Large Synoptic Survey Telescope
Data Management Applications Design

Mario Jurić[*], R.H. Lupton, T. Axelrod, J.F. Bosch, P.A. Price,
G.P. Dubois-Felsmann, Ž. Ivezić, A.C. Becker, J. Becla,
A.J. Connolly, J. Kantor, K-T Lim, D. Shaw,
for the LSST Data Management

Monday 20th June, 2016, 18:50hrs

---

[*] Please direct comments to <mjuric@lsst.org>.

Abstract

The LSST Science Requirements Document (the LSST SRD) specifies a set of data product guidelines, designed to support science goals envisioned to be enabled by the LSST observing program. Following these guidlines, the details of these data products have been described in the LSST Data Products Definition Document (DPDD), and captured in a formal flow-down from the SRD via the LSST System Requirements (LSR), Observatory System Specifications (OSS), to the Data Management System Requirements (DMSR). The LSST Data Management subsystem's responsibilities include the design, implementation, deployment and execution of software pipelines necessary to generate these data products. This document, in conjunction with the UML Use Case model (LDM-134), describes the design of the scientific aspects of those pipelines.

Contents

1 Preface

The purpose of this document is to describe the design of pipelines belonging to the Applications Layer of the Large Synoptic Survey Telescope (LSST) Data Management system. These include most of the core astronomical data processing software that LSST employs.

The intended audience of this document are LSST software architects and developers. It presents the baseline architecture and algorithmic selections for core DM pipelines. The document assumes the reader/developer has the required knowledge of astronomical image processing algorithms and solid understanding of the state of the art of the field, understanding of the LSST Project goals and concepts, and has read the LSST Science Requirements (SRD) as well as the LSST Data Products Definition Document (DPDD).

This document should be read in conjunction with the LSST DM Applications Use Case Model (LDM-134). They are intended to be complementary, with the Use Case model capturing the detailed (inter)connections between individual pipeline components, and this document capturing the overall goals, pipeline architecture, and algorithmic choices.

Though under strict change control[1], this is a ***living document***. Firstly, as a consequence of the "rolling wave" LSST software development model, the designs presented in this document will be refined and made more detailed as particular pipeline functionality is about to be implemented. Secondly, the LSST will undergo a period of construction and commissioning lasting no less than seven years, followed by a decade of survey operations. To ensure their continued scientific adequacy, the overall designs and plans for LSST data processing pipelines will be periodically reviewed and updated.

---

[1] LSST Docushare handle for this document is `LDM-151`.

2 Introduction

2.1 LSST Data Management System

To carry out this mission the Data Management System (DMS) performs the following major functions:

- Processes the incoming stream of images generated by the camera system during observing to produce transient alerts and to archive the raw images.
- Roughly once per year, creates and archives a Data Release ("DR"), which is a static self-consistent collection of data products generated from all survey data taken from the date of survey initiation to the cutoff date for the Data Release. The data products (described in detail in the DPDD), include measurements of the properties (shapes, positions, fluxes, motions, etc.) of all detected objects, including those below the single visit sensitivity limit, astrometric and photometric calibration of the full survey object catalog, and limited classification of objects based on both their static properties and time-dependent behavior. Deep coadded images of the full survey area are produced as well.
- Periodically creates new calibration data products, such as bias frames and flat fields, that will be used by the other processing functions, as necessary to enable the creation of the data products above.
- Makes all LSST data available through interfaces that utilize, to the maximum possible extent, community-based standards such as those being developed by the Virtual Observatory ("VO"), and facilitates user data analysis and the production of user-defined data products at Data Access Centers ("DAC") and at external sites.

The overall architecture of the DMS is discussed in more detail in the Data Management System Design (DMSD) document. The overall architecture of the DMS is shown in Figure 1.

This document discusses the role of the Applications layer in the first three functions listed above (the functions involving *science pipelines*). The fourth is discussed separately in the SUI Conceptual Design Document (SUID).

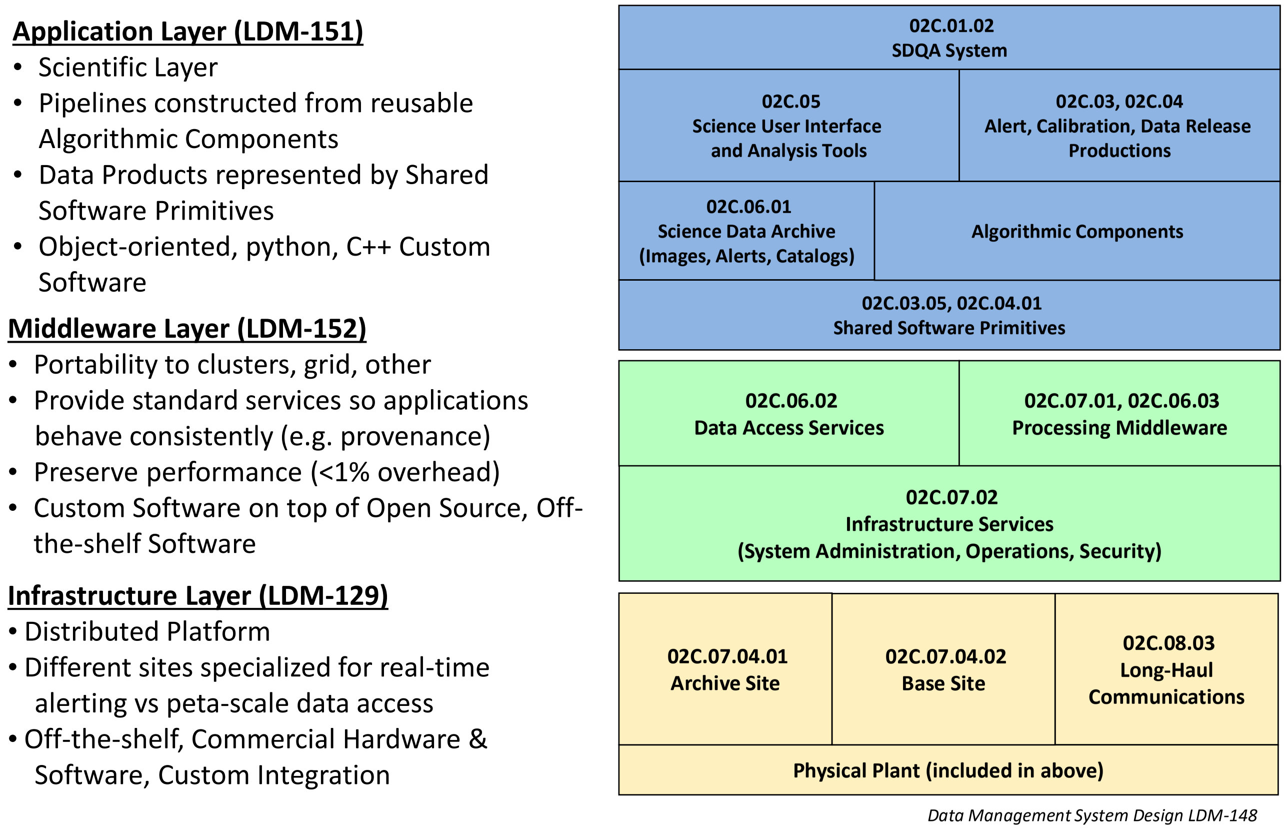

Application Layer (LDM-151)

- Scientific Layer
- Pipelines constructed from reusable Algorithmic Components
- Data Products represented by Shared Software Primitives
- Object-oriented, python, C++ Custom Software

Middleware Layer (LDM-152)

- Portability to clusters, grid, other
- Provide standard services so applications behave consistently (e.g. provenance)
- Preserve performance (<1% overhead)
- Custom Software on top of Open Source, Off-the-shelf Software

Infrastructure Layer (LDM-129)

- Distributed Platform
- Different sites specialized for real-time alerting vs peta-scale data access
- Off-the-shelf, Commercial Hardware & Software, Custom Integration

