

The Vera C. Rubin Observatory Data Butler and Pipeline Execution System

Tim Jenness¹, James F. Bosch², Andrei Salnikov³, Nate B. Lust², Nathan M. Pease³, Michelle Gower⁴, Mikolaj Kowalik⁴, Gregory P. Dubois-Felsmann⁵, Fritz Mueller³, and Pim Schellart²

¹Rubin Observatory Project Office, 950 N. Cherry Ave., Tucson, AZ 85719, USA

²Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA

³SLAC National Accelerator Laboratory, 2575 Sand Hill Rd., Menlo Park, CA 94025, USA

⁴NCSA, University of Illinois at Urbana-Champaign, 1205 W. Clark St., Urbana, IL 61801, USA

⁵IPAC, California Institute of Technology, MS 100-22, Pasadena, CA 91125, USA

ABSTRACT

The Rubin Observatory’s Data Butler is designed to allow data file location and file formats to be abstracted away from the people writing the science pipeline algorithms. The Butler works in conjunction with the workflow graph builder to allow pipelines to be constructed from the algorithmic tasks. These pipelines can be executed at scale using object stores and multi-node clusters, or on a laptop using a local file system. The Butler and pipeline system are now in daily use during Rubin construction and early operations.

Keywords: Data Management, Rubin Observatory, Legacy Survey of Space and Time, Databases

1. INTRODUCTION

The Vera C. Rubin Observatory’s Legacy Survey of Space and Time (LSST)¹ will image the entire southern sky every three days and consist of tens of petabytes of raw image data and associated calibration data. All these files must be tracked, along with the intermediate datasets and output products from pipeline processing, and depending on where the processing occurs the files will be stored on either POSIX file systems or object stores. The LSST Data Management System (DMS)² is responsible for transferring raw files off the mountain and storing them at the US Data Facility (USDF). The datasets are then processed by pipelines^{3,4} that can be run at different data centers and the results integrated into a unified data release. This paper will discuss the part of the DMS that abstracts data access from the pipeline algorithms, builds the execution workflow graphs, and allows the processing jobs to be run in large batch processing systems.

2. THE DATA BUTLER

The Data Butler (hereafter the “Butler”) is the system that abstracts the data access details from the pipeline developers. The top requirements for the Butler are:

- The pipeline writer should not have to know where the data are being read from or written to or what file formats are being used, or even if a file is involved at all.
- The pipeline writer only has to deal with Python objects.
- Users should be able to locate the relevant data using common astronomical concepts such as observation identifier, physical filter, patch of sky, or telescope.

The current implementation of the Butler is colloquially known as the “Generation 3” Butler since it is the third implementation during the evolution of the LSST DMS, a project that began nearly 20 years ago.⁵ The “Generation 2” Butler was used for many years but certain limitations in the design and implementation, especially in terms of its relationship with the evolving demands of the pipeline design and need for more flexibility, led to it being completely replaced with a brand new system that shares some concepts but no code with previous versions.

The two main components of the Butler are the Registry and the Datastore. Registry organizes datasets conceptually and associates them with astronomical concepts as stated above, but it has no idea where or how datasets are persisted. Datastore is responsible for serializing a Python object to a storage system and reading data back from that storage system and recreating the Python object. There is then a thin layer of code (the responsibility of the `Butler` Python class itself) providing a unified interface to the user that coordinates the Datastore and Registry interaction to ensure consistency.

2.1 Defining a Dataset

A cornerstone of the Butler design is that any dataset can be found by specifying a coordinate in its dimensional space. These dimensions are generally quantities that can be understood by an astronomer as being relevant for the particular dataset. For example, raw data can be addressed using dimensions `instrument`, `detector`, and `exposure`. A specific dataset must then be specified using a coordinate within that dimensional space, for example: `{instrument="LSSTCam", detector=1, exposure=2024050100023}` will uniquely represent the second detector of the 23rd LSSTCam observation taken on 2024-05-01.*

In order to determine which dimensions are relevant for a specific dataset, Butler has the concept of a “dataset type”. A dataset type gives a name to the combination of the list of relevant dimensions, an abstraction around the Python type it corresponds to, and whether or not it is a calibration dataset. A calibration dataset is a dataset that is associated with a data coordinate (for example an instrument and detector) but can be used for calibrating datasets, such as those used for instrument signature removal, that is additionally associated with a temporal validity range. The name of the dataset type itself does not matter, but it is important that the names are used consistently because they are critical when constructing pipelines. Raw data are commonly given the dataset type `raw` and the LSST science pipelines use dataset type names such as `diffim` (difference image), `defects` (pixel defect masks), and `bias` (processed biases), with the intent being for the name to be descriptive enough for a science user to have a good idea as to what type of data it represents.

The abstraction around the Python type is called a “storage class”. The Butler’s main job is to store Python objects somewhere and to retrieve them again. We use storage classes to map a name to a Python type and allow additional configuration to be associated with the Python type such as whether Butler can treat the type as a composite and what parameters can be included in the retrieval to modify the returned object. Composites can be important in that some pipelines may, for example, only need to take a WCS or a pixel mask as input and do not want the entire calibrated exposure to be loaded. A storage class can declare the components it allows and then declare the associated storage classes of each component.

The data model used by Butler is currently designed specifically for astronomical imaging data. The model itself is specified in a YAML text file and is therefore easily extended or modified to suit specific needs. There is a tension between trying to develop one unified model that will support all astronomy data versus an approach where every observatory specifies their own dimension universe for every instrument. A unified model is needed in scenarios where datasets from different instruments should be processed in the same pipeline. There is, though, no problem with multiple distinct Butler repositories being configured with their own dimension universes when the data are very different and then requiring the user to choose the relevant Butler repository. Currently there are two projects experimenting with the Butler using non-imaging survey data, and both Subaru’s Prime Focus Spectrograph (PFS) team⁶ and NASA’s SPHEREx team⁷ have decided to alter the default dimension universe data model.

*Embedding the date in the integer exposure ID is an LSSTCam convention and is not required. Any integer can be used here so long as it uniquely identifies a specific observation.

2.2 Collections

A dataset can be specified by its data coordinate and dataset type, but the Butler needs one final piece of information to uniquely locate the desired dataset. In order to allow different pipeline execution runs to store their outputs without over-writing datasets from previous runs, each dataset must be associated with a single RUN collection. Every dataset tracked in a Butler Registry can be uniquely located by knowing the run collection, data coordinate, and dataset type.

