁰Department of Physics, University of California, San Diego, La Jolla, CA 92093, USA

¹¹Department of Physics, Faculty of Science, Kyoto University, Kyoto 606-8502, Japan

¹²University of Southern California, 825 Bloomwalk, Los Angeles, CA, 90089, USA

ABSTRACT

The Simons Observatory (SO) is a Cosmic Microwave Background experiment located in the Atacama Desert in Chile. SO consists of three small aperture telescopes (SATs) and one large aperture telescope (LAT) with a total of 60,000 detectors in six frequency bands.¹ As an observatory, SO encompasses hundreds of hardware components simultaneously running at different readout rates—all separate from its 60,000 detectors on-sky and their metadata. We provide an overview of commissioning SO’s data acquisition software system for SAT-MF1, the first SAT deployed to the Atacama site. Additionally, we share insights from deploying data access software for all four telescopes, detailing how performance limitations affected data loading and quality investigations, which led to site-compatible software improvements.

Keywords: Cosmic Microwave Background, Observatory Control System, control software, data acquisition, monitoring, data access software

1. INTRODUCTION

The Simons Observatory (SO) is a newly commissioned ground-based Cosmic Microwave Background (CMB) experiment in the Atacama Desert, Chile. The Simons Observatory features three 0.5 m diameter Small Aperture Telescopes (SATs) and one 6 m diameter Large Aperture Telescope (LAT), all designed for precision measurements of cosmological signals across a wide range of angular scales. Measurements from the LAT target a

Further author information: (Send correspondence to Sanah Bhimani)

Sanah Bhimani: E-mail: sanah.bhimani@yale.edu

constraint on the sum of neutrino masses and the effective number of relativistic species, while the SAT aims to measure or constrain the primordial tensor-to-scalar ratio.¹ To accomplish these goals, the observatory fields approximately 60,000 cryogenic bolometers across the four telescopes, and requires a large set of diverse instruments for cryogenics, telescope platform operations, calibration, and instrument and observatory health. SO recently deployed all four telescopes to the site, starting with a SAT sensitive to two mid-frequency bands centered at 90 and 150 GHz, hereafter referred to as SAT-MF1.

Similar to any observatory, SO requires a software system to acquire data and provide control and monitoring of all systems and instruments. This system must allow us to follow a designated scan-strategy, coordinate commanding and acquisition across multiple subsystems, and send alerts when subsystems are in an unhealthy state.² We developed the Observatory Control System (`ocs`) to meet these data acquisition, control, and monitoring needs.³ In addition, accessing ancillary data, referred to as “housekeeping” (HK) data, is challenging due to SO’s serialized, asynchronous data structure. To enable fast and efficient exploration of HK data for the detector analysis pipeline, we developed `G3tHK`, an SQLite database with an HK data loading mechanism.

In these proceedings, we focus on the software deployment for SAT-MF1. In Section 2, we describe the `ocs` deployment for SAT-MF1, including the software and physical hardware systems for the platform and site. In Section 3 we describe the deployment of `G3tHK`, its use cases, lessons learned, and future plans.

2. SAT-MF1 OCS SOFTWARE DEPLOYMENT

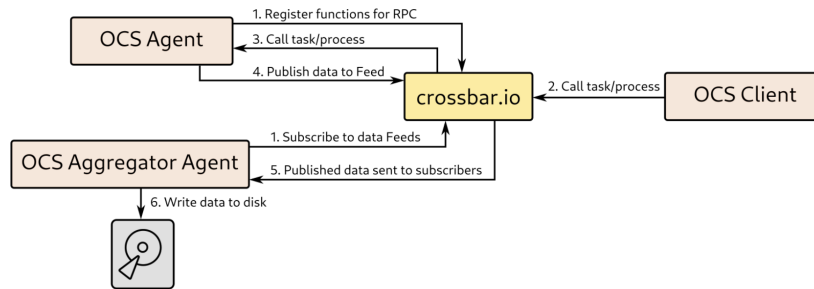


Figure 1: The `ocs` architecture, including major components, is shown. The `crossbar` router (yellow) facilitates communications between all other elements of the system. An `ocs` Client (right) is used to call a start, stop, or status for an operation. The operation is run by an `ocs` Agent (top left). Agents can facilitate `ocs` control and data acquisition for hardware units. If an operation collects data, that data is published to a Feed, which is accessed by an Aggregator Agent (bottom left). The Aggregator Agent writes the data to disk. Image by Brian Koopman.³

`ocs` ensures commanding and data acquisition of the instrumentation, initiating and ending processes like telescope scanning, tuning detectors onto their superconducting transitions, or adjusting calibration equipment. The `ocs` architecture comprises two primary components: agents and clients, as shown in Figure 1. `ocs` agents are long-running software servers that interface with hardware or software components to coordinate data acquisition.³ Agents have operations classified as Tasks and Processes:

- **Tasks** run for a finite amount of time and have a natural completion, such as moving the telescope platform or setting the heater range for servoing a temperature sensor in a cryostat.
- **Processes**, like the continuous “acq” process for data acquisition, run until stopped.³

`ocs` clients are scripts that orchestrate the actions of one or more `ocs` agents over the network. For example, an `ocs` client can instruct a managing agent to move a telescope platform to a specific position, after which the agent executes the command and returns the final position data. Crossbar.io facilitates interactions between agents and clients. In addition, this system provides a real-time visualization for slow rate data via Grafana,³

as described in Section 2.1. Detailed specifics on the alerts based on this system are described in Nguyen et al. 2024.²

In this section, we describe the deployment of `ocs` for the SAT-MF1 platform and receiver, the first receiver deployed in SO. We begin with a discussion of the agents required for the SAT-MF1 receiver, platform, and site. We also include where the agents are running, and provide a discussion of the user interfaces and real-time visualization for the system.

2.1 SAT-MF1 Agents

SAT-MF1 is one of two mid-frequency SATs and the first to begin observations, which commenced in October 2023. It consists of approximately 12,000 transition-edge sensor (TES) bolometers cooled to 90mK,⁴⁻⁷ a spinning cryogenic half-wave plate at 1K temperatures, and an extensive set of housekeeping components.

Currently, there are a total of 69 SAT-MF1 `ocs` agents running on-site. Of these, 8 are ‘core `ocs`’ agents: the Aggregator, the Registry, the InfluxDB Publisher, and 5 Host Manager agents. Detailed descriptions of these agents can be found in Koopman et. al. 2024, but we briefly describe each for context:

- Aggregator – subscribes to ‘slow’ data rate feeds, referred to here as housekeeping (HK) data, and writes to disk. Slow data comes from all components that are not the fast-cadence (200 Hz) detector data.
- Registry – monitors all active agents, tracking the status of each agent’s Tasks and Processes.
- InfluxDB Publisher – similar to the Aggregator agent, but publishes the HK data to an InfluxDB server for viewing on Grafana³ instead of writing to disk.
- The Host Manager – provides a way to start and stop agents running on a given machine or host.