Figure 1: Architecture of the Data Management System

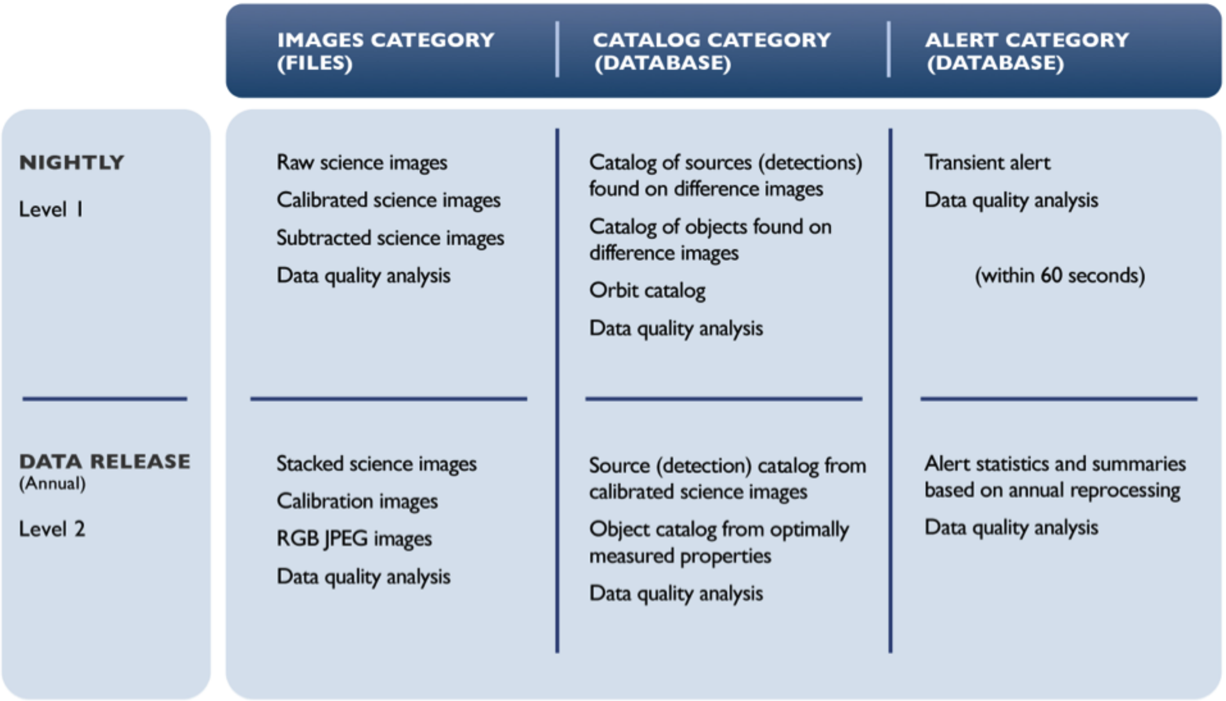

	IMAGES CATEGORY (FILES)	**CATALOG CATEGORY (DATABASE)**	**ALERT CATEGORY (DATABASE)**
NIGHTLY Level 1	Raw science images Calibrated science images Subtracted science images Data quality analysis	Catalog of sources (detections) found on difference images Catalog of objects found on difference images Orbit catalog Data quality analysis	Transient alert Data quality analysis (within 60 seconds)
DATA RELEASE (Annual) Level 2	Stacked science images Calibration images RGB JPEG images Data quality analysis	Source (detection) catalog from calibrated science images Object catalog from optimally measured properties Data quality analysis	Alert statistics and summaries based on annual reprocessing Data quality analysis

Figure 2: Organization of LSST Data Products

2.2 Data Products

The LSST data products are organized into three groups, based on their intended use and/or origin. The full description is provided in the Data Products Definition Document (DPDD); we summarize the key properties here to provide the necessary context for the discussion to follow.

- **Level 1** products are intended to support timely detection and follow-up of time-domain events (variable and transient sources). They are generated by near-real-time processing the stream of data from the camera system during normal observing. Level 1 products are therefore continuously generated and / or updated every observing night. This process is of necessity highly automated, and must proceed with absolutely minimal human interaction. In addition to science data products, a number of related Level 1 "SDQA"[2] data products are generated to assess quality and to provide feedback to the Observatory Control System (OCS).

- **Level 2** products are generated as part of a Data Release, generally performed yearly, with an additional data release for the first 6 months of survey data. Level 2 includes data products for which extensive computation is required, often because they combine information from many exposures. Although the steps that generate Level 2 products will be automated, significant human interaction may be required at key points to ensure the quality of the data.

- **Level 3** products are generated on any computing resources anywhere and then stored in an LSST Data Access Center. Often, but not necessarily, they will be generated by users of LSST using LSST software and/or hardware. LSST DM is required to facilitate the creation of Level 3 data products by providing suitable APIs, software components, and computing infrastructure, but will not by itself create any Level 3 data products. Once created, Level 3 data products may be associated with Level 1 and Level 2 data products through database federation. Where appropriate, the LSST Project, with the agreement of the Level 3 creators, may incorporate user-contributed Level 3 data product pipelines into the DMS production flow, thereby promoting them to Level 1 or 2.

---

[2] Science Data Quality Analysis

The organization of LSST Data Products is shown in Figure 2.

Level 1 and Level 2 data products that have passed quality control tests will be accessible to the public without restriction. Additionally, the source code used to generate them will be made available, and LSST will provide support for builds on selected platforms.

The pipelines used to produce these public data products will also produce many intermediate data products that may not be made publically available (generally because they are fully superseded in quality by a public data product). Intermediate products may be important for QA, however, and their specification is an important part of describing the pipelines themselves.

2.3 Data Units

In order to describe the components of our processing pipelines, we first need standard nomenclature for the units of data the pipeline will process.

The smallest data units are those corresponding to individual astrophysical entities. In keeping with LSST conventions, we use “object” to refer to the astrophysical entity itself (which typically implies aggregation of some sort over all exposures), and “source” to refer to the realization of an object on a particular exposure. In the case of blending, of course, these are just our best attempts to define distinct astrophysical objects, and hence it is also useful to define terms that represent this process. We use “family” to refer to group of blended objects (or, more rarely, sources), and “child” to refer to a particular deblended object within a family. A “parent” is also created for each family, representing the alternate hypothesis that the blend is actually a single object. Blends may be hierarchical; a child at one level may be a parent at the level below.

LSST observations are taken as a pair of 15-second “snaps”; together these constitute a “visit”. Because snaps are typically combined early in the processing (and some special programs and survey modes may take only a single snap), visit is much more frequently used as a unit for processing and data products. The image data for to a visit is a set of 189 “CCD” or “sensor” images. CCD-level data from the camera is further data divided across the 16 amplifiers within a CCD, but these are also combined at an early stage, and the 3×3 CCD “rafts” that play an important role in the hardware design are relativley unimportant for the pipeline. This leaves visit and CCD the main identifiers of most exposure-level data products and pipelines.

Our convention for defining regions on the sky is deliberately vague; we hope to build a codebase capable of working with virtually any pixelization or projection scheme (though different schemes may have different performance or storage implications). Our approach involves two region concepts: "tracts" and "patches". A tract is a large region with a single Cartesian coordinate system; we assume it is larger than the LSST field of view, but its maximum size is essentially set by the point at which distortion in the projection becomes significant enough to affect the processing (by e.g. breaking the assumption that the PSF is well-sampled on the pixel grid). Tracts are divided into patches, all of which share the tract coordinate system. Most image processing is perfomed at the patch level, and hence patch sizes are chosen largely to ensure that patch-level data products and processing fit in memory. Both tracts and patches are defined such that each region overlaps with its neighbors, and these overlap regions must be large enough that any individual astronomical object is wholly contained in at least one tract and patch. In a patch overlap region, we expect pixel values to be numerically equivalent (i.e. equal up to floating point round-off errors) on both sides; in tract overlaps, this is impossible, but we expect the results to be scientifically consistent. Selecting larger tracts and patches thus reduces the overall fraction of the area that falls in overlap regions and must be processed multiple times, while increasing the computational load for processing individual tracts and patches.