Whilst every dataset must be stored in a RUN collection there are also other types of collections that can be used:

- A **TAGGED** collection is a collection of arbitrary datasets.
- A **CHAINED** collection is a collection of other collections and those other collections can be of any type (chains of chains are allowed). A **CHAINED** collection does not itself contain datasets. The order of collections in a chain is important and determines how datasets are found.
- A **CALIBRATION** collection is a special type of collection where a calibration dataset (one whose dataset type is marked as a calibration) can be associated with a validity timespan.

All these collections are important when building and executing pipelines.

2.3 The Registry

Registry can be thought of as a database that can be queried to find out what datasets are available. We define Registry to have an abstract interface to allow different implementations to co-exist, although we do assume that a SQL-like syntax is supported when constraining queries (effectively a **WHERE** clause), even if that does not result in SQL being executed in the specific Registry implementation.

The primary Registry implementation uses SQLAlchemy^{8†} to talk to PostgreSQL[‡] or SQLite[§] backend databases. We do not make use of the Object Relational Mapping interface, and instead use the low-level SQLAlchemy “Core” interface to create and interact with tables.

The Registry stores all the dimension information, collection and dataset type definitions, and the associations between these and the actual datasets. All dimension values must be pre-defined before a dataset can be stored with that value, and the dimension records have additional metadata that can be used to constrain dataset queries or to provide more detail as to what a dimension really represents. For example, in the default schema, which is currently used by LSST cameras, Hyper Suprime-Cam⁹ and DECam,¹⁰ the **exposure** dimension contains metadata including the time of the exposure, the elevation and azimuth, the observation type, and the exposure time, whereas the **detector** dimension contains metadata including the full name, the raft name, and the detector role. The tables that hold this dimension metadata exist independently of the datasets they correspond to; for example, while the metadata associated with the **exposure** dimension metadata is usually derived from the header of a **raw** dataset, that **exposure** is also associated with the dataset types that result from processing done on that exposure. The dimension tables – which often also include foreign key relationships, such as the **physical_filter** associated with an **exposure** – thus form a data model “skeleton” of sorts for the actual datasets, which have no relationships of their own. In some cases those dimension relationships are spatial (some dimensions are associated with regions on the sky) or temporal (associated with time spans).

These dimension relationships allow the Registry to provide a rich query system, based on a custom expression language based on SQL boolean expressions. These are parsed and translated by Registry into actual SQL queries, and we delegate optimization and execution of these queries almost entirely to the underlying database. But in contrast to raw SQL the user essentially has to only provide the Registry the equivalent of the **WHERE** clause, because the **SELECT** clause is based on what dataset types or dimensions are being requested, and the multi-join **FROM** clause can be worked out internally according to the known dimension relationships, despite those

[†]<https://www.sqlalchemy.org>

[‡]<https://www.postgresql.org>

[§]<https://www.sqlite.org/index.html>

relationships being encoded in YAML configuration rather than code. The query system also isolates the actual database schema from the user, allowing it to be changed for performance reasons without breaking user code.

2.4 The Datastore

Datastores are responsible for serializing a Python object to a storage system and reading the data back from the storage system and returning a Python object. On `put`, Datastore is given a reference to the dataset (encapsulated in the Python class `DatasetRef`) by Registry (this is the data coordinate, the dataset type, the run collection, and the unique identifier, a UUID,¹¹ assigned by Registry) and the Python object. To retrieve a dataset the Datastore is passed in the dataset reference. By design the only connection between Registry and Datastore is the dataset reference, although in some implementations the Datastore can make use of the shared PostgreSQL database to store relevant information.

The Butler user does not have to know how data are serialized or where it is coming from, indeed there is no requirement for files to be involved at all. In the current system there are three Datastores defined: `FileDatastore` serializes to files; `InMemoryDatastore` stores Python objects in an in-memory cache; and `ChainedDatastore` is a Datastore consisting of other Datastores. All Datastores support configuration-based constraints that can be used to decide whether a specific dataset type or storage class should be accepted or rejected by the Datastore.

2.4.1 File Datastore

The `FileDatastore` is the most fundamental of the Datastore implementations in that it is responsible for serializing datasets to files and reading them back in again. A file Datastore is defined by a URI pointing to the area where files are under the control of the Datastore. A URI is used rather than a file path to allow file I/O itself to be abstracted to allow use of remote storage or POSIX file systems. To achieve this abstraction we use a unified URI handler class, `ResourcePath` from the `lsst-resources` package¹, which currently supports S3 storage, Google Cloud Storage, WebDAV, and POSIX, as well as new storage systems which can easily be added.

Within the file Datastore the class that is responsible for creating the file and reading the file is called the `Formatter`. The Datastore configuration system includes a look up table that matches the dataset type or storage class of the relevant dataset to a corresponding formatter implementation. This formatter is then given the Python object and the destination location and told to write the file. Similarly, when a dataset is requested, the formatter is looked up in the internal Datastore database and the formatter is told to read the file. We do not look up the formatter from configuration when reading since it is possible that the configuration may have changed since the file was stored.

Whereas the Registry only deals with datasets, the Datastore understands composites. When a dataset is associated with a composite storage class, the Datastore can be configured to disassemble the dataset and write the contents as distinct files. For example, an astronomical image could be disassembled into the pixel image data, a variance plane, a mask, and a header before being written out as four files. If a user requests just the header component, that file can be read without needing to look at the other files, and if the full dataset is requested all components will be read and the composite will be reassembled using the helper class declared in the storage class definition. Of course, if a user requests a component when the dataset was not disassembled, this will still work, although it may require the entire file be downloaded from a remote object store in order to extract a small subset.

When the datasets are stored on a slow network disk or a remote object store, it is inefficient to continually re-read the files if the pipeline algorithm is reading subsets or if the user is requesting individual components. The file Datastore overcomes this by implementing a caching system where the file is stored in a more local cache directory on first retrieval. The cache can expire files based on number of files, total cache size, or number of datasets (noting that a disassembled composite will consist of multiple files). In some situations this can significantly improve performance with remote Datastores.

¹<https://github.com/lsst/resources>

2.4.2 In-Memory Datastore

This is an in-memory caching Datastore. Since it caches Python objects in a single process its main purpose is to support intermediate products that are to be passed from one task to another without needing to include the overhead for serialization. In most scenarios an in-memory Datastore is combined with another Datastore in a chain. Care must be taken when returning datasets from the in-memory Datastore since Python can not prevent the caller from modifying the object, which would cause confusion if the object is retrieved a second time. Options being considered are to always deep copy the object on return, or consider allowing the Datastore to be configured to remove the object from the cache when retrieved.