The remaining 61 agents within the SAT-MF1 `ocs` system are integral to the operation of the SAT’s receiver and telescope platform: the telescope control unit (ACU); the cryogenics (Bluefors dilution refrigerator, thermometry, vacuum components); cryogenic half-wave-plate and wiregrid calibrators; and SMuRF (detector readout) agents for monitoring, controlling, and data transfer.

Additionally, the site `ocs` system and supporting software infrastructure facilitate remote operations for the entire observatory. Separate from SAT-MF1’s `ocs` setup, there are 7 additional site agents crucial for assessing observatory conditions and ensuring effective telescope operations, noted in Table 1: 5 generator agents that monitor fuel levels for the site’s diesel generators; a radiometer agent that monitors atmospheric precipitable water vapor (PWV) to assess whether conditions are suitable for observations on any given day; and a weather monitor that tracks local site conditions such as wind speed, wind direction, humidity, and more.

The salient features of SAT-MF1 agents and their data characteristics are summarized in Table 2, including data rate, number of fields, data type, and the development origin of each agent. In total, the 76 agents (69 SAT-MF1 + 7 site) generate data at a rate of approximately 3.74 Mbps. This data rate includes 2074 `ocs` fields, each representing a specific data component collected by an agent, such as temperature, voltage, or status information. For example, both temperature and voltage readings for each of the 11 Lakeshore 372 channels would be represented by 22 separate `ocs` fields. Of the 2074 fields, SAT-MF1 contributes 1524 fields with a data rate of 3.58 Mbps, while components of the site `ocs` sub-system mentioned in this proceedings contribute 550 fields with a data rate of 0.162 Mbps. Many of the agents described in Table 2 were developed by hardware experts across the SO collaboration, aligning with one of the primary goals of the `ocs` architecture.

Figure 2 provides the current layout for SAT-MF1 HK agents, including a subset of the site software system on the bottom right. The site setup provides the framework for the observatory for alerting (Campana²), commanding (`ocs-web`), and monitoring `ocs` data and logs (Grafana, Loki).⁸ For details surrounding the complete site layout, refer to Koopman et al. 2024.

The fast-cadence data from the detectors is acquired separately by the SMuRF readout system.^{9,10} The SMuRF electronics readout is managed by four `ocs` agents: Pysmurf Monitor, Pysmurf Controller, Crate, and Magpie. Since real-time monitoring of detector streams is valuable but unfeasible with our ‘slow’ HK acquisition

streams, the Pysmurf Monitor Agent can send a small subset of downsampled detector timestreams to the HK aggregation system for viewing through Grafana.

Subsystem	Site ocs Agent	Data Rate (kbps)	No. of Fields	Data Type	Where Developed
Environment	UCSC Radiometer	0.015	1	Float	DAQ
	Generator	160.4	474	Float, Int	DAQ
	Weather Monitor	1.8	75	Float, Int	DAQ
Total	–	0.162 Mbps	550	–	–

Table 1: A table describing data characteristics of a small set of ocs agents for the site. The total data rate for this system is 0.162 Mbps, with 550 fields across 7 agents.

Subsystem	SAT-MF1 ocs Agent	Data Rate (kbps)	No. of Fields	Data Type	Where Developed
Telescope	ACU	2520	197	Float	DAQ
Cryo	Bluefors	288	45	Float	DAQ
	Pfeiffer TC400	0.384	6	Float	UCSD
	Cryomech CPA	5.44	34	Float	UCSD
	Lakeshore 372	7.04	11	Float	DAQ
	Lakeshore 240	8.32	52	Float	DAQ
	Lakeshore 425	0.064	1	Float	UTokyo
	Flowmeter	0.015	12	Float	DAQ
Labjack	Labjack	7.42	80	Float	UCSD
Power	ibootbar	1.24	144	Float, Int, Str	DAQ
	UPS	0.91	105	Float, Int, Str	DAQ
Calibration	HWP PCU	0.024	1	Str	UTokyo
	HWP PID	1.28	4	Int	UCSD/UTokyo
	HWP PMX	5.312	6	Float, Int, Str	UTokyo
	HWP Gripper	0.31	24	Float, Int	DAQ/UCSD
	HWP Encoder	10.8	42	Float, Int	UTokyo
	Wiregrid Encoder	4.16	13	Float, Int	Kyoto Uni.
	Wiregrid Kikusui	0.32	5	Float	Kyoto Uni.
	Wiregrid Actuator	0.83	13	Int	Kyoto Uni.
ocs Core	Registry	157.3	365	Float, Int, Str	DAQ
	Aggregator	–	–	–	DAQ
	Host Manager	–	–	–	DAQ
	InfluxDB Publisher	–	–	–	DAQ
Smurf	Pysmurf Monitor	8.512	133	Float	Detector/DAQ
	SupRsync	1.728	20	Float, Int, Str	Detector/DAQ
	Pysmurf Controller	544.32	126	Float, Int, Str	Detector/DAQ
	Crate	2.016	84	Str	Detector/DAQ
Total	–	3.58 Mbps	1524	–	–

Table 2: A table describing data characteristics of ocs agents setup for SAT-MF1, divided by subsystem. The number of fields represents those currently online for all instances of that agent for SAT-MF1. For instance, the 12 flowmeter data fields in the cryo subsystem correspond to 6 active flowmeter agents, each with 2 fields (flow in liters/min and temperature in Celsius). The total data rate for the instrument is 3.58 Mbps, with 1524 queryable data fields from the telescope’s 69 agents. The registry agent, described in Section 2.1, monitors all active agents and the status of their corresponding Tasks and Processes. It currently tracks 365 fields, contributing the most to the data rate. Data rates for other core ocs agents are not noted, as their characteristics do not require data fields being written to disk.




Figure 2: `ocs` deployment for SAT-MF1 at the site, referred to in the software as `satp1`. This layout breaks down the distribution of agents across the instrument by host, further detailed in Section 2.3. A subset of the site software system, including `ocs`, `ocs-web`, and open-source services is shown on the bottom right. Site agents monitor site-wide hardware relevant to observatory operations, which includes monitoring precipitable water vapor (PWV) values, weather conditions, and generator fuel levels. Details on the open-source software and `ocs-web` can be found in Koopman et al. 2020.