2.4 Science Pipelines Organization

As shown in Figure 1, the Applications Layer is itself split into three levels. In sections 3, 4, and 5, we describe the Alert Production, Calibration Products Production, and Da ta Release Production (respectively), breaking them down into *pipelines*. In this document, a pipeline is a high-level combination of algorithms that is intrinsically tied to its role in the production in which it is run. For instance, while both Alert Production and Data Release Production will include a pipeline for single-visit processing, these two pipelines are *distinct*, because the details of their design depend very much on the context in which they are run. Section 6 describes the Science Data Quality Analysis System, a collection of pipelines and mini-productions designed to assess and continuously validate the quality of both the data and the processing system. The SDQA System is not a single production; its components are either directly integrated into other productions or part of a set of multiple mini-productions run on different cadences.

Pipelines are largely composed of Algorithmic Components: mid-level algorithmic code that we expect to reuse (possibly with different configuration) across different productions. These components constitute the bulk of the new code and algorithms to be developed for Alert Production and Data Release Production, and are discussed in section 8. Most algorithmic components are applicable to any sort of astronomical imaging data, but some will be customized for LSST.

The lowest level in the Applications Layer is made up of our shared software primitives: libraries that provide important data structures and low-level algorithms, such as images, tables, coordinate transformations, and nonlinear optimizers. Much (but not all) of this content is astronomy-related, but essentially none of it is specific to LSST, and hence we can and will make use of third-party libraries whenever possible. These primitives also play an important role in connecting the Science User Interface Toolkit and Level 3 processing environment with Level 1 and Level 2 data products, as they constitute the programmatic representation of those data products. Shared software primitives are discussed in section 9.

3 Level 1 Pipelines

3.1 Single Frame Processing Pipeline (WBS 02C.03.01)

3.1.1 Key Requirements

Single Frame Processing (SFM) Pipeline is responsible for reducing raw image data to *calibrated exposures*, and detection and measurement of `Sources` (using the components functionally a part of the Object Characterization Pipeline).

SFM pipeline functions include:

- Assembly of per-amplifier images to an image of the entire CCD;
- Instrumental Signature Removal;
- Cosmic ray rejection and snap combining;
- Per-CCD determination of zeropoint and aperture corrections;
- Per-CCD PSF determination;
- Per-CCD WCS determination and astrometric registration of images;
- Per-CCD sky background determination;
- Source detection and measurement

Calibrated exposure produced by the SFM pipeline must possess all information necessary for measurement of source properties by single-epoch Object Characterization algorithms.

It shall be possible to run this pipeline in two modes: a "fast" mode needed in nightly operations for Level 1 data reductions where no source characterization is done beyond what's required for zero-point, PSF, sky, and WCS determination (image reduction); and a "full" mode that will be run for Level 2 data reductions.

3.1.2 Baseline Design

Single Frame Processing pipeline will be implemented as a flexible framework where different data can be easily treated differently, and new processing steps can be added without modifying the stack code.

It will consist of three primary components:

- A library of useful methods that wrap a small number of atomic operations (e.g., `interpolateFromMask`, `overscanCorrection`, `biasCorrection`, etc.)
- A set of classes (`Tasks`) that perform higher level jobs (e.g., `AssembleCcdTask`, or `FringeTask`), and a top level class to apply corrections to the input data in the proper order. This top level class can be overridden in the instrument specific `obs_*` packages, making the core SFM pipeline camera agnostic.
- A top-level Task to run the SFM pipeline.

In the paragraphs to follow, we describe the adopted baseline for key SFM algorithms. If not discussed explicitly, the algorithmic baseline for all other functionallity is assumed to be the same as that used by SDSS *Photo* pipeline [18].

Output information for OCS telemetry: ACTION clarify OCS interactions

3.1.2.1 Instrumental Signature Removal: Clarify interaction with butler

Input Data

- Camera corrected (crosstalk, overscan, linearity) images
- Sensor defect lists
- Metadata including electronic parameters (saturation limits, readnoise, electronic footprint)

Output Data

- Calexp images

Ancillary Products?

- Source detection and measurements
- ICExp background subtracted images
- Post ISR exposure

Actions in case of failure?
Actions in case camera data are not available due to network outage longer than buffer of data at summit **Alternative procedures?**

Subtasks:

- Mask defects and saturation
- Assembly
- Full frame corrections: Dark, Flats (includes fringing)
- Pixel level corrections: Brighter fatter, static pixel size effects
- **QUESTION is this run prior to pixel level corrections** Interpolation of defects and saturation
- CR rejection
- Generate snap difference
- Snap combination

3.1.2.2 PSF determination and background determination:

Input Data?
Output Data?
Anscillary Products?
Actions in case of failure?
Alternative procedures?

Subtasks:
Iterate till convergence (convergence criteria TBD)

- Background estimation
- Source detection
- Selection of PSF candidate stars
- PSF determination

3.1.2.3 Source measurement:

Input Data?
Output Data?
Anscillary Products?
Actions in case of failure?
Alternative procedures?

Subtasks:

- Source measurement - Single Visit Measurement
- Aperture correction

3.1.2.4 Photometric and Astrometric calibration:

Input Data?
DRP's internal reference catalog **Output Data?**
OCS PSF, WCS, metadata (TBD) **Anscillary Products?**
Actions in case of failure?
Alternative procedures?

Subtasks:

- Source association
- CCD level photometric solution
- Visit level photometric solution
- Remove known astrometric distortions
- Fit remaining residual
- Single visit composed astrometric solution

- Output information for OCS telemetry: WCS ACTION clarify OCS interactions

OUTPUT: Calibrated Exposure and Calibrated Catalog

3.1.3 Prototype Implementation

The prototype codes are available in the following repositories: `https://github.com/lsst/ip_isr`, `https://github.com/lsst/meas_algorithms`, `https://github.com/lsst/meas_astrom`, `https://github.com/lsst-dm/legacy-meas_mosaic`, `https://github.com/lsst/pipe_tasks`.

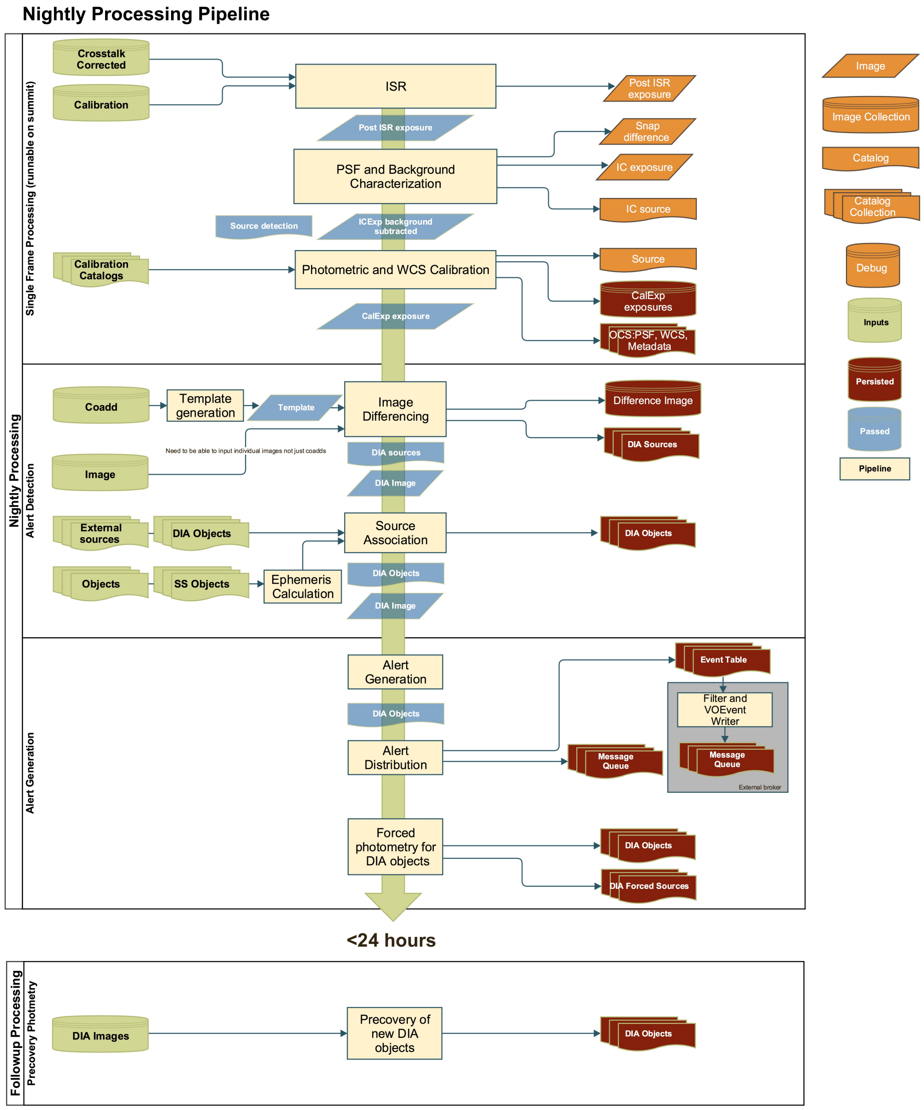

Figure 3: The nightly processing flowchart describing the flow of images and data through single frame processing, image differencing, alert generation and production