2.4.3 Chained Datastore

The chained Datastore implementation does not itself store any datasets. It consists of multiple Datastores, each with its own configuration and allowed to be of any type. When writing a dataset the dataset is presented to each Datastore in turn and everything is okay so long as one Datastore accepts the dataset. When reading a dataset each Datastore is asked for the dataset in turn and the Python object is returned from the first matching Datastore. Combined with per-Datastore dataset type constraints, this can allow some datasets to be stored in an in-memory Datastore and file Datastore but allow the in-memory dataset to be retrieved efficiently, whilst other datasets are only stored in the file Datastore. Alternatively a chain of two file Datastores can be used to allow one Datastore to be read-only (for example a Datastore containing a validated self-contained data release) and a second Datastore to accept derived products from users.

2.4.4 Other Datastores

The Datastore interface does not require datasets to be persisted as files. For example, we are considering storing pipeline metrics directly into our metrics database.^{12,13} The Butler user would not know whether a metric was being stored as a JSON file in a file Datastore or stored directly into a metrics database (or even stored in both places).

2.5 Client/Server Butler

The Butler is a Python library that, by default, is set up to use SQLAlchemy to talk to a SQL Registry, and will use AWS or Google Cloud credentials to talk to object stores. This is not a convenient interface if Python is not being used, or if the authorization required to access the database or object store directly is not available to the people trying to access the Butler. The latter situation is likely to be the situation in the Rubin Science Platform in the cloud where users will be logged in with Rubin accounts but will not be issued cloud credentials.^{14,15}

To solve this problem we are working on a `https` client/server Butler.¹⁶ The client user will present their authentication token and the server will then determine if that user is authorized to retrieve the requested dataset. If they are allowed, they will be returned a signed URL that their client can then use to retrieve the dataset.

2.6 Command Line Tooling

There are many common operations that should be achievable without needing to write any Python code. To serve those use cases we provide command-line tooling based on the `click` Python package.^{||} Examples of the current subcommands and options are shown in Fig. 1. We have defined a pluggable architecture where packages other than the core `daf_butler` package can register their own Butler subcommands. This allows, for example, specialist commands such as `ingest-photodiodes` and `make-discrete-skymap` to appear in the subcommand listing even though they are not at all generic functionality.

^{||}<https://click.palletsprojects.com>

```

$ butler --help
Options:
  --log-level LEVEL|COMPONENT=LEVEL ...
                                The logging level. Without an explicit
                                logger name, will only affect the default
                                root loggers (lsst). To modify the root
                                logger use '.=LEVEL'. Supported levels are [
                                CRITICAL|ERROR|WARNING|INFO|VERBOSE|DEBUG|TR
                                ACE]
  --long-log
                                Make log messages appear in long format.
  --log-file FILE ...
                                File(s) to write log messages. If the path
                                ends with '.json' then JSON log records will
                                be written, else formatted text log records
                                will be written. This file can exist and
                                records will be appended.
  --log-tty / --no-log-tty
                                Log to terminal (default). If false logging
                                to terminal is disabled.
  --log-label TEXT ...
                                Keyword=value pairs to add to MDC of log
                                records.
  --progress / --no-progress
                                Show a progress bar for slow operations when
                                possible.
  -h, --help
                                Show this message and exit.

Commands:
  associate
                                Add existing datasets to a tagged collection.
  certify-calibrations
                                Certify calibrations in a repository.
  collection-chain
                                Define a collection chain.
  config-dump
                                Dump butler config to stdout.
  config-validate
                                Validate the configuration files.
  convert
                                Convert a gen2 repo to gen3.
  create
                                Create an empty Gen3 Butler repository.
  define-visits
                                Define visits from exposures.
  export-calibs
                                Export calibrations from the butler for later import.
  import
                                Import data into a butler repository.
  ingest-files
                                Ingest files from table file.
  ingest-photodiode
                                Ingest photodiode data.
  ingest-raws
                                Ingest raw frames.
  make-discrete-skymap
                                Define a discrete skymap from calibrated exposures.
  prune-datasets
                                Remove datasets.
  query-collections
                                Search for collections.
  query-data-ids
                                List the data IDs in a repository.
  query-dataset-types
                                Get the dataset types in a repository.
  query-datasets
                                List the datasets in a repository.
  query-dimension-records
                                Query for dimension information.
  register-dataset-type
                                Register a new dataset type with this...
  register-dcr-subfilters
                                Add subfilters for chaotic modeling.
  register-instrument
                                Add an instrument definition to the repository
  register-skymap
                                Make a SkyMap and add it to a repository.
  remove-collections
                                Remove one or more non-RUN collections.
  remove-dataset-type
                                Remove a dataset type definition from a repository.
  remove-runs
                                Remove one or more RUN collections.
  retrieve-artifacts
                                Retrieve file artifacts from a Butler.
  transfer-datasets
                                Transfer datasets from one butler to another.
  write-curated-calibrations
                                Add an instrument's curated calibrations.

```

Figure 1. The butler command line options and subcommands.

2.7 Data Ingest

A Butler repository is not very useful without containing any datasets. Facilities are provided for ingesting externally-generated files, and for the generic ingesting tooling** care must be taken to define the correct data coordinates for each dataset and to ensure that those dimension values have already been defined. The `obs_base` package†† provides a command specifically targeted at ingesting raw imaging data. The software scans a directory tree for files matching the provided glob, uses the `astro_metadata_translator`‡‡ infrastructure to read the headers and translate the contents to a standard form, creates `exposure` records as needed, and then ingests the files by constructing a data coordinate from the translated metadata. It can handle raw files that store all the detector images in a single file or as one file per detector, and Butler does allow a single file artifact to be associated with multiple datasets in the Registry. In situations where the files may be ingested into Butler repositories multiple times, a facility is also provided to use “sidecar” JSON files containing extracted metadata, or even a JSON index file containing metadata for multiple files, since it is much quicker to parse JSON than to extract and parse a FITS header, especially if that FITS file is in an object store.

2.8 Calibrations

As noted above, a `CALIBRATION` collection is a special type of collection that associates datasets with a validity range. When a calibration is requested the data coordinate of the dataset being calibrated is used and any time-based dimensions (such as `exposure`) specify which calibration should be chosen.