2.2 User Interfaces

The tools and interfaces used by and designed for SO provide essential functions for monitoring, commanding, and visualizing the status and data of various HK agents. Below is an overview of the user interface with `ocs` and supporting software for SAT-MF1, now fully operational and observing:

- The ‘sequencer’ – a tool with a graphical front end designed to run `ocs` clients.⁸ It utilizes another library to enable the sequential command of multiple clients, ensuring efficient and coordinated operations.^{8,11}
- Grafana – open-source time series display for live monitoring, see Figure 3.
- `ocs-web` – a collection of `ocs` clients, presented as a graphical user interface in a web browser and implemented using JavaScript, is primarily used for direct interaction with a given agent for specific commands;⁸ see Figure 4.
- Campana – an alarms and alerts system² based on Grafana.
- Loki – a log aggregation tool developed by Grafana Labs, allowing users to follow real-time logs or query past logs directly within the Grafana interface.

A more comprehensive description on the supporting software infrastructure can be found in Koopman et al. 2024; more details on the alarms and alerting system can be found in Nguyen et al. 2024.

Since commissioning began, both on-site and remote users have constantly interacted with `ocs` at the site. SO uses Remote Observing Coordinators (ROCs) to remotely monitor the telescope and follow the health of the instrument. Now transitioning to the operations phase, ROCs follow a routine set of actions, typically performing the following tasks during each shift:

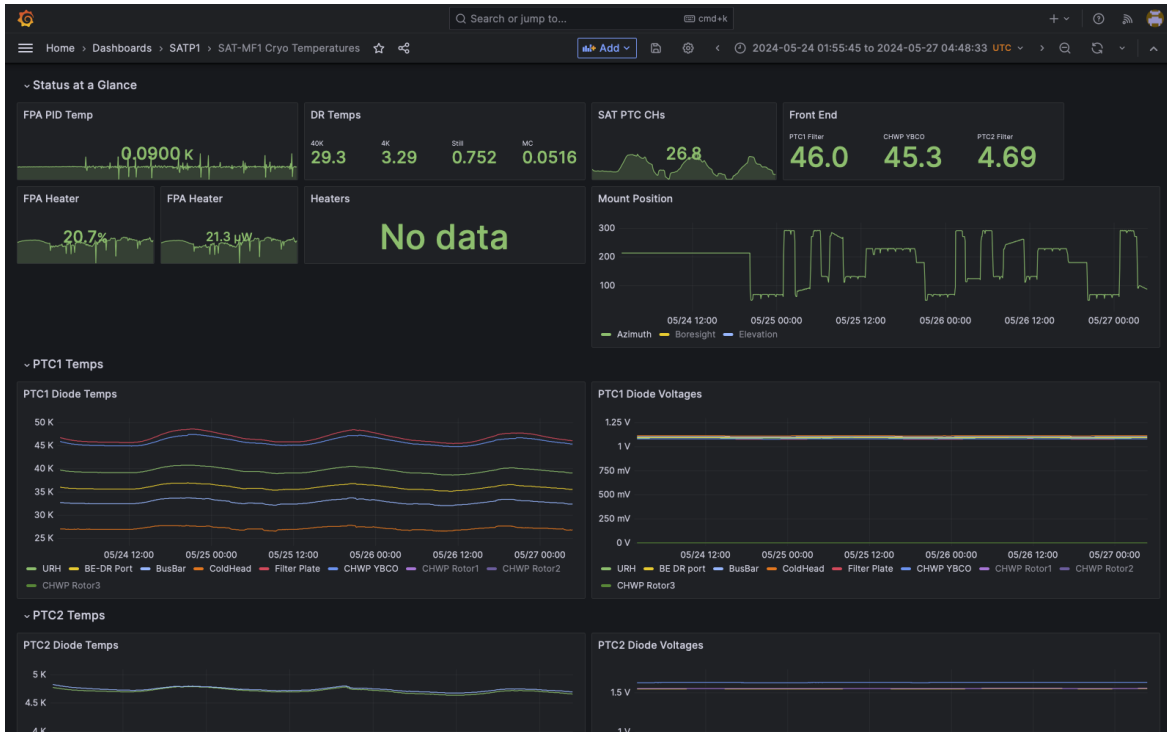


Figure 3: Screenshot of the the SAT-MF1 Grafana dashboard for monitoring the state of the dilution refrigerator and cryogenic temperatures, as well as brief information about the state of the telescope platform.

1. **Check PWV Values:** First, the ROC checks water vapor conditions on Grafana for a real-time weather estimate. If the values are favorable (less than 2mm), they proceed to the next step.
2. **Deploy Schedule:** The ROC deploys a pre-made schedule (sequence of commands) onto `nextline`,¹¹ the front-end feature of the sequencer. `nextline` runs the schedule.
3. **Monitor Alarms:** While the schedule runs, the ROC monitors the Grafana alarms page. If any alarms go off, they receive notifications via text, email, and/or phone (using Campana to choose their preferred notification methods²). They then check the status of the specific alarm on Grafana.
4. **Check Agent Logs if Necessary:** For agent communication-related alarms, the ROC may use Loki on Grafana to view logs of the failing agent.
5. **Pause Schedule:** If the alarm is severe enough, the ROC needs to put the telescope in a specific state, which requires pausing the deployed schedule. This is done using the `nextline` front-end feature.
6. **Command Hardware:** After the schedule is paused, the ROC commands the telescope using `ocs-web`. This can include commanding instruments like the cryogenic half-wave plate, thermometry readout devices, and/or the telescope ACU. By clicking any `ocs` Task on the right-hand side of an agent panel in `ocs-web`, they effectively run an `ocs` client. See Figure 4.
7. **Handle Site Communications:** Even with telescope operations running smoothly, the ROC may receive notifications from SO's remote operations channel (`satp1-remote-ops` on Slack for SAT-MF1). This may require pausing the schedule and changing the state of the instrument for site testing.