3.2 Alert Detection (WBS 02C.03.04)

3.2.1 Key Requirements

The alert detection pipeline shall difference a visit image against a deeper template, and detect and characterize sources in the difference image in the time required to achieve the 60 second design goal for Level 1 alert processing (current timing allocation: 24 seconds). The algorithms employed by the pipeline shall result in purity and completeness of the sample as required by the DMSR . Image differencing shall perform as well in crowded as in uncrowded fields.

3.2.2 Baseline Design

3.2.2.1 Template Generation

Input Data?
Coadded CalExps or Series of CalExp from which to interpolate a template.
Output Data?
Anscillary Products?
Actions in case of failure?
Alternative procedures?

Subtasks:

- Determine appropriate template to use
- Generate template for observation

3.2.2.2 Image differencing

Input Data?
Internal reference catalog for CalExp from DRP PSF for science image **Output Data?**
Anscillary Products?
Actions in case of failure?
Alternative procedures?

Subtasks:

- match DRP sources and sources from SFP

- Determine relative astrometric solution
- Warp template and measurements to science image frame
- Correlate science image with science PSF (pre-convolution)
- Determine appropriate PSF matching sources
- Compute PSF matching kernel and spatial model using ZOGY approach
- Difference science and template images
- Apply correction for correlated noise
- Difference image source detection
- Difference image source measurement: dipole fit, trailed source measurement
- Measure flux on snap difference for all DIASources

3.2.2.3 Real-Bogus classification

Input Data?
Internal reference catalog for CalExp from DRP PSF for science image **Output Data?**
Anscillary Products?
Actions in case of failure?
Alternative procedures?

Subtasks:

- Application of random forest or other classification algorithm
- Update DIASources with probabilistic classification
- Filter DIASource list based on classifier

3.2.2.4 Ephemeris Calculation

Input Data?
Output Data?
Anscillary Products?
Actions in case of failure?
Alternative procedures?

Subtasks:

- Calculate positions for all solar system objects that may overlap the current exposure.

3.2.2.5 Source Association

Input Data?
Output Data?
Anscillary Products?
Actions in case of failure?
Alternative procedures?

Subtasks:

- Match all DIASources to predicted Solar System object positions and DIAObject catalog positions
- Perform forced photometry of un-associated DIAObjects.(Maybe not if we force photometer all DIAObjects?). SSObjects will not be force photometered because the precision of the prediction will not be good enough. Force photometry for external DIAObjects?
- Update associated DIAObjects with aggregate quantities: e.g. parallax, proper motion, and variability metrics
- New spuriousness calculation?

3.2.3 Prototype Implementation

The prototype code is available at `https://github.com/lsst/ip_diffim`. The current prototype, while functional, will require a partial redesign to be transfered to construction to address performance and extensibility concerns.

3.3 Alert Generation Pipeline (WBS 02C.03.03)

3.3.1 Key Requirements

Alert Generation Pipeline shall take the newly discovered `DIASources` and all associated metadata as described in the DPDD, and generate alert packets in `VOEvent` format. It will transmit these packets to VO Event Brokers, using standard IVOA protocols (eg., VOEvent Transport Protocol; VTP). End-users will primarily use these brokers to classify and filter events for subsets fitting their science goals.

To directly serve the end-users, the Alert Generation Pipeline shall provide a basic, limited capacity, alert filtering service. This service will run at the LSST U.S. Archive Center (at NCSA). It will let astronomers create simple filters that limit what alerts are ultimately forwarded to them. These *user defined filters* will be possible to specify using an SQL-like declarative language, or short snippets of (likely Python) code.

3.3.2 Baseline Design

3.3.2.1 Alert generation

Input Data?
Output Data?
Anscillary Products?
Actions in case of failure?
Alternative procedures?

Subtasks:

- Generate postage stamps for all DIASources: direct image and difference image
- Push alert records to alert database

3.3.2.2 Alert Distribution

Input Data?
Output Data?
Anscillary Products?
Actions in case of failure?
Alternative procedures?

Subtasks:

- Filter event records (for content as well as for events)
- Author VOEvent
- Push to messaging queue

3.3.2.3 Forced Photometry on all DIAObjects

Input Data?
Output Data?
Anscillary Products?
Actions in case of failure?
Alternative procedures?

Subtasks:

- Compute forced photometry on all DIAObjects in the field. This does not end up in the alerts.

3.3.3 Prototype Implementation

3.4 Precovery Photometry Pipeline

3.4.1 Key Requirements

Within 24 hrs.

3.4.1.1 Precovery of new DIAObjects

Input Data?
Output Data?
Anscillary Products?
Actions in case of failure?
Alternative procedures?

Subtasks:

- Force photometer in difference images for all new DIAObjects for the past 30 days.

3.5 Moving Object Pipeline (WBS 02C.03.06)

3.5.1 Key Requirements

The Moving Object Pipeline System (MOPS) has two responsibilities within LSST Data Management:

- First, it is responsible for generating and managing the Solar System[3] data products. These are Solar System objects with associated Keplerian orbits, errors, and detected `DIASources`. Quantitatively, it shall be capable of detecting 95% of all Solar System objects that meet the findability criteria as defined in the OSS . The software components implementing this function are known as ***DayMOPS***.
- The second responsibility of the MOPS is to predict future locations of moving objects in incoming images so that their sources may be associated with known objects; this will reduce the number of spurious transient detections and appropriately flag alerts to detections of known Solar System objects. The software components implementing this function are known as ***NightMOPS***.

3.5.2 Baseline Design

3.5.2.1 Generate Tracklets

Output Data?
Anscillary Products?
Actions in case of failure?
Alternative procedures?

Subtasks:

- Make all tracklet pairs
- Merge multiple chained observation into single longer tracklets
- Purge any tracklets inconsistent with the merged tracklets

---

[3] Also sometimes referred to as 'Moving Object'

3.5.2.2 Attribution and precovery

Output Data?
Anscillary Products?
Actions in case of failure?
Alternative procedures?

Subtasks:

- Predict locations of known Solar System objects
- Match tracklet observation to predicted ephimerides taking into account velocity
- Update SSObjects
- Possibly iterate

3.5.2.3 Fit Orbits

Output Data?
Anscillary Products?
Actions in case of failure?
Alternative procedures?

Subtasks:

- Merge unassociated tracklets into tracks.
- Fit orbits to all tracks.
- Purge unphysical tracks.
- Update SSObjects
- Possibly iterate

3.5.2.4 Association and Precovery: New SSObjects

Output Data?
Anscillary Products?
Actions in case of failure?
Alternative procedures?

Subtasks:

- Do association and precovery just for SSObjects just found
- Update SSObjects

3.5.2.5 Merge Orbits

Output Data?
Anscillary Products?
Actions in case of failure?
Alternative procedures?

Subtasks:

- Merge orbits with high probability of being the same orbit into a single SSObject

3.5.3 Prototype Implementation

Prototype MOPS codes are available at `https://github.com/lsst/mops_daymops` and `https://github.com/lsst/mops_nightmops`. We expect it will be possible to transfer a significant fraction of the existing code into Construction. Current DayMOPS prototype already performs within the computational envelope envisioned for LSST Operations, though it does not yet reach the required completeness requirement.