Not all calibrations are calculated from datasets available to a Butler. For example, for LSSTCam the QE curves are calculated by the camera team and given to the pipelines team as part of the camera delivery. Additionally, some calibrations, such as defect masks, are relatively static, and these can be represented in a compact text format, and are useful to people outside of a Butler. We call such datasets “curated calibrations” and store them in Git repositories using text file formats and a directory layout that makes the validity ranges clear. It can be very convenient when setting up a new Butler for an instrument to be able to seed it with these curated calibrations without having to locate raw data files and reconstruct them from scratch, and infrastructure is provided to enable that, for example using the `butler write-curated-calibrations` command line tool. Generally, when these curated calibrations are stored the Butler converts them from text format to a binary format to make reading them more efficient; this is all handled by the formatter configured for the particular Datastore.

3. THE PIPELINE SYSTEM

Processing data is done with reusable units of code called `Tasks`. Each `Task` has a specialized configuration object attached to it (from the `pex_config` package*) and must provide a `run()` method that is the method that implements the algorithm. `Tasks` can be nested within each other in a hierarchy to create high level algorithms. However, the top level `Task` of any algorithm must satisfy an additional interface, defined as a `PipelineTask`, which allows it to interact with the `Butler` for i/o and ordering within a processing pipeline.

A `PipelineTask` is special because it requires the author to declare the task’s “connections”: the datasets the task will consume as inputs and produce as outputs. The interface also requires the definition of the task’s dimensions, which defines the unit of processing over which the task runs. The task’s dimensions do not need to match those of its input or output datasets; for example, a task with dimensions `{instrument, exposure}` that takes an input with dimensions `{instrument, exposure, detector}` is a “gather” step that processes all detectors from that exposure together, but processes each exposure independently (and hence possibly in parallel). Other examples of `PipelineTask` dimensions are `detector` and `exposure` which operates on data from an exposure of a single detector taken from the telescope, whilst another could be `band`, `tract`, and `patch` (sky regions) which would operate on many intermediate data products produced from previous processing stages that overlap the specified region. The input and output datasets are tied in to the Butler by declaring the dataset

**`butler ingest-files`.

††https://github.com/lsst/obs_base

‡‡<https://astro-metadata-translator.lsst.io>

*https://github.com/lsst/pex_config


```

description: A demo pipeline.
instrument: lsst.obs.subaru.HyperSuprimeCam
tasks:
  calibrate:
    class: lsst.pipe.tasks.calibrate.CalibrateTask
    config:
      astrometry.matcher.maxOffsetPix: 300
  characterizeImage: lsst.pipe.tasks.characterizeImage.CharacterizeImageTask
  isr:
    class: lsst.ip.isr.IsrTask
    config:
      doVignette: true
      vignetteValue: 0.0

```

Figure 2. An example minimalist HSC pipeline describing single frame processing.

type name and associated storage class. The storage class is a proxy for the Python type and thus allows the pipeline author to guarantee that the correct Python type will be given (converting it from another type if required and supported) – if the storage classes associated with datasets of that dataset type in the target Butler are not compatible with those required by the pipeline the pipeline will not run.

To support pipeline execution the `run()` method of a `PipelineTask` must take parameters that match the input connections and must return a dict-like data structure where the keys match the expected outputs. A `PipelineTask` then has enough information to be able to pull the required datasets from a Butler and store the outputs, although task authors can override that default behavior using the associated config class to allow for more complexity in the Butler interaction. Connections have options to allow the dataset loading to be deferred (something that is important when co-adding hundreds of images) or to allow multiple datasets to be provided if the dimensionality implies that.

Each `PipelineTask` provides one data processing step, several of which may be integrated together into a processing pipeline. We define pipelines using a YAML text file; an example is shown in Fig. 2. The example pipeline consists of three tasks, with labels `calibrate`, `characterizeImage`, and `isr`, referencing the corresponding `PipelineTask` Python classes. The order listed in the YAML file is not important because the classes implementing each task know which dataset types they need and the pipeline builder arranges them such that the output connections of one task are associated with the corresponding input connections of one or more other tasks. Each task can have configurations specified in the pipeline file which override the default values, or overrides may be given through the pipeline executor command line interface. This configuration not only supports algorithmic values, but also allows configuring the names used for the datasets in a task’s connections. The ability to configure dataset types provides flexibility to reuse the same `PipelineTask` with different configurations multiple times in a pipeline, with each configuration outputting a distinct dataset. Additionally, it also provides flexibility in re-using pipelines. Different pipelines can be derived from a common base pipeline, but tasks can be added or removed, and the datasets in the task connections can be configured such all the tasks can still be joined into a connected pipeline graph. When a pipeline specifies an instrument explicitly, the `Instrument` class can be used to apply specific configuration overrides to any of the tasks in the pipeline. Pipeline definitions can be significantly more complex than in the example. To simplify this process, pipelines may be built up using the common recipe and ingredients paradigm^{17,18} to allow for pipeline reuse. It is also possible for a user to specify that a subset of a pipeline be executed rather than the entire thing.

Fig. 3a shows a visualization of the demonstration pipeline shown in Fig. 2. The pipeline graph shows all the input and output datasets (including calibrations) and the order of execution. What it does not show is any specific datasets attached to each step. To determine what is to be processed a graph building algorithm queries a Butler and allocates datasets to tasks. All of the datasets which correspond to a unique set of task dimensions are bundled together into a unit of work that we call a “quantum”. The set of all quanta to be executed and the relations between them and their tasks are called a “Quantum Graph.” The user provides an input Butler collection to search for datasets, in addition to a Registry query expression that affects not