2.3 Nodes and Hosts Setup

The `ocs` agents run on computing resources at the site. Many agents can communicate via HTTP or TCP interfaces, allowing them to run on any computer within the site network. Each telescope platform is assigned

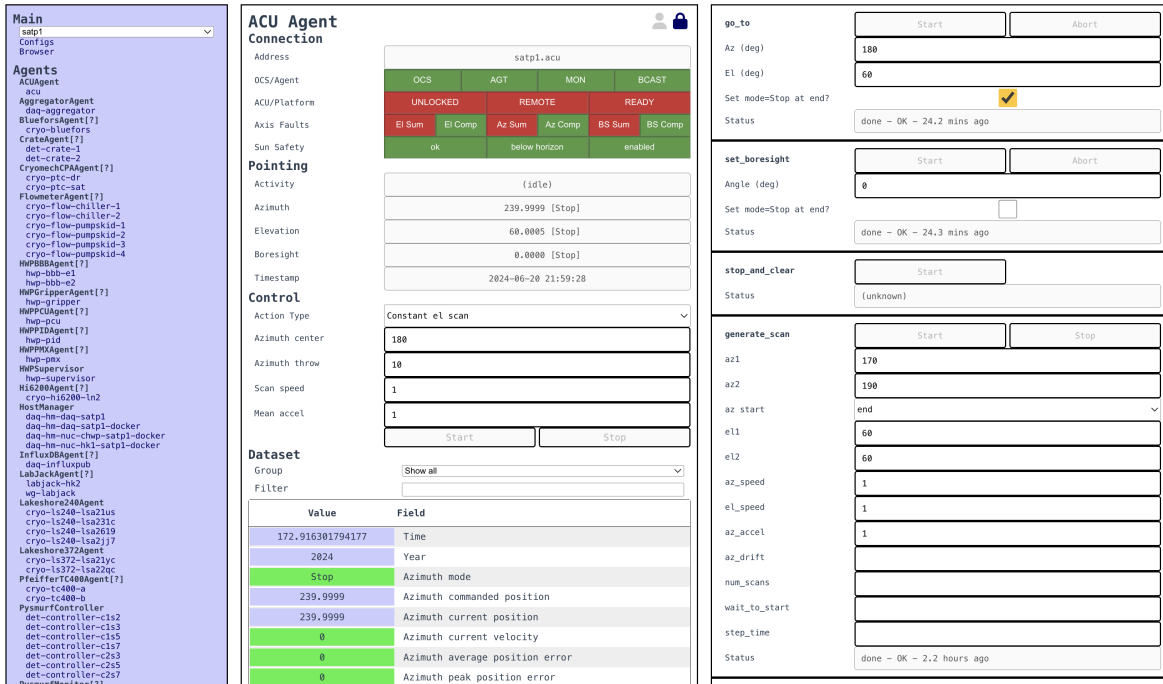


Figure 4: Screenshot of ocs-web for SAT-MF1’s telescope platform agent, the ACU. The right-hand side displays several tasks a user can command the ACU to perform via ocs-web. For example, if weather conditions deteriorate quickly during a scheduled observation, a remote observer will use the `go_to` Task to move the ACU to weather-safe azimuth and elevation positions. Entering parameters for `go_to` on the ocs-web panel runs an `ocs Client`.

a virtual machine hosted on a Linux server.⁸ In this proceedings, we will refer to this virtual machine, or any similar computing unit within a network, as a node.

`daq-satp1` serves as the dedicated site-managed node for SAT-MF1, running 36 out of the 69 agents. Interface limitations or impracticality of network communication for some hardware components required nodes to be physically installed closer to the hardware. Since `daq-satp1` is hosted on a Linux server that is located in the site office, it could not host agents for these types of hardware, necessitating the integration of additional nodes into SAT-MF1’s `ocs` system. Most other nodes on the platform are Intel NUCs* (mini-PCs), although can also be anything with a processor (BeagleBone[†], Raspberry Pi[‡], etc.). The `crossbar` router is configured on `daq-satp1`, allowing all agents, whether on `daq-satp1` or other nodes, to connect to it.⁸ A dedicated port on the router ensures that SAT-MF1 operations remain isolated from those of other telescopes.

2.3.1 SAT-MF1 Hosts

The agents described in Section 2.1 can be grouped by hosts, denoted in Figure 2. A host is a physical or virtual environment that provides the necessary infrastructure for running various applications and services. It can be a bare-metal server or a virtualized instance, such as a Docker host, offering computational resources, storage, and network connectivity. A host may manage multiple containers, each functioning as an independent computing environment within the larger system. We predominantly use Docker for deploying `ocs` within SO as it allows us to package the software, libraries, and configuration files for each agent into distributable containers, creating isolated and reproducible environments that are invaluable for site deployment.

With respect to SAT-MF1’s `ocs` infrastructure (see Figure 2):

*<https://simplynuc.com/onyx/>

[†]<https://www.beagleboard.org>

[‡]<https://www.raspberrypi.com>

- `daq-satp1` is a node,
- `daq-satp1-docker` is the Docker host on `daq-satp1`. This host is where the 36 agents on `daq-satp1` reside.
- each agent within the `daq-satp1-docker` host is its own container, i.e., 36 containers

This structure applies across all nodes and hosts within the SAT-MF1 system.

In Figure 2, we see that `daq-satp1` also contains a second Docker host that manages the sequencer containers, separate from the 36 agents. In addition, there is a third Docker host (not pictured) on `daq-satp1`, which manages the `crossbar` Docker container to house the router that interfaces with all of SAT-MF1’s agents. It is isolated from any host on the network that contains the agents. This separation ensures that restarting SAT-MF1-specific hardware and agents will not disrupt long-term services like `crossbar.io`.³

2.3.2 Enclosures and Node Setup

SAT-MF1 houses detector readout and housekeeping hardware in an enclosure system of 12 units. These enclosures are mounted on the telescope and are categorized as: SMuRF readout enclosures, cold half-wave plate and wiregrid enclosures, the azimuth enclosure, the coolant control enclosure, thermometer readout enclosures, the gas handling system enclosure, and the turbo pump enclosure. They also include power supplies such as iBootBars and UPS units, as well as timing hardware. Table 3 includes a list of SAT-MF1 nodes at the site and their corresponding enclosures.

The NUC computer housing the `nuc-chwp-satp1-docker` host for the cold half-wave plate (CHWP) readout and the NUC containing the `nuc-hk1-satp1-docker` host for thermometry readout via the Lakeshore Cryogenics devices (Lakeshore 240s) are both on Linux machines and operate similarly to the larger `daq-satp1-docker` host.

The Bluefors NUC is a Windows computer running the Bluefors 3He/4He LD400 dilution refrigerator readout. The Bluefors agent queries the logs generated by the Bluefors software on the NUC and publishes the data to both the aggregator agent and InfluxDB for Grafana visualization and alarms. Note that this Bluefors setup runs within a Docker container, but is the only one without a Host Manager agent, as the Host Manager is not Windows-compatible.

All three NUCs—regardless of OS—publish feeds, query data, command agents as applicable, and write data to disk via the aggregator agent through the `crossbar` connection on `daq-satp1`.