4 Calibration Products Production

4.1 Calibration Products Pipeline (WBS 02C.04.02)

4.1.1 Key Requirements

The work performed in this WBS serves two complementary roles:

- It will enable the production of calibration data products as required by the Level 2 Photometric Calibration Plan (LSE-180) and other planning documents [20][4]. This includes both characterization of the sensitivity of the LSST system (optics, filters and detector) and the transmissivity of the atmosphere.
- It will characterize of detector anomalies in such a way that they can be corrected either by the instrument signature removal routines in the Single Frame Processing Pipeline (WBS 02C.03.01) or, if appropriate, elsewhere in the system;
- It will manage and provide a catalog of optical ghosts and glints to other parts of the system upon demand.

4.1.2 Baseline Design

4.1.2.1 Instrumental sensitivity We expect laboratory measurements of the filter profiles. We further baseline the development of a procedure for measuring the filter response at 1 nm resolution using the approach described in [20].

We baseline the following procedure for creating flat fields:

1. Record bias/dark frames;
2. Use “monochromatic” (1 nm) flat field screen flats with no filter in the beam to measure the per-pixel sensitivity;
3. Use a collimated beam projector (CBP) to measure the quantum efficiency (QE) at a set of points in the focal plane, dithering those points to tie them together;

---

[4] Resolving contradictions between these documents is out of scope here.

4. Combine the screen and CBP data to determine the broad band (10–100 nm) QE of all pixels;
5. Fold in the filter response to determine the 1 nm resolution effective QE of all pixels.

This WBS is responsible for the development of the data analysis algorithms and software required and the ultimate delivery of the flat fields. Development and commissioning of the CBP itself, together with any other infrastructure required to perform the above procedure, lies outwith Data Management (see 04C.08 *Calibration System*).

4.1.2.2 Atmospheric transmissivity Measurements from the auxiliary instrumentation—to include the 1.2 m "Calypso" telescope, a bore-sight mounted radiometer and satellite-based measurement of atmospheric parameters such as pressure and ozone—will be used to determine the atmospheric absorption along the line of sight to standard stars. The atmospheric transmission will be decomposed into a set of basis functions and interpolated in space in time to any position in the LSST focal plane.

This WBS will develop a pipeline for accurate spectrophotometric measurement of stars with the auxiliary telescope. We expect to repurpose and build upon publicly available code e.g. from the PFS[5] project for this purpose.

This WBS will construct the atmospheric model, which may be based either on MODTRAN (as per LSE-180) or a PCA-like decomposition of the data (suggested by [20]).

This WBS will define and develop the routine for fitting the atmospheric model to each exposure from the calibration telescope and providing estimates of the atmospheric transmission at any point in the focal plane upon request.

4.1.2.3 Detector effects An initial cross-talk correction matrix will be determined by laboratory measurements on the Camera Calibration Optical Bench (CCOB). However, to account for possibile instabilities, this WBS will develop an on-telescope method. We baseline this as being based on measurement with the CBP, but we note the alternative approach based on cosmic rays adopted by HSC [13].

Multiple reflections between the layers of the CCD give rise to spatial variability with fine scale structure in images which may vary with time [20,

---

[5] Subaru's Prime Focus Spectrograph; `http://sumire.ipmu.jp/pfs/`.

§2.5.1]. These can be characterized by white light flat-fields. Preliminary analysis indicates that these effects may be insignificant in LSST [23]; however, the baseline calls for a a routine developed in this WBS to analyse the flat field data and generate fringe frames on demand. This requirement may be relaxed if further analysis (outside the scope of thie WBS) demonstrates it to be unnecessary.

This WBS will develop algorithms to characterize and mitigate anomalies due to the nature of the camera's CCDs.

> **Note:**
> There's a complex inter-WBS situation here: the actual mitigation of CCD anomalies will generally be performed in SFM (WBS 02C.03.01), based on products provided by this WBS which, in turn, may rely on laboratory based research which is broadly outside the scope of DM. We baseline the work required to develop the corrective algorithms here. We consider moving it to WBS 02C.03.01 in future.

The effects we anticipate include:

- QE variation between pixels;
- Static non-uniform pixel sizes (e.g. "tree rings" [27]);
- Dynamic electric fields (e.g. "brighter-fatter" [2]);
- Time dependent effects in the camera (e.g. hot pixels, changing cross-talk coefficients);
- Charge transfer (in)efficiency (CTE).

Laboratory work required to understand these effects is outwith the scope of this WBS. In some cases, this work may establish that the impact of the effect may be neglected in LSST. The baseline plan addresses these issues through the following steps:

- Separate QE from pixel size variations[6] and model both as a function of position (and possibly time);

---

[6] Refer to work by Rudman.

- Learn how to account for pixel size variation over the scale of objects (e.g. by redistributing charge);
- Develop a correction for the brighter-fatter effect and develop models for any features which cannot be removed;
- Handle edge/bloom using masking or charge redistribution;
- Track defects (hot pixels);
- Handle CTE, including when interpolating over bleed trails.

4.1.2.4 Ghost catalog The Calibration Products Pipeline must provide a catalog of optical ghosts and glints which is available for use in other parts of the system. Detailed characterization of ghosts in the LSST system will only be possible when the system is operational. Our baseline design therefore calls for this system to be prototyped using data from precursor instrumentation; we note that ghosts in e.g. HSC are well known and more significant than are expected in LSST.

Note:
It is not currently clear where the responsibility for characterizing ghosts and glints in the system lies. We assume it is outwith this WBS.

4.1.3 Constituent Use Cases and Diagrams

Produce Master Fringe Exposures; Produce Master Bias Exposure; Produce Master Dark Exposure; Calculate System Bandpasses; Calculate Telescope Bandpasses; Construct Defect Map; Produce Crosstalk Correction Matrix; Produce Optical Ghost Catalog; Produce Master Pupil Ghost Exposure; Determine CCOB-derived Illumination Correction; Determine Optical Model-derived Illumination Correction; Create Master Flat-Spectrum Flat; Determine Star Raster Photometry-derived Illumination Correction; Create Master Illumination Correction; Determine Self-calibration Correction-Derived Illumination Correction; Correct Monochromatic Flats; Reduce Spectrum Exposure; Prepare Nightly Flat Exposures;

4.1.4 Prototype Implementation

While parts of the Calibration Products Pipeline have been prototyped by the LSST Calibration Group (see the LSE-180 for discussion), these have not been written using LSST Data Management software framework or coding standards. We therefore expect to transfer the know-how, and rewrite the implementation.

4.2 Photometric Calibration Pipeline (WBS 02C.03.07)

4.2.1 Key Requirements

The Photometric Calibration Pipeline is required to internally calibrate the relative photometric zero-points of every observation, enabling the Level 2 catalogs to reach the required SRD precision.

4.2.2 Baseline Design

The adopted baseline algorithm is a variant of “ubercal” [22, 25]. This baseline is described in detail in the Photometric Self Calibration Design and Prototype Document (UCAL).

4.2.3 Constituent Use Cases and Diagrams

Perform Global Photometric Calibration;

4.2.4 Prototype Implementation

Photometric Calibration Pipeline has been fully prototyped by the LSST Calibration Group to the required level of accuracy and performance (see the UCAL document for discussion).

As the prototype has not been written using LSST Data Management software framework or coding standards, we assume a non-negligible refactoring and coding effort will be needed to convert it to production code in LSST Construction.

4.3 Astrometric Calibration Pipeline (WBS 02C.03.08)

4.3.1 Key Requirements

The Astrometric Calibration Pipeline is required to calibrate the relative and absolute astrometry of the LSST survey, enabling the Level 2 catalogs to reach the required SRD precision.

4.3.2 Baseline Design

Algorithms developed for the Photometric Calibration Pipeline (WBS 02C.03.07) will be repurposed for astrometric calibration by changing the relevant functions to minimize. This pipeline will further be aided by WCS and local astrometric registration modules developed as a component of the Single Frame Processing pipeline (WBS 02C.03.01).

Gaia standard stars will be used to fix the global astrometric system. It is likely that the existence of Gaia catalogs may make a separate Astrometric Calibration Pipeline unnecessary.