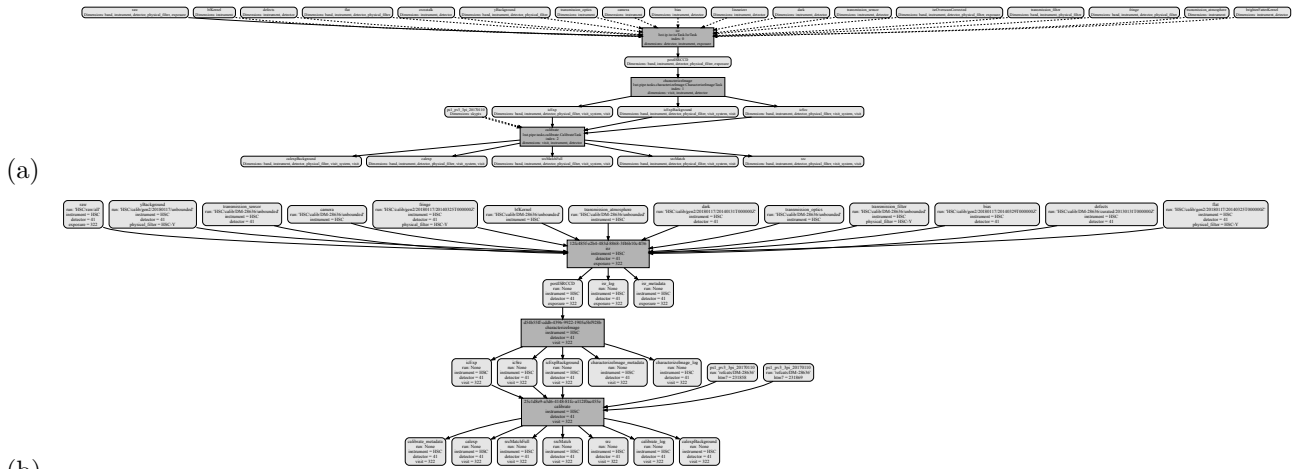


Figure 3. Top is a visualization of the pipeline shown in Fig. 2. Bottom is the Quantum Graph for a single frame processing of one detector from one HSC observation showing the same pipeline with specific datasets.

just the input datasets but the dimensions of the tasks and output datasets as well. Having the Registry’s dimension table skeleton populated before the Quantum Graph is generated allows us to query the database for both input datasets and predict future output datasets in much the same way. The graph builder then allocates each matching input dataset to a specific Quantum and determines all the expected datasets that will be created during pipeline processing. During processing the outputs are written to a `RUN` collection along with additional datasets containing all the log messages written by a task, software package versions, and processing metadata (including execution times and CPU usage for `Tasks`). In normal usage we then create a new `CHAINED` collection that combines the input collections and the new output `RUN` collection to allow a single collection to be queried for all the input and output datasets in that processing. This procedure produces chains of collections, as processing campaigns proceed, which represent the final data products. The Quantum Graph itself provides provenance information and is tied to the datasets by preserving the predicted output dataset UUIs in the Butler when they are stored during processing.

4. INTEGRATION WITH BATCH SYSTEMS

A Quantum Graph describes the work that the pipeline has to do in order to process all the data, but it does not say how the work should be scheduled or any resources that are required. To run a Quantum Graph it has to be converted to something that can be executed. On a single compute node (including a laptop) the `ctrl_mpexec` package[†] provides a multi-processing executor, commonly run via the `pipetask` command line, that can run every quantum in the correct order and track all the responses.

When a Quantum Graph consists of thousands of quanta, a single compute node is no longer sufficient. At that scale we need to use workflow systems such as HTCondor,¹⁹ Pegasus,²⁰ Parsl,²¹ or PanDA.²² For these workflow systems the Quantum Graph must be translated into the form expected by the target system. Multiple prototypes were developed in parallel^{23–25} but it was soon realized that many of the concepts for translating a Quantum Graph to a specialist graph are common and we developed a new Batch Production System (`ctrl_bps`)[‡] that would provide a framework to allow the aforementioned workflow systems to be treated as plugin code.

The BPS system provides a facility for a Quantum Graph to be converted to a generic form supporting workflow system concepts, called the Workflow Graph, before being converted to the final form. This intermediate graph includes submission configuration parameters such as those that specify special resource requirements (CPU or memory) required by specific tasks. The BPS submission process also allows quanta to be clustered for efficiency. Some quanta can be quick, executing in tens of seconds, and not all batch systems are designed

[†]https://github.com/lst/ctrl_mpexec

[‡]https://github.com/ctrl_bps

to efficiently handle such small payloads. Clustering allows for multiple quanta to be combined into a single execution job; this can be particularly efficient if the clustering is designed such that the output of the first quantum is the only input to the next quantum (which is a common scenario for the early stages of many pipelines when a single detector is being characterized).

In the initial experiments, each quantum would talk to the main Butler repository when retrieving data at the start of the quantum and then when writing the results out at the end. It soon became clear that this approach was not tenable once thousands of jobs are running concurrently and they are all trying to ask the Registry where their specific datasets are located. We fixed this problem by changing the way that batch jobs interact with a Butler. Initially we created a standalone read-only SQLite Butler Registry containing the relevant information for the entire workflow. That did fix the problem with the large numbers of simultaneous connections and updates. However, as a SQLite file it had to be copied to each job. For very large workflows the SQLite file became large enough that the copying caused its own problems on some systems and at minimum caused time delays. A better solution was to use the Quantum Graph itself since the graph can be stored in a shared location and only the appropriate subset of the graph needs to be read by each job. The Quantum Graph knows all the input datasets and all the expected output datasets; it is therefore possible for a Butler to be constructed that uses the Quantum Graph itself as the Registry whilst using the main Datastore. To enable this it was necessary to transfer Datastore records for existing datasets into the Quantum Graph. Running large workflows in this manner means that the main Registry is not involved at all during the bulk of the processing. Once the main workflow completes, a final job is executed that checks the Datastore to see which of the expected files were produced, and then ensures that the associated Registry entries are transferred back to the main Registry. This merge job runs even if the workflow fails, to ensure that all the datasets that were generated are known to the main Registry. These changes significantly improved the scalability of the bulk processing. Note though that in this mode we are still writing to the main Datastore file system or object store and not using any file management facilities of the workflow system. This decision is not hard-coded into BPS though and is entirely controlled by configuration. In the future, we may decide to have jobs read and write to a job-local file system Datastore (whilst matching all the formatter and file naming configuration) and have the workflow system move files between the main Datastore and the job's Datastore.

5. DEVELOPMENT PROCESS

All the code described in this paper is written in Python (requiring at least Python version 3.9²⁶) and is open-source using the BSD 3-clause license for some packages and GPLv3 license for others. The code is available from GitHub at <https://github.com/lstt> and we follow the Rubin Observatory Data Management development process.^{*27} We use the `black` [†] tool to automatically format the code, along with `isort` [‡] to order the Python imports.

Much of the code uses Python type annotations which are verified using the Mypy package.[§] We had a fairly large amount of code written before we decided to use type annotations. It is known to be difficult to add type annotations to an existing project,²⁸ and that was our experience. There are still parts of the system that lack annotations and in some places it is extremely difficult to add them. In particular, the more flexible an API is the harder annotations become, and this does begin to drive API design decisions. We have, though, found that annotations do help once they have been added, and in particular code refactoring is less dangerous.