2.4 Enhanced Site Deployment Structures

One goal while developing `ocs` was for the agents to be developed by hardware experts in the lab, such that deploying the software required minimal changes from their lab experience. Nonetheless, deployment to the site did include a few notable changes: the use of a centralized configuration file management system;⁸ the integration of the Host Manager agent;^{3,8} expansion of agents covered by `ocs-web`, and now the primary use for commanding each individual agent at the site; significant refactoring of critical calibration agents; and the implementation of Campana for alarms and alerts.²

Site configuration files (SCFs) are a core component to the deployment of `ocs` agents—configuration files are required on each computer running `ocs`. A SCF contains information such as the `crossbar` router address, a list of agents grouped by host, and a unique ‘instance-id’ for each agent instance to differentiate between multiple copies (e.g., four Lakeshore 240s). Prior to site deployment, SCFs were managed by each lab developing SO’s receivers. At the site, SCF management requires telescope experts to push to a private GitHub repository, `ocs-deployment-configs` which uses a GitHub Actions runner that executes an Ansible playbook to deploy SCF changes to the site computers. This allows the configuration file to be automatically deployed and updated when a new agent has been added or when an existing agent parameter is added to the SCF.

While critical agents for the telescopes were initially developed and integrated in SO labs before commissioning, new agent development or updates to existing agents are not restricted to off-site locations. A `daq-dev` node exists to enable agent development at the site and to test directly on the site hardware.⁸ This node allows for software updates to be tested without introducing test data sets into the full SO data stream.

Agents for the half-wave plate were refactored and tested during lab integration to find an optimal setup for the site, resulting in the breakdown of calibrator actions into separate, related half-wave plate agents.

Campana has proven indispensable for determining the exact issues when alarms sound, which is crucial in a large observatory. For further information, refer to Nguyen et al., 2024.

3. HOUSEKEEPING DATA ACCESS TOOLS

At the time of writing, across the entire observatory (four telescopes + the site `ocs` systems), we have over 250 agents and over 5000 slow-rate housekeeping data fields. It is effective to have a system that allows efficient investigation of these fields without the need to repeatedly read full data files for each analysis. In this section, we briefly describe the data formats used for SO, a database we developed to query HK data, and highlight some of the most important applications of this database.

3.1 Data packaging

To capture all observatory data, the Simons Observatory has adopted the `.g3` file format, originally developed by the South Pole Telescope as part of their `spt3g` software[§]. A `.g3` file is a series of *frames*, where each frame is a mapping from strings to custom serializable data types. Each frame typically contains either metadata, or data taken by a data provider over a period of time. The format has been adapted to meet SO’s specific requirements and is referred to as `so3g`. For SO, the data files are heterogeneous: housekeeping files are written hourly, while detector `.g3` files are written every 10 minutes, with specific detector setup operations recorded separately from the detector timestream data. For SO, typical CMB observations are ~65 minutes long whereas planet scans can last approximately 1.5 to 3 hours, further highlighting the variability in data file durations. Hence, SO’s data packaging follows three levels of data production:

- Level 1: Initial acquisition of data into memory.
- Level 2: Level 1 data are written to disk, with a specific folder structure, becoming ‘aggregated data’ known as `.g3` files.
- Level 3: Data files are grouped together into the Level 3 format, hereafter known as books. Books are essentially directories of `.g3` files which have been re-framed and reorganized to reflect detector observations for user-friendly analysis. Books are our final storage format, to be used in analysis both at the site and off-site. For more details on data packaging, refer to Guan et al., 2024^{.12}

While the reshuffling from `.g3` files to books is critical for organizing detector data by telescope observation (e.g., Planet scan vs. CMB scan), housekeeping books remain unchanged from their Level 2 `.g3` counterparts.

3.2 G3tHK Database

We developed `G3tHK`, which consists of two components: a database and a data loader. `G3tHK` uses an SQLite database and organizes data into three tables: `.g3` files, HK `ocs` agents, and the agents’ corresponding fields. For example, for SAT-MF1, a given `.g3` HK file has data from all 69 agents, and 1524 fields, as described in Section 2.1. The `.g3` files table stores brief information about each HK file with attributes such as: unique file id, filename, file start time, file location, and source of information (i.e., `satp1`, `satp2`, `satp3`, `lat`, `site`); see Figure 6.

The HK agents table in Figure 7 contains information about each agent within an HK file. Attributes in this table include: the same file ID in the `.g3` files table, a unique agent ID number, the instance-id, and file start and end timestamps. The HK fields table lists all fields for each agent in every HK file with attributes including the unique file ID, the unique agent ID, field name, and field start and end time (which may be different from the file start and end time). We note that this database is populated for every file, regardless of observation or detector status, thus keeping a record of all HK data recorded at the site, even if we were not observing the sky.

[§]https://cmb-s4.github.io/spt3g_software/

There are two separate instances of this database structure. One is used by the bookbinder, as described in Section 3.2.1, to transform Level 2 data to Level 3. This database contains only the information relevant for data packaging from Level 2 to Level 3. For SAT-MF1, this means the database includes 153 fields per HK .g3 file (compared to 1524), reducing the need to query a much larger database during the bookbinding process. The second database instance is used for HK data exploration and is streamlined to include only the agents most commonly queried for contextualizing detector behavior. It primarily excludes information from core ocs agents and power subsystem agents, such as the ibootbar and UPS, which display extensive status information on power supplies.

3.2.1 For Bookbinding

SAT-MF1 contains approximately 12,000 detectors distributed across 7 silicon wafers, with .g3 data files organized by wafer. The bookbinder requires precise start and end times for each wafer, as well as information on which wafers were streaming data, to re-frame the detector data. These times are obtained from the .g3 file itself and stored in the G3tHK instance developed for the bookbinder.

In Figure 5, two of the seven wafers for SAT-MF1 are shown: mv22 and mv29, along with their corresponding streaming times in a single hour long HK file. The bookbinder uses the G3tHK instance by querying the pysmurf monitor for any of the streaming wafers to accurately reframe the .g3 files. The pysmurf monitor tracks when each wafer was streaming by using the field that contains `AMCcSmurfProcessorS0Streamopen_g3stream`, which indicates which pysmurf monitors were in control of which wafers during specific periods. This ensures that timestreams and metadata from each wafer are grouped correctly for permanent archiving via the bookbinder.

file_id	agent_id	field	start	stop
1	3	satp1.det-monitor-so3.feeds.ufm_mv22_meta. AMCcSmurfProcessorS0Streamopen_g3stream	1719018929.05	1719020429.49
1	3	satp1.det-monitor-so3.feeds.ufm_mv29_meta. AMCcSmurfProcessorS0Streamopen_g3stream	1719018915.94	1719020416.36
1	4	satp1.daq-sync-smurf-so3.feeds.archive_stats. finalized_until	1719018767.48	1719020555.30