4.3.3 Constituent Use Cases and Diagrams

Perform Global Astrometric Calibration;

4.3.4 Prototype Implementation

The Astrometric Calibration Pipeline has been partially prototyped by the LSST Calibration Group, but outside of LSST Data Management software framework. We expect to transfer the know-how, and rewrite the implementation.

5 Data Release Production

TODO:
Update figure to reflect merging BootstrapJointCal and RefineJointCal into just StandardJointCal (flow is the same, but these pipelines do exactly the same thing, so there's no need for different names).

A Data Release Production is run every year (twice in the first year of operations) to produce a set of catalog and image data products derived from all observations from the beginning of the survey to the point the production began. This includes running a variant of the difference image analysis run in Alert Production, in addition to direct analysis of individual exposures and coadded images. The data products produced by a Data Release Production are summarized in table 1.

From a conceptual standpoint, data release production can be split into five groups of pipelines, executed in approximately the following order:

1. We characterize and calibrate each exposure, estimating point-spread functions, background models, and astrometric and photometric calibration solutions. This iterates between processing individual exposures independently and jointly fitting catalogs derived from multiple overlapping exposures. These steps are described more fully in section 5.1.

2. We alternately combine images and subtract them, using differences to find artifacts and time-variable sources while building coadds that produce a deeper view of the static sky. Coaddition and difference imaging is described in section 5.2.

3. We detect and deblend on coadds, while associating these detection with detections from difference imaging to define objects. We then merge catalogs in the overlap regions between patches and tracts to produce a single contiguous catalog over the full sky. This is described in section 5.3.

4. We measure objects on coadds and visit-level direct and difference images in object characterization, as described section 5.4.

5. After all image processing is complete, we run additional catalog-only pipelines to fill in additional object properties. Unlike previous stages, this postprocessing is not localized on the sky, as it may use statistics computed from the full data release to improve our characterization of individual objects. Postprocessing pipelines are described in section 5.5.

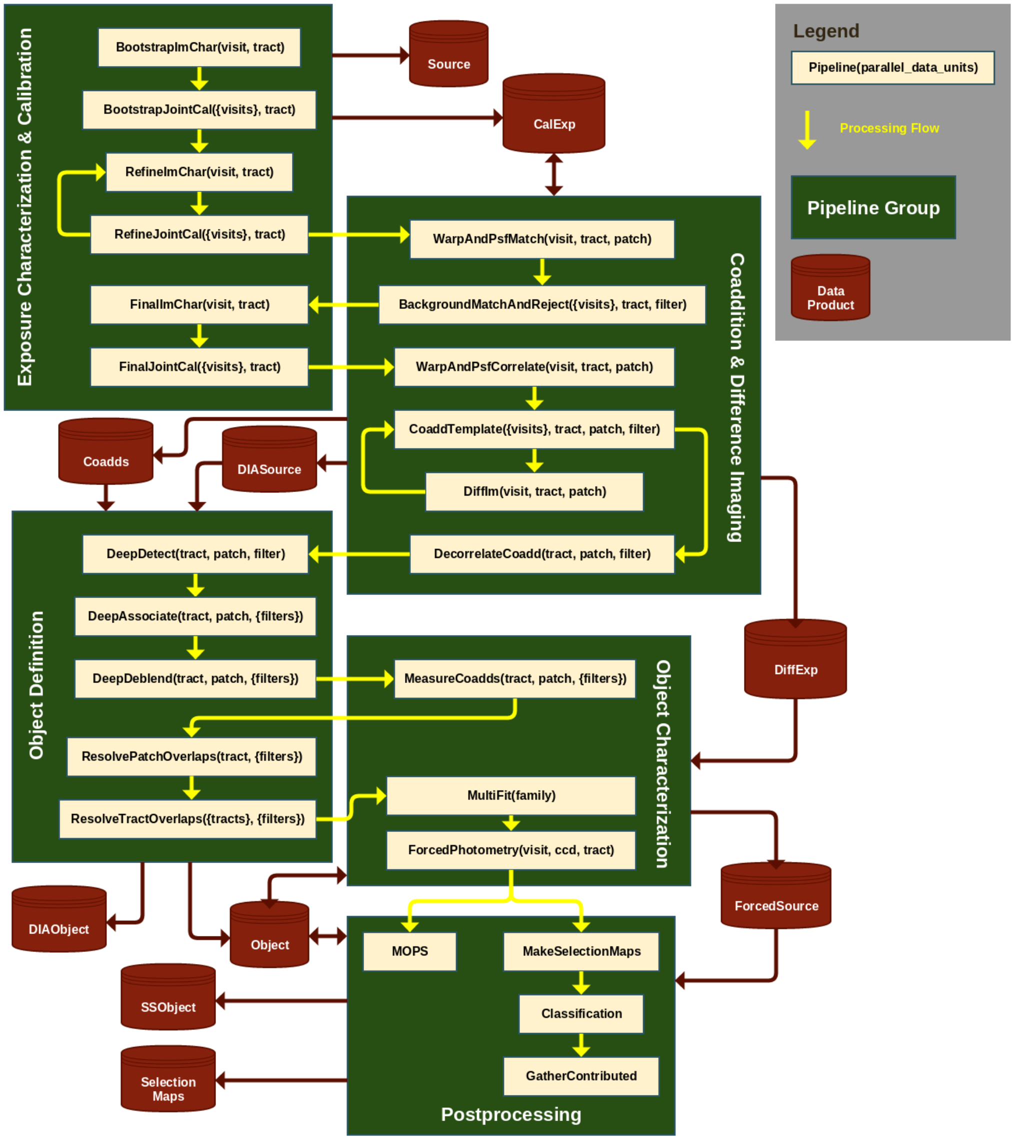

Figure 4: Summary of the Data Release Production processing flow. Processing is split into multiple pipelines, which are conceptually organized into the groups discussed in sections 5.1-5.5.

Name	**Availability**	**Description**
Source	Stored	Measurements from direct analysis of individual exposures.
DIASource	Stored	Measurements from difference imagine analysis of individual exposures.
Object	Stored	Measurements for a single astrophysical object, derived from all available information, including coadd measurements, simultaneous multi-epoch fitting, and forced photometry. Does not include solar system objects.
DIAObject	Stored	Aggregate quantities computing by associating spatially colocated DIASources.
ForcedSource	Stored	Flux measurements on each direct and difference image at the position of every Object.
SSObject	Stored	Solar system objects derived by associating DIASources and inferring their orbits.
CalExp	Regenerated	Calibrated exposure images for each CCD/visit (sum of two snaps).
DiffExp	Regenerated	Difference between CalExp and PSF-matched template coadd.
DeepCoadd	Stored	Coadd image with a reasonable combination of depth and resolution.
EpochRangeCoadd	Renegerated	Coadd image that cover only a limited range of epochs.
BestSeeingCoadd	Regenerated	Coadd image built from only the best-seeing images.
PSFMatchedCoadd	Regenerated	Coadd image with a constant, predetermined PSF.

Table 1: Table of public data products produced during a Data Release Production. A full description of these data products can be found in the Data Products Definition Document (LSE-163).

This conceptual ordering is an oversimplification of the actual processing flow, however; as shown in Figure 4, pipeline groups are actually interleaved.

Each pipeline in this the diagram represents a particular piece of code excuted in parallel on a specific unit of data, but pipelines may contain additional (and more complex) parallelization to further subdivide that data unit. The processing flow also includes the possibility of iteration between pipelines, indicated by cycles in the diagram. The number of iterations in each cycle will be determined (via tests on smaller productions) before the start of the production, allowing us to remove these cycles simply by duplicating some pipelines a fixed number of times. The final data release production processing can thus be described as a directed acyclic graph (DAG) to be executed by the orchestration middleware, with pipelines as edges and (intermediate) data products as vertices. Most of the graph will be generated by applications code before the production begins, using a format and/or API defined by the orchestration middleware. Howver, some parts of the graph must be generated on-the-fly; this will be discussed further in section 5.4.2.

5.1 Image Characterization and Calibration

[**ImChar/JointCal Diagram:** Extract ImChar/JointCal pipelines from "DRP Top-Level Overview" on confluence and expand detail to show data flow and ordering of "Task/Process" boxes.]