Documentation is built using Sphinx, based on the `documenteer` tooling[¶] and is integrated into the LSST Science Pipelines documentation.^{||}

We use GitHub Actions in all the repositories to ensure that code is formatted correctly, the type annotations are correct, and that all the tests pass and documentation builds. We also automatically publish all the packages

*developer.lsst.io

†<https://black.readthedocs.io/>

‡<https://pycqa.github.io/isort/>

§<https://mypy-lang.org>

¶<https://documenteer.lsst.io>

||<https://pipelines.lsst.io>

to the Python Package Index** when a tag is added to a repository. These packages do not require the entire LSST Science Pipelines software system to be installed.

6. CONCLUSIONS

In this paper we have described a flexible system for abstracting data access from pipeline algorithmic code and for constructing complex pipelines. The Butler system allows scientists and pipelines to have no knowledge of file formats or data locations, and allows for pipelines to be built in a manner where the pipeline builder can determine what data will be processed and how the individual algorithmic tasks will be combined. This system has been successfully demonstrated at scale with the reprocessing of the DESC DC2 data²⁹ on the LSST Interim Data Facility at Google¹⁵ as part of LSST Data Preview 0.2³⁰ during late 2021 and early 2022.

ACKNOWLEDGMENTS

We thank Eli Rykoff and Kian-Tat Lim for their reviews of the manuscript. We also thank Kian-Tat Lim for his advice and feedback during the development of this system. This material or work is supported in part by the National Science Foundation through Cooperative Agreement AST-1258333, Cooperative Support Agreement AST-1202910, and Cooperative Support Agreement AST-1836783 managed by the Association of Universities for Research in Astronomy (AURA), and the Department of Energy under Contract No. DE-AC02-76SF00515 with the SLAC National Accelerator Laboratory managed by Stanford University. Additional Rubin Observatory funding comes from private donations, grants to universities, and in-kind support from LSSTC Institutional Members.

REFERENCES

- [1] Ivezić, Ž., Kahn, S. M., Tyson, J. A., Abel, B., Acosta, E., Allsman, R., Alonso, D., AlSayyad, Y., Anderson, S. F., Andrew, J., Angel, J. R. P., Angeli, G. Z., Ansari, R., Antilogus, P., Araujo, C., Armstrong, R., Arndt, K. T., Astier, P., Aubourg, É., Auza, N., Axelrod, T. S., Bard, D. J., Barr, J. D., Barrau, A., Bartlett, J. G., Bauer, A. E., Bauman, B. J., Baumont, S., Bechtol, E., Bechtol, K., Becker, A. C., Becla, J., Beldica, C., Bellavia, S., Bianco, F. B., Biswas, R., Blanc, G., Blazek, J., Blandford, R. D., Bloom, J. S., Bogart, J., Bond, T. W., Booth, M. T., Borgland, A. W., Borne, K., Bosch, J. F., Boutigny, D., Brackett, C. A., Bradshaw, A., Brandt, W. N., Brown, M. E., Bullock, J. S., Burchat, P., Burke, D. L., Cagnoli, G., Calabrese, D., Callahan, S., Callen, A. L., Carlin, J. L., Carlson, E. L., Chandrasekharan, S., Charles-Emerson, G., Chesley, S., Cheu, E. C., Chiang, H.-F., Chiang, J., Chirino, C., Chow, D., Ciardi, D. R., Claver, C. F., Cohen-Tanugi, J., Cockrum, J. J., Coles, R., Connolly, A. J., Cook, K. H., Cooray, A., Covey, K. R., Cribbs, C., Cui, W., Cutri, R., Daly, P. N., Daniel, S. F., Daruich, F., Daubard, G., Daues, G., Dawson, W., Delgado, F., Dellapenna, A., de Peyster, R., de Val-Borro, M., Digel, S. W., Doherty, P., Dubois, R., Dubois-Felsmann, G. P., Durech, J., Economou, F., Eifler, T., Eracleous, M., Emmons, B. L., Fausti Neto, A., Ferguson, H., Figueroa, E., Fisher-Levine, M., Focke, W., Foss, M. D., Frank, J., Freemon, M. D., Gangler, E., Gawiser, E., Geary, J. C., Gee, P., Geha, M., Gessner, C. J. B., Gibson, R. R., Gilmore, D. K., Glanzman, T., Glick, W., Goldina, T., Goldstein, D. A., Goodenow, I., Graham, M. L., Gressler, W. J., Gris, P., Guy, L. P., Guyonnet, A., Haller, G., Harris, R., Hascall, P. A., Haupt, J., Hernandez, F., Herrmann, S., Hileman, E., Hoblitt, J., Hodgson, J. A., Hogan, C., Howard, J. D., Huang, D., Huffer, M. E., Ingraham, P., Innes, W. R., Jacoby, S. H., Jain, B., Jammes, F., Jee, M. J., Jenness, T., Jernigan, G., Jevremović, D., Johns, K., Johnson, A. S., Johnson, M. W. G., Jones, R. L., Juramy-Gilles, C., Jurić, M., Kalirai, J. S., Kallivayalil, N. J., Kalmbach, B., Kantor, J. P., Karst, P., Kasliwal, M. M., Kelly, H., Kessler, R., Kinnison, V., Kirkby, D., Knox, L., Kotov, I. V., Krabbenham, V. L., Krughoff, K. S., Kubánek, P., Kuczewski, J., Kulkarni, S., Ku, J., Kurita, N. R., Lage, C. S., Lambert, R., Lange, T., Langton, J. B., Le Guillou, L., Levine, D., Liang, M., Lim, K.-T., Lintott, C. J., Long, K. E., Lopez, M., Lotz, P. J., Lupton, R. H., Lust, N. B., MacArthur, L. A., Mahabal, A., Mandelbaum, R., Markiewicz, T. W., Marsh, D. S., Marshall, P. J., Marshall, S., May, M., McKercher, R., McQueen, M., Meyers, J., Migliore, M., Miller, M., Mills, D. J., Miraval, C., Moeyens, J., Moolekamp, F. E., Monet, D. G., Moniez,