Figure 5: An example of a G3tHK HK fields table is shown, listing two fields from the `daq-monitor-so3` pysmurf monitor and one `daq-sync-smurf-so3` suprsync agent. The field names containing `AMCcSmurfProcessorS0Streamopen_g3stream` have start and stop times that indicate which detector wafers were active and streaming data. The field `daq-sync-smurfso3.feeds.archive_stats.finalized_until` indicates the last time the suprsync agent for a SMuRF server successfully checked the files and copied them to `daq-satp1`, making them available for bookbinding.

id	filename	global_start_time	path	aggregator
1	1717737416.g3	1717737416	/so/level2-daq/satp1/hk/17177/1717737416.g3	satp1

Figure 6: A G3tHK HK files table. The `global_start_time` is derived from the `ctime` in the filename and is provided as a separate column in the SQLite table to facilitate queries for files within a specific time range. In this table, SAT-MF1 is noted as `satp1` under the `aggregator` column. At the observatory site, the aggregator serves as the identifier for the agents assigned to each telescope, as well as the site system. The aggregator name is also part of the full data field, as shown in Figure 5.

The detector readout system have their own servers where data is framed into .g3 files. The suprsync agent manages these files, ensuring they are transferred from the server to the site-managed DAQ node (in this case, `daq-satp1`) where the .g3 files are supposed to reside. For bookbinding, the stop time for the field `daq-sync-smurf-so3.feeds.archive_stats.finalized_until` indicates the last time the `daq-sync-smurf-so3` suprsync agent successfully checked the files on the SMuRF server and copied them to `daq-satp1`, making them

id	file_id	instance_id	start	stop
1	1	daq-sync-smurf-srv20	1717719258.48711	1717721371.14252
2	1	daq-sync-timestreams-so3	1717719174.79869	1717721316.99255
4	1	daq-sync-smurf-so3	1717719341.73984	1717721144.89606
7	1	daq-sync-timestreams-srv20	1717719186.00708	1717721306.23980

Figure 7: An example of an G3tHK HK agents table. The instance-id’s are unique identifiers for each agent instance for the suprsync agent, where there is a separate agent instance for detector timestreams and detector metadata.

available for bookbinding. This finalized time and the streaming times from the pysmurf monitors are critical to the bookbinder.

3.2.2 For Aggregate HK Properties

As noted above, G3tHK also functions as a data loader; it has also been developed for Level 3 data analysis, enabling efficient exploration of specific HK data sets. G3tHK uses Level 3 data to summarize instrument information (mean, median, minimum, maximum, and standard deviation) for each field in an HK file, allowing users to query specific properties without reading entire files. At Level 3, G3tHK focuses on a reduced number of agents. In a SAT-MF1 G3tHK instance, ~ 14 of SAT-MF1’s 69 agents are saved per HK file, reducing the number of fields from ~ 1500 to ~ 400 . This subset includes only essential agents needed to diagnose the telescope’s cryogenic, calibrator, and platform statuses, such as Lakeshore 240s, Lakeshore 372s, Bluefors dilution refrigerator, half-wave plate encoder, wiregrid encoder, and the ACU. By excluding hundreds of fields from power supply agents and core ocs agents like the registry agent, it prevents the database from growing excessively large at a rapid rate. Access to this data is critical for hardware experts to debug and fine-tune telescope parameters, as well as to explore systematics, as the database continuously collects information regardless of whether an observation is occurring. For example, one might investigate the median focal plane temperature from a Lakeshore 372 in relation to scan speed from the ACU using the values provided in the database.

file	agent	field	start	stop	med	mean	min	max	stdev
1	1	satp1.cryo-ls372-lsa21yc.fe eds.temperatures.Channel_02_R	1717938677	1717942271	21416	21414	21309	21444	11.87
1	1	satp1.cryo-ls372-lsa21yc.fe eds.temperatures.Channel_02_T	1717938677	1717942271	89.99	89.99	89.91	90.31	4E-5
3	6	satp1.cryo-bluefors.feeds. pressure_ch1	1717945871	1717948270	6.23E-6	6.22E-6	6E-6	6.44E-6	9E-8

Figure 8: Level 3 SAT-MF1 G3tHK example of HK fields. Here, aggregate properties are provided for 2 fields from a Lakeshore 372 device (temperature and resistance for channel 2) and a Bluefors dilution refrigerator pressure sensor (for channel 1). These aggregate properties are calculated per HK file, and can be used for quick analysis without the need to explore .g3 files. However, they were most commonly used to load data under specific conditions such as when the scan speed of the telescope was a specific value and the maximum focal plane temperature was ~ 95 K, for example.

Note that multiple databases can be queried to improve and refine data exploration. For example, a common query involves checking the telescope’s azimuth position using the SAT-MF1 G3tHK instance when atmospheric water vapor conditions were less than 2mm, which requires querying the site G3tHK instance. This type of query can cover a time range from a few weeks to several months, returning data only from files that meet the specified conditions. With SAT-MF1 being one of two mid-frequency SATs deployed for SO, G3tHK is also highly beneficial for quickly expanding analysis across telescopes under specific conditions, enabling efficient cross-analysis.

file_id	agent_id	field	start	stop	med	mean	min	max	stdev
1	2	site.env-radiometer-class.feeds.pwvs.pwv	1717934895.73	1717938437.28	1.35	1.34	1.21	1.46	0.07
1	1	site.env-vantage.feeds.weather_data.wind_speed	1717938526.69	1717942118.03	8.00	8.35	3.00	16.00	2.43

Figure 9: Level 3 site G3tHK example where the top entry describes atmospheric water vapor readings from the site radiometer, and the bottom field identifies wind speed from the weather monitor. These parameters, especially PWV, can be used in conjunction with other G3tHK instances to refine data exploration.

3.2.3 HK Data Loading

Loading SO HK data is cumbersome due to the serialized, asynchronous data structure, which provides flexibility but results in a non-seekable format. Data sources can be sporadic (e.g., calibration in the lab) or consistently present (e.g., cold and warm thermometry). G3tHK’s data loader was developed for lab use, but its reading time increases with the number of .g3 files. Consequently, after five months of commissioning data from SAT-MF1, data loading became prohibitively slow.

To address this, a new data loading mechanism called `hkdb` was developed to replace the one most commonly used on the backend of G3tHK, enabling fast and efficient exploration of HK data for the detector analysis pipeline. This mechanism uses a supplemental database to index each .g3 frame. For each data frame, this index tracks the .g3 file in which it is located, the byte offset, start and stop time of the frame, the `ocs` agent instance-id, and the feed³ name writing the data. At the site, this index is implemented using PostgreSQL for compatibility with the NFS and is deployed per telescope and the site system.