The first steps in a Data Release Production characterize the properties of individual exposures, by iterating between pixel-level processing of individual visits ("ImChar", or "Image Characterization" steps) and joint fitting of all catalogs overlapping a tract ("JointCal", or "Joint Calibration" steps). All ImChar steps involve fitting the PSF model and measuring Sources (gradually improving these as we iterate), while JointCal steps fit for new astrometric (WCS) and photometric solutions while building new reference catalogs for the ImChar steps. Iteration is necessary for a few reasons:

- The PSF and WCS must have a consistent definition of object centroids. Celestial positions from a reference catalog are transformed via the WCS to set the positions of stars used to build the PSF model, but the PSF model is then used to measure debiased centroids that feed the WCS fitting.

- The later stages of photometric calibration and PSF modeling require secure star selection and colors to infer their SEDs. Magnitude and morphological measurements from ImChar stages are aggregated the reference catalog in the subsequent JointCal stage, allowing these colors and classifications to be used for PSF modeling in the following ImChar stage.

The ImChar and JointCal iteration is itself interleaved with background matching, described in section 5.2. This allows the best backgrounds and masks to be defined in the BackgroundMatchAndReject before the final Source measurements, image characterizations, and calibrations.

Each ImChar pipeline runs on a single visit, and each JointCal pipeline runs simultaneously on all visits within a single tract, allowing tracts to be run entirely independently.

The final output data products of the ImChar/JointCal iteration are the Source table and CalExp (calibrated exposure) images. The latter includes Image, Mask, Variance, Background, PSF, WCS, and PhotoCalib components that we will track separately.

There are also several intermediate versions of the Source and CalExp data products passed between the ImChar/JointCal pipelines, as well as...

[**TODO:**
finish discussing data products.]

5.1.1 BootstrapImChar

The BootstrapImChar pipeline is the first thing run on each science exposure in a data release. It has the difficult task of bootstrapping multiple quantities (PSF, WCS, photometric calibration, background model, etc.) that each normally require all of the others to be specified when one is fit. As a result, while the algorithmic components to be run in this pipeline are generally clear, their ordering and specific requirements are not; algorithms that are run early will have a harder task than algorithms that are run later, and some iteration will almost certainly be necessary.

A plausible (but by no means certain) high-level algorithm for this pipeline is given below in pseudocode. Highlighted terms are described in more detail below the pseudocode block.

```
def BootstrapImChar(raw, reference):
    # Some data products components are visit-wide and some are per-CCD;
    # these imaginary data types lets us deal with both.
    # VisitExposure also has components; most are self-explanatory, and
    # {mi} == {image,mask,variance} (for "MaskedImage").
    calexp = VisitExposure()
    sources = VisitCatalog()
    snaps = VisitMaskedImageList()  # holds both snaps, but only {image,mask,variance}
    parallel for ccd in ALL_SENSORS:
        snaps[ccd] = [RunISR(raw[ccd]) for snap in SNAP_NUMBERS]
        snaps[ccd].mask = SubtractSnaps(snaps[ccd])
        calexp[ccd].mi = CombineSnaps(snaps[ccd])
    calexp.psf = FitWavefront(calexp[WAVEFRONT_SENSORS].mi)
    calexp.{image,mask,variance,background}
        = SubtractBackground(calexp.mi)
    parallel for ccd in ALL_SENSORS:
        sources[ccd] = DetectSources(calexp.{mi,psf})
    sources[ccd] = DeblendSources(sources[ccd], calexp.{mi,psf})
    sources[ccd] = MeasureSources(sources[ccd], calexp.{mi,psf})
    matches = MatchSemiBlind(sources, reference)
    while not converged:
        SelectStars(matches, exposures)
        calexp.wcs = FitWCS(matches, sources, reference)
        calexp.psf = FitPSF(matches, sources, calexp.{mi,wcs})
        WriteDiagnostics(snaps, calexp, sources)
        parallel for ccd in ALL_SENSORS:
            snaps[ccd] = SubtractSnaps(snaps[ccd], calexp[ccd].psf)
            calexp[ccd].mi = CombineSnaps(snaps[ccd])
            calexp[ccd].mi = SubtractStars(calexp[ccd].{mi,psf}, sources[ccd])
        calexp.{mi,background} = SubtractBackground(calexp.mi)
        parallel for ccd in ALL_SENSORS:
            sources[ccd] = DetectSources(calexp.{mi,psf})
            calexp[ccd].mi, sources[ccd] =
                ReinsertStars(calexp[ccd].{mi,psf}, sources[ccd])
            sources[ccd] = DeblendSources(sources[ccd], calexp.{mi,psf})
            sources[ccd] = MeasureSources(sources[ccd], calexp.{mi,psf})
        matches = MatchNonBlind(sources, reference)
    calexp.psf.apcorr = FitApCorr(matches, sources)
    parallel for ccd in SCIENCE_SENSORS:
        sources[ccd] = ApplyApCorr(sources[ccd], calexp.psf)
    return calexp, sources
```

5.1.1.1 Input Data Product: Raw Raw amplifier images from science and wavefront CCDs, spread across one or more snaps. Needed telescope telemetry (seeing estimate, approximate pointing) is assumed to be included in the raw image metadata.

5.1.1.2 Input Data Product: Reference A full-sky catalog of reference stars derived from both external (e.g. Gaia) and LSST data.

The StandardJointCal pipeline will later define a deeper reference catalog derived from this one and the new data being processed, but the origin and depth of the initial reference catalog is largely TBD. It will almost certainly include Gaia stars, but it may also include data from other telescopes, LSST special programs, LSST commissioning observations, and/or the last LSST data release. Decisions will require some combination of negotation with the LSST commissioning team, specification of the special programs, quality analysis and experimentation with the Gaia catalog, and policy decisions from DM leadership on the degree to which data releases are required to be independent. Depending on the choices selected, it could also require a major separate processing effort using modified versions of the data release production pipelines.

5.1.1.3 Output Data Product: Source A preliminary version of the Source table. This could contain all of the columns in the DPDD Source schema if the MeasureSources is appropriately configured, but some of these columns are likely unnecessary in its role as an intermediate data product that feeds StandardJointCal, and it is likely that other non-DPDD columns will be present for that role.

BootstrapImChar also has the capability to produce even earlier versions of the Source table for diagnostic purposes (see WriteDiagnostics). These tables are not associated with any photometric calibration or aperture correction, and some may not have any measurements besides centroids, and hence are never substitutable for the final Source table.

5.1.1.4 Output Data Product: CalExp A preliminary version of the CalExp (calibrated direct exposure). CalExp is an Exposure object, and hence it has several components. BootstrapImChar is the only pipeline that actually updates all of them. Some CalExp components are determined at the scale of a full FoV and hence should probably be persisted at the visit level (PSF, WCS, PhotoCalib, Background), while others are straightforward CCD-level data products (Image, Mask, Variance).

5.1.1.5 RunISR Delegate to the ISR algorithmic component to perform standard detrending as well as brighter-fatter correction and interpolation for pixel-area variations (Warping Irregularly-Sampled Images). It is possible that these corrections will require a PSF model, and hence must be backed-out and recorrected at a later stage when an improved PSF model is available.

We assume that the applied flat field is appropriate for background estimation.

5.1.1.6 SubtractSnaps

Delegate to the Snap Subtraction algorithmic component to mask artifacts in the difference between snaps. If passed a PSF (as in the second call), also interpolate them by delegating to the Artifact Interpolation algorithmic component.

We assume here that the PSF modeled on the combination of the two Snaps is sufficient for interpolation on the Snaps individually; if this is not true, we can just mask and interpolate both Snaps when an artifact appears on either of them (or we could do per-Snap PSF estimation, but that's a lot more work for very little gain).

5.1.1.7 CombineSnaps

Delegate to the Image Coaddition algorithmic component to combine the two Snaps while handling masks appropriately.

We assume there is no warping involved in combining snaps. If this is needed, we should instead advocate for dropping snaps in favor of a a single longer exposure.

5.1.1.8 FitWavefront

Delegate to the Wavefront Sensor PSF algorithmic component to generate an approximate PSF using only data from the wavefront sensors and observational metadata (e.g. reported seeing).