**For example the Butler can be found at <https://pypi.org/project/lsst-daf-butler/>.

- M., Monkewitz, S., Montgomery, C., Morrison, C. B., Mueller, F., Muller, G. P., Muñoz Arancibia, F., Neill, D. R., Newbry, S. P., Nief, J.-Y., Nomerotski, A., Nordby, M., O'Connor, P., Oliver, J., Olivier, S. S., Olsen, K., O'Mullane, W., Ortiz, S., Osier, S., Owen, R. E., Pain, R., Palecek, P. E., Parejko, J. K., Parsons, J. B., Pease, N. M., Peterson, J. M., Peterson, J. R., Petravick, D. L., Libby Petrick, M. E., Petry, C. E., Pierfederici, F., Pietrowicz, S., Pike, R., Pinto, P. A., Plante, R., Plate, S., Plutchak, J. P., Price, P. A., Prouza, M., Radeka, V., Rajagopal, J., Rasmussen, A. P., Regnault, N., Reil, K. A., Reiss, D. J., Reuter, M. A., Ridgway, S. T., Riot, V. J., Ritz, S., Robinson, S., Roby, W., Roodman, A., Rosing, W., Roucelle, C., Rumore, M. R., Russo, S., Saha, A., Sassolas, B., Schalk, T. L., Schellart, P., Schindler, R. H., Schmidt, S., Schneider, D. P., Schneider, M. D., Schoening, W., Schumacher, G., Schwamb, M. E., Sebag, J., Selvy, B., Sembroski, G. H., Seppala, L. G., Serio, A., Serrano, E., Shaw, R. A., Shipsey, I., Sick, J., Silvestri, N., Slater, C. T., Smith, J. A., Smith, R. C., Sobhani, S., Soldahl, C., Storr-Lombardi, L., Stover, E., Strauss, M. A., Street, R. A., Stubbs, C. W., Sullivan, I. S., Sweeney, D., Swinbank, J. D., Szalay, A., Takacs, P., Tether, S. A., Thaler, J. J., Thayer, J. G., Thomas, S., Thornton, A. J., Thukral, V., Tice, J., Trilling, D. E., Turri, M., Van Berg, R., Vanden Berk, D., Vetter, K., Virieux, F., Vucina, T., Wahl, W., Walkowicz, L., Walsh, B., Walter, C. W., Wang, D. L., Wang, S.-Y., Warner, M., Wiecha, O., Willman, B., Winters, S. E., Wittman, D., Wolff, S. C., Wood-Vasey, W. M., Wu, X., Xin, B., Yoachim, P., and Zhan, H., "LSST: From Science Drivers to Reference Design and Anticipated Data Products," *ApJ* **873**, 111 (Mar 2019).
- [2] Jurić, M., Kantor, J., Lim, K. T., Lupton, R. H., Dubois-Felsmann, G., Jenness, T., Axelrod, T. S., Aleksić, J., Allsman, R. A., AlSayyad, Y., Alt, J., Armstrong, R., Basney, J., Becker, A. C., Becla, J., Biswas, R., Bosch, J., Boutigny, D., Kind, M. C., Ciardi, D. R., Connolly, A. J., Daniel, S. F., Daves, G. E., Economou, F., Chiang, H. F., Fausti, A., Fisher-Levine, M., Freemon, D. M., Gris, P., Hernandez, F., Hoblitt, J., Ivezić, Z., Jammes, F., Jevremović, D., Jones, R. L., Kalmbach, J. B., Kasliwal, V. P., Krughoff, K. S., Lurie, J., Lust, N. B., MacArthur, L. A., Melchior, P., Moeyens, J., Nidever, D. L., Owen, R., Parejko, J. K., Peterson, J. M., Petravick, D., Pietrowicz, S. R., Price, P. A., Reiss, D. J., Shaw, R. A., Sick, J., Slater, C. T., Strauss, M. A., Sullivan, I. S., Swinbank, J. D., Van Dyk, S., Vujčić, V., Withers, A., and Yoachim, P., "The LSST Data Management System," in [*Astronomical Data Analysis Software and Systems XXV*], Lorente, N. P. F., Shortridge, K., and Wayth, R., eds., *ASP Conf. Ser.* **512**, 279 (Dec 2017).
- [3] Bosch, J., AlSayyad, Y., Armstrong, R., Bellm, E., Chiang, H.-F., Eggl, S., Findeisen, K., Fisher-Levine, M., Guy, L. P., Guyonnet, A., Ivezić, Ž., Jenness, T., Kovács, G., Krughoff, K. S., Lupton, R. H., Lust, N. B., MacArthur, L. A., Meyers, J., Moolekamp, F., Morrison, C. B., Morton, T. D., O'Mullane, W., Parejko, J. K., Plazas, A. A., Price, P. A., Rawls, M. L., Reed, S. L., Schellart, P., Slater, C. T., Sullivan, I., Swinbank, J. D., Taranu, D., Waters, C. Z., and Wood-Vasey, W. M., "An Overview of the LSST Image Processing Pipelines," in [*Astronomical Data Analysis Software and Systems XXVII*], Teuben, P. J., Pound, M. W., Thomas, B. A., and Warner, E. M., eds., *Astronomical Society of the Pacific Conference Series* **523**, 521 (2019).
- [4] Bosch, J., Armstrong, R., Bickerton, S., Furusawa, H., Ikeda, H., Koike, M., Lupton, R., Mineo, S., Price, P., Takata, T., Tanaka, M., Yasuda, N., AlSayyad, Y., Becker, A. C., Coulton, W., Coupon, J., Garmilla, J., Huang, S., Krughoff, K. S., Lang, D., Leauthaud, A., Lim, K.-T., Lust, N. B., MacArthur, L. A., Mandelbaum, R., Miyatake, H., Miyazaki, S., Murata, R., More, S., Okura, Y., Owen, R., Swinbank, J. D., Strauss, M. A., Yamada, Y., and Yamanoi, H., "The Hyper Suprime-Cam software pipeline," *PASJ* **70**, S5 (Jan. 2018).
- [5] Kantor, J. P., "Managing the evolution of the LSST data management system," in [*Advanced Software and Control for Astronomy*], Lewis, H. and Bridger, A., eds., *Proc. SPIE* **6274**, 62740P (June 2006).
- [6] Wang, S.-Y., Huang, P.-J., Hsin-Yo, C., Kimura, M., Wen, C.-Y., Yan, C.-H., Karr, J., et al., "Prime Focus Spectrograph (PFS): the prime focus instrument," in [*Ground-based and Airborne Instrumentation for Astronomy VIII*], *Proc. SPIE* **11447**, 114477V (Dec. 2020).
- [7] Crill, B. P., Werner, M., Akeson, R., Ashby, M., Bleem, L., Bock, J. J., Bryan, S., et al., "SPHEREx: NASA's near-infrared spectrophotometric all-sky survey," in [*Space Telescopes and Instrumentation 2020: Optical, Infrared, and Millimeter Wave*], *Proc. SPIE* **11443**, 114430I (Dec. 2020).
- [8] Myers, J. and Copeland, R., [*Essential SQLAlchemy: Mapping Python to Databases*], O'Reilly Media, Inc. (2015).