4. CONCLUSION

This proceedings described the deployment of `ocs` for SAT-MF1, the first of SO’s SATs commissioned and observing at the Atacama site. We provided an overview of the software, nodes, and host architecture at the site, as well as the data characteristics for the receiver. We also discussed a smaller subset of the site software system that enables successful telescope operations for SAT-MF1. This subset also includes context on the use of open-source software hosted on the site system, ensuring compatibility across all telescopes. For SAT-MF1, `ocs` successfully acquires data from 69 agents with a total of 1524 HK fields at a rate of approximately 3.58 Mbps, as well as detector data from 7 wafers totaling 12,000 detectors.

We also described the data access software designed for efficiently exploring housekeeping data to contextualize detector behavior, eliminating the tedious task of arbitrarily loading .g3 files. We concluded with a description of HK data reading tools for both data packaging and providing aggregate information for a HK file, as well as the development of an improved HK data loading tool.

As of this writing, `hkdb` has been deployed to the site and is transitioning to become the default data loading mechanism. One potential improvement for G3tHK is to use `hkdb` to replace its existing data loader, allowing for quick database queries and data loading without excluding any agents for Level 3 data. This ensures that information about any agent can be accessed when needed, even if queried rarely. Another feature to consider is expanding the database to include observation-level summary information, making it easier to search the database using observation IDs while also retaining the ability to explore data outside of observation periods. Finally, there may be value in converting G3tHK into a PostgreSQL database for compatibility with the NFS.

5. APPENDIX A. ENCLOSURES NODE SETUP

Location	Node	Description
Azimuth Enclosure (SAT-MF1 Platform)	NUC (Linux)	Contains network and power interfaces to the observatory site.
DR Gas Handling System (SAT-MF1 Platform)	None	Contains the gas-handling system for the cryostat's DR insert, and includes electrical, vacuum, helium mixture, compressed air, and network interfaces.
Housekeeping Enclosure 1 (SAT-MF1 Platform)	None	Housekeeping 1 only contains two Lakeshore 372s for cryogenic thermometry readout and heater control for dilution refrigerator.
Housekeeping Enclosure 2 (SAT-MF1 Platform)	None	Housekeeping 2 contains 4 Lakeshore 240s, network interfaces and multifunction DAQ devices.
ATCA Crates (SAT-MF1 Platform)	None	The SMuRF electronics for detector readout, of which there are 2.
Coolant Control (SAT-MF1 Platform)	IFM IO-Link	Provides coolant flow to the ATCA crates as well as telescope receiver components. IFM IO-Link used to connect to the flowmeters in the enclosure to read out coolant temperature and flow.
CHWP Gripper Enclosures (SAT-MF1 Platform)	None	Together, they contain the 3 actuators for the cryostat's cold half-wave plate grippers, and include only electrical interfaces.
CHWP and Grid Loader Enclosure (SAT-MF1 Platform)	NUC (Linux)	Contains the electronics to operate and read out both the spinning cold half wave plate, and the sparse wire grid loader, and includes electrical and network interfaces.
Turbo Pump Enclosure (SAT-MF1 Platform)	None	Contains 2 turbo pumps used by the DR GHS to operate the cryostat's DR; includes electrical, vacuum, helium mixture, compressed air, and network interfaces.
Site Office	daq-satp1	Site-managed daq node for SAT-MF1 for Agent Management. Is the main node, where 36 of the 69 SAT-MF1 OCS Agents live.
Site Office	SMuRF Servers	Run the Agents and SMuRF-related software to readout SAT-MF1 detector arrays.
Site Container	NUC (Windows)	One of several containers at the site; this particular one houses the NUC that interfaces with SAT-MF1's dilution refrigerator. The Bluefors NUC is located on a rack near the large ACU control unit in this container which is used to communicate with SAT-MF1's telescope platform.

Table 3: Table listing the SAT-MF1 enclosure setup at the site, detailing the purpose of each enclosure and identifying those that contain nodes for network connections that cannot be supported directly by the `daq-satp1` node in the site office.

ACKNOWLEDGMENTS

This work was supported in part by a grant from the Simons Foundation (Award #457687, B.K.). This work was also supported by the National Science Foundation (UEI GM1XX56LEP58).

REFERENCES

- [1] The Simons Observatory Collaboration, “The simons observatory: science goals and forecasts,” *Journal of Cosmology and Astroparticle Physics* **2019**, 056–056 (feb 2019).
- [2] Nguyen, D. V., “The Simons Observatory: Alarms and Detector Quality Monitoring,” in [*International Society for Optics and Photonics*], SPIE (in press).
- [3] Koopman, B. J., Lashner, J., Saunders, L. J., Hasselfield, M., Bhandarkar, T., Bhimani, S., Choi, S. K., Duell, C. J., Galitzki, N., Harrington, K., Hincks, A. D., Ho, S.-P. P., Newburgh, L., Reichardt, C. L., Seibert, J., Spisak, J., Westbrook, B., Xu, Z., and Zhu, N., “The Simons Observatory: overview of data acquisition, control, monitoring, and computer infrastructure,” in [*Software and Cyberinfrastructure for Astronomy VI*], Guzman, J. C. and Ibsen, J., eds., **11452**, 1145208, International Society for Optics and Photonics, SPIE (2020).
- [4] Galitzki, N., Tsan, T., Spisak, J., Randall, M., Silva-Feaver, M., Seibert, J., Lashner, J., Adachi, S., Adkins, S. M., Alford, T., Arnold, K., Ashton, P. C., Austermann, J. E., Baccigalupi, C., Bazarko, A., Beall, J. A., Bhimani, S., Bixler, B., Coppi, G., Corbett, L., Crowley, K. D., Crowley, K. T., Day-Weiss, S., Dicker, S., Dow, P. N., Duell, C. J., Duff, S. M., Gerras, R. G., Groh, J. C., Gudmundsson, J. E., Harrington, K., Hasegawa, M., Healy, E., Henderson, S. W., Hubmayr, J., Iuliano, J., Johnson, B. R., Keating, B., Keller, B., Kiuchi, K., Kofman, A. M., Koopman, B. J., Kusaka, A., Lee, A. T., Lew, R. A., Lin, L. T., Link, M. J., Lucas, T. J., Lungu, M., Mangu, A., McMahon, J. J., Miller, A. D., Moore, J. E., Morshed, M., Nakata, H., Nati, F., Newburgh, L. B., Nguyen, D. V., Niemack, M. D., Page, L. A., Sakaguri, K., Sakurai, Y., Rao, M. S., Saunders, L. J., Shroyer, J. E., Sugiyama, J., Tajima, O., Takeuchi, A., Bua, R. T., Teply, G., Terasaki, T., Ullom, J. N., Lanen, J. L. V., Vavagiakis, E. M., Vissers, M. R., Walters, L., Wang, Y., Xu, Z., Yamada, K., and Zheng, K., “The simons observatory: Design, integration, and testing of the small aperture telescopes,” (2024).
- [5] Zhu, N., Bhandarkar, T., Coppi, G., Kofman, A. M., Orłowski-Scherer, J. L., Xu, Z., Adachi, S., Ade, P., Aiola, S., Austermann, J., Bazarko, A. O., Beall, J. A., Bhimani, S., Bond, J. R., Chesmore, G. E., Choi, S. K., Connors, J., Cothard, N. F., Devlin, M., Dicker, S., Dober, B., Duell, C. J., Duff, S. M., Dünner, R., Fabbian, G., Galitzki, N., Gallardo, P. A., Golec, J. E., Haridas, S. K., Harrington, K., Healy, E., Ho, S.-P. P., Huber, Z. B., Hubmayr, J., Iuliano, J., Johnson, B. R., Keating, B., Kiuchi, K., Koopman, B. J., Lashner, J., Lee, A. T., Li, Y., Limon, M., Link, M., Lucas, T. J., McCarrick, H., Moore, J., Nati, F., Newburgh, L. B., Niemack, M. D., Pierpaoli, E., Randall, M. J., Sarmiento, K. P., Saunders, L. J., Seibert, J., Sierra, C., Sonka, R., Spisak, J., Sutariya, S., Tajima, O., Teply, G. P., Thornton, R. J., Tsan, T., Tucker, C., Ullom, J., Vavagiakis, E. M., Vissers, M. R., Walker, S., Westbrook, B., Wollack, E. J., and Zannoni, M., “The simons observatory large aperture telescope receiver,” *The Astrophysical Journal Supplement Series* **256**, 23 (Sept. 2021).
- [6] Gudmundsson, J. E., Gallardo, P. A., Puddu, R., Dicker, S. R., Adler, A. E., Ali, A. M., Bazarko, A., Chesmore, G. E., Coppi, G., Cothard, N. F., Dachlythra, N., Devlin, M., Dünner, R., Fabbian, G., Galitzki, N., Golec, J. E., Patty Ho, S.-P., Hargrave, P. C., Kofman, A. M., Lee, A. T., Limon, M., Matsuda, F. T., Mausekopf, P. D., Moodley, K., Nati, F., Niemack, M. D., Orłowski-Scherer, J., Page, L. A., Partridge, B., Puglisi, G., Reichardt, C. L., Sierra, C. E., Simon, S. M., Teply, G. P., Tucker, C., Wollack, E. J., Xu, Z., and Zhu, N., “The simons observatory: modeling optical systematics in the large aperture telescope,” *Applied Optics* **60**, 823 (Jan. 2021).
- [7] McCarrick, H., Healy, E., Ahmed, Z., Arnold, K., Atkins, Z., Austermann, J. E., Bhandarkar, T., Beall, J. A., Bruno, S. M., Choi, S. K., Connors, J., Cothard, N. F., Crowley, K. D., Dicker, S., Dober, B., Duell, C. J., Duff, S. M., Dutcher, D., Frisch, J. C., Galitzki, N., Gralla, M. B., Gudmundsson, J. E., Henderson, S. W., Hilton, G. C., Ho, S.-P. P., Huber, Z. B., Hubmayr, J., Iuliano, J., Johnson, B. R., Kofman, A. M., Kusaka, A., Lashner, J., Lee, A. T., Li, Y., Link, M. J., Lucas, T. J., Lungu, M., Mates, J. A. B., McMahon, J. J., Niemack, M. D., Orłowski-Scherer, J., Seibert, J., Silva-Feaver, M., Simon, S. M., Staggs, S., Suzuki,

- A., Terasaki, T., Thornton, R., Ullom, J. N., Vavagiakis, E. M., Vale, L. R., Van Lanen, J., Vissers, M. R., Wang, Y., Wollack, E. J., Xu, Z., Young, E., Yu, C., Zheng, K., and Zhu, N., “The simons observatory microwave squid multiplexing detector module design,” *The Astrophysical Journal* **922**, 38 (Nov. 2021).
- [8] Koopman, B. J., “The Simons Observatory: Deployment of the Observatory Control System and supporting infrastructure,” in [*Software and Cyberinfrastructure for Astronomy VIII*], International Society for Optics and Photonics, SPIE (in press).
- [9] Henderson, S. W., Ahmed, Z., Austermann, J., Becker, D., Bennett, D. A., Brown, D., Chaudhuri, S., Cho, H.-M. S., D’Ewart, J. M., Dober, B., Duff, S. M., Dusatko, J. E., Fatigoni, S., Frisch, J. C., Gard, J. D., Halpern, M., Hilton, G. C., Hubmayr, J., Irwin, K. D., Karpel, E. D., Kernasovskiy, S. S., Kuenstner, S. E., Kuo, C.-L., Li, D., Mates, J. A. B., Reintsema, C. D., Smith, S. R., Ullom, J., Vale, L. R., Winkle, D. D. V., Vissers, M., and Yu, C., “Highly-multiplexed microwave SQUID readout using the SLAC Microresonator Radio Frequency (SMuRF) electronics for future CMB and sub-millimeter surveys,” in [*Millimeter, Sub-millimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*], Zmuidzinas, J. and Gao, J.-R., eds., **10708**, 1070819, International Society for Optics and Photonics, SPIE (2018).
- [10] Kernasovskiy, S. A., Kuenstner, S. E., Karpel, E., Ahmed, Z., Van Winkle, D. D., Smith, S., Dusatko, J., Frisch, J. C., Chaudhuri, S., Cho, H. M., Dober, B. J., Henderson, S. W., Hilton, G. C., Hubmayr, J., Irwin, K. D., Kuo, C. L., Li, D., Mates, J. A. B., Nasr, M., Tantawi, S., Ullom, J., Vale, L., and Young, B., “Slac microresonator radio frequency (smurf) electronics for read out of frequency-division-multiplexed cryogenic sensors,” *Journal of Low Temperature Physics* **193**, 570–577 (May 2018).
- [11] Sakuma, T., Guan, Y., Hasselfield, M., Koopman, B., Newburgh, L., and Nguyen, D., “Nextline.” <https://doi.org/10.5281/zenodo.11451619> (2024).
- [12] Guan, Y., “Simons Observatory: Observatory Scheduler and Automated Data Processing,” in [*Software and Cyberinfrastructure for Astronomy VIII*], International Society for Optics and Photonics, SPIE (in press).