- [9] Miyazaki, S., Komiyama, Y., Kawanomoto, S., Doi, Y., Furusawa, H., Hamana, T., Hayashi, Y., et al., “Hyper Suprime-Cam: System design and verification of image quality,” *PASJ* **70**, S1 (Jan. 2018).
- [10] Flaughner, B., Diehl, H. T., Honscheid, K., Abbott, T. M. C., Alvarez, O., Angstadt, R., Annis, J. T., et al., “The Dark Energy Camera,” *AJ* **150**, 150 (Nov. 2015).
- [11] Leach, P., Mealling, M., and Salz, R., “RFC 4122: A Universally Unique Identifier (UUID) URN Namespace.” IETF, <http://www.ietf.org/rfc/rfc4122.txt> (2005).
- [12] Sick, J. and Fausti, A., “LSST Verification Framework API Demonstration,” Vera C. Rubin Observatory SQuaRE Technical Note, SQR-019 <https://sqr-019.lsst.io/> (Apr. 2018).
- [13] Jenness, T., “Tracking Metrics in Butler,” Vera C. Rubin Observatory Data Management Technical Note, DMTN-203 <https://dmtn-203.lsst.io/> (Aug. 2021).
- [14] Allbery, R., “Possible authorization approaches for Butler,” Vera C. Rubin Observatory Data Management Technical Note, DMTN-182 <https://dmtn-182.lsst.io/> (Mar. 2021).
- [15] O’Mullane, W., Economou, F., Huang, F., Speck, D., Chiang, H.-F., Graham, M. L., Allbery, R., Banek, C., Sick, J., Thornton, A. J., Masciarelli, J., Lim, K.-T., Mueller, F., Padolski, S., Jenness, T., Krughoff, K. S., Gower, M., Guy, L. P., and Dubois-Felsmann, G. P., “Rubin Science Platform on Google: the story so far,” in [*Astronomical Data Analysis Software and Systems XXXI*], *ASP Conf. Ser.* **in press**, arXiv:2111.15030 (Nov. 2021).
- [16] Jenness, T., “A client/server Butler,” Vera C. Rubin Observatory Data Management Technical Note, DMTN-176 <https://dmtn-176.lsst.io/> (Mar. 2021).
- [17] Jenness, T. and Economou, F., “ORAC-DR: A generic data reduction pipeline infrastructure,” *Astronomy and Computing* **9**, 40–48 (Mar. 2015).
- [18] Labrie, K., Cardenas, R., Anderson, K., Simpson, C., and Turner, J., “DRAGONS: One Pipeline to Rule them All,” in [*Astronomical Data Analysis Software and Systems XXVII*], Ballester, P., Ibsen, J., Solar, M., and Shortridge, K., eds., *Astronomical Society of the Pacific Conference Series* **522**, 583 (Apr. 2020).
- [19] Thain, D., Tannenbaum, T., and Livny, M., “Distributed computing in practice: the Condor experience,” *Concurrency and computation: practice and experience* **17**(2-4), 323–356 (2005).
- [20] Deelman, E., Vahi, K., Juve, G., Rynge, M., Callaghan, S., Maechling, P. J., Mayani, R., Chen, W., Ferreira da Silva, R., Livny, M., and Wenger, K., “Pegasus, a workflow management system for science automation,” *Future Generation Computer Systems* **46**, 17–35 (2015).
- [21] Babuji, Y., Woodard, A., Li, Z., Katz, D. S., Clifford, B., Kumar, R., Lacinski, L., Chard, R., Wozniak, J. M., Foster, I., Wilde, M., and Chard, K., “Parsl: Pervasive Parallel Programming in Python,” in [*Proceedings of the 28th International Symposium on High-Performance Parallel and Distributed Computing*], *HPDC ’19*, 25–36, Association for Computing Machinery, New York, NY, USA (2019).
- [22] Maeno, T., De, K., Wenaus, T., Nilsson, P., Stewart, G. A., Walker, R., Stradling, A., Caballero, J., Potekhin, M., and and, D. S., “Overview of ATLAS PanDA workload management,” *Journal of Physics: Conference Series* **331**, 072024 (Dec. 2011).
- [23] Bektesevic, D., Chiang, H.-F., Lim, K.-T., Miller, T. L., Thain, G., Jenness, T., Bosch, J., Salmikov, A., and Connolly, A., “A Gateway to Astronomical Image Processing: Vera C. Rubin Observatory LSST Science Pipelines on AWS,” in [*Gateways 2020, October 12-23, 2020*], **Online**, <https://osf.io/2rqfb> (arXiv:2011.06044), Science Gateways Community Institute (Oct. 2020).
- [24] Gower, M. and Lim, K.-T., “Batch Production Services Design,” Vera C. Rubin Observatory Data Management Technical Note, DMTN-123 <https://dmtn-123.lsst.io/> (Aug. 2019).
- [25] Chiang, H.-F., Bektesevic, D., and the AWS-PoC team, “AWS Proof of Concept Project Report,” Vera C. Rubin Observatory Data Management Technical Note, DMTN-137 <https://dmtn-137.lsst.io/> (Jan. 2020).
- [26] Jenness, T., “Modern Python at the Large Synoptic Survey Telescope,” in [*Astronomical Data Analysis Software and Systems XXVII*], Ballester, P., Ibsen, J., Solar, M., and Shortridge, K., eds., *Astronomical Society of the Pacific Conference Series* **522**, 541 (Apr. 2020).
- [27] Jenness, T., Economou, F., Findeisen, K., Hernandez, F., Hoblitt, J., et al., “LSST data management software development practices and tools,” in [*Software and Cyberinfrastructure for Astronomy V*], *Proc. SPIE* **10707**, 1070709 (July 2018).

- [28] Jin, W., Zhong, D., Ding, Z., Fan, M., and Liu, T., “Where to Start: Studying Type Annotation Practices in Python,” in [2021 36th IEEE/ACM International Conference on Automated Software Engineering (ASE)], 529–541 (2021).
- [29] LSST Dark Energy Science Collaboration (LSST DESC), Abolfathi, B., Alonso, D., Armstrong, R., Aubourg, É., Awan, H., Babuji, Y. N., et al., “The LSST DESC DC2 Simulated Sky Survey,” *ApJS* **253**, 31 (Mar. 2021).
- [30] O’Mullane, W., “Data Preview 0: Definition and planning.,” Vera C. Rubin Observatory Technical Note, RTN-001 <https://rtn-001.lsst.io/> (Sept. 2021).