Processing the wavefront sensors will likely require some form of detection and measurement; we currently consider this to be part of the Wavefront Sensor PSF code, though it may delegate to e.g. Source Detection and/or Single Visit Measurement.

The required quality of this PSF estimate is TBD; setting preliminary requirements will involve running a version of BootstrapImChar with at least mature detection and PSF-modeling algorithms on precursor data taken in crowded fields, and final requirements will require proceessing full LSST camera data in crowded fields. However, robustness to poor data quality and crowding is much more important than accuracy; this stage need only provide a good enough result for subsequent stages to prcoeed.

5.1.1.9 SubtractBackground

Delegate to the Background Estimation algorithmic component to model and subtract the background consistently over the full field of view.

The multiple backgrounds subtracted in BootstrapImChar may or may not be cumulative (i.e. we may or may not add the previous background back in before estimating the latest one).

5.1.1.10 DetectSources Delegate to the Source Detection algorithmic component to find above-threshold regions (Footprints) and peaks within them in a PSF-correlated version of the image.

In crowded fields, each iteration of detection will decrease the threshold, increasing the number of objects detection. Because this will treat fluctuations in the background due to undetected objects as noise, we may need to extend PSF-correlation to the appropriate filter for an image with correlated noise and characterize the noise field from the image itself.

Because we will use wavefront data to constrain the PSF, we also run detection on the wavefront sensors. It is possible that this will require a different algorithmic component if we cannot just treat the wavefront sensors as science sensors with an out-of-focus PSF.

5.1.1.11 DeblendSources Delegate to the Single Visit Deblending algorithmic component to split Footprints with multiple peaks into deblend families.

Because we will use wavefront data to constrain the PSF, we also run deblending on the wavefront sensors. It is possible that this will require a different algorithmic component if we cannot just treat the wavefront sensors as science sensors with an out-of-focus PSF, and we need deblending to extract wavefront information.

5.1.1.12 MeasureSources Delegate to the Single Visit Measurement algorithmic component to measure source properties.

In BootstrapImChar, we anticipate using the Neighbor Noise Replacement approach to deblending, with the following plugin algorithms:

- Centroids
- Second-Moment Shapes
- Pixel Flag Aggregation
- Aperture Photometry (but only for one or two radii)

- Static Point Source Models

Because we will use wavefront data to constrain the PSF, we also run measurement on the wavefront sensors (but probably without any flux measurement algorithms, and perhaps with modified versions of other algorithms). It is possible that this will require a different algorithmic component if we cannot just treat the wavefront sensors as science sensors with an out-of-focus PSF.

5.1.1.13 MatchSemiBlind Delegate to the Single Visit Reference Matching algorithmic component to match source catalogs to a global reference catalog. This occurs over the full field of view, ensuring robust matching even when some CCDs have no matchable stars due to crowding, flux limits, or artifacts.

"Semi-Blind" refers to the fact that the WCS is not yet well known (all we have is what is provided by the observatory), so the matching algorithm must account for an unknown (but small) offset between the WCS-predicted sources positions and the reference catalog positions.

5.1.1.14 SelectStars Use reference catalog classifications and source flags to select a clean sample stars to use for later stages.

If we decide not to rely on a pre-existing reference catalog to separate stars from galaxies and other objects, we will need a new algorithmic component to select stars based on source measurements.

5.1.1.15 FitWCS Delegate to the Single Visit Astrometric Fit algorithmic component to determine the WCS of the image.

We assume this works by fitting a simple mapping from the visit's focal plane coordinate system to the sky and composing it with the (presumed fixed) mapping between CCD coordinates and focal plane coordinates. This fit will be improved in later pipelines, so it does not need to be exact; <0.05 arcsecond accuracy should be sufficient.

As we iterate in crowded fields, the number of degrees of freedom in the WCS should be allowed to slowly increase.

5.1.1.16 FitPSF Delegate to the Full Visit PSF Modeling algorithmic component to construct an improved PSF model for the image.

Because we are relying on a reference catalog to select stars, we should be able to use colors from the reference catalog to estimate SEDs an include wavelength dependence in the fit. If we do not use a the reference catalog early in BootstrapImChar, PSF estimation here will not be wavelength-dependent. In either case the PSF model will be further improved in later pipelines.

PSF estimation at this stage must include some effort to model the wings of bright stars, even if this is tracked and constrained separately from the model for the core of the PSF.

As we iterate in crowded fields, the number of degrees of freedom in the PSF model should be allowed to slowly increase.

5.1.1.17 WriteDiagnostics If desired, the current state of the `source`, `calexp`, and `snaps` variables may be persisted here for diagnostic purposes.

5.1.1.18 SubtractStars Subtract all detected stars above a flux limit from the image, using the PSF model. In crowded fields, this should allow subsequent SubtractBackground and DetectSources steps to push fainter by removing the brightest stars in the image.

Sources classified as extended are never subtracted.

5.1.1.19 ReinsertStars Add stars removed in SubtractStars back into the image, and merge corresponding Footprints and peaks into the source catalog.

5.1.1.20 MatchNonBlind Match a single-CCD source catalog to a global reference frame, probably by delegating to the same matching algorithm used in JointCal pipelines. A separate algorithm component may be needed for efficiency or code maintenance reasons; this is a simple limiting case of the multi-way JointCal matching problem that may or may not merit a separate simpler implementation.

"Non-Blind" refers to the fact that the WCS is now known well enough that there is no significant offset between WCS-projected source positions and reference catalog positions.

5.1.1.21 FitApCorr Delegate to the Aperture Correction algorithmic component to construct a curve of growth from aperture photometry measurements and build an interpolated mapping from other fluxes to the predicted integrated flux at infinity.

5.1.1.22 ApplyApCorr Delegate to the Aperture Correction algorithmic component to apply aperture corrections to flux measurements.

5.1.2 StandardJointCal

In StandardJointCal, we jointly process all of the Source tables produced by running BootstrapImChar on each visit in a tract. There are four steps:

1. We match all sources and the reference catalog by delegating to JointCalMatching. This is a non-blind search; we assume the WCSs output by BootstrapImChar are good enough that we don't need to fit for any additional offsets between images at this stage. Some matches will not include a reference object, as the sources will almost certainly extend deeper than the reference catalog.

2. We classify matches to select a clean sample of low-variability stars for later steps, delegating to JointCalClassification. This uses morphological and possibly color information from source measurements as well as reference catalog information (where available). This step also assigns an inferred SED to each match from its colors; for matches associated with a reference object, whether this supersedes SEDs or colors in the reference catalog is depends on our approach to absolute calibration.

3. We fit simultaneously for improved astrometric solution by requiring each star in a match to have the same position. This may need to correct (perhaps approximately) for centroid shifts due to DCR and/or proper motion; if it does not, it must be robust against these shifts (perhaps via outlier rejection). The models and parameters to fit must be determined by experimentation, but they will represent further perturbation of the WCS fit in BootstrapImChar. This fit generates a new WCS component for each CalExp.

4. We fit simultaneously for photometric zeropoints by requiring each star in a match to have the same flux after applying smoothed monochromatic flat fields produced by the calibration products pipeline. There is a small chance this fit will also be used to further constrain those monochromatic

flat fields. This fit generates a new PhotoCalib component for each CalExp.

In addition to updating the CalExp WCS and PhotoCalib, StandardJointCal generates a new Reference dataset containing the joint-fit centroids and fluxes for each of its match groups as well as their classifications and inferred SEDs.

StandardJointCal may be iterated with RefineImChar to ensure the PSF and WCS converge on the same centroid definitions. StandardJointCal is always run immediately after BootstrapImChar, but RefineImChar or StandardJointCal may be the last step in the iteration run before proceding with WarpAndPsfMatch.

5.1.3 RefineImChar

RefineImChar performs an incremental improvement on the measurements and PSF model produced by BootstrapImChar, using the improved reference catalog, WCS, and PhotoCalib produced by StandardJointCal. Its steps are thus a strict subset of those in BootstrapImChar. A pseudocode description of RefineImChar is given below, but all steps refer to back to the descriptions in 5.1.1:

```
def RefineImChar(calexp, sources, reference):
    matches = MatchNonBlind(sources, reference)
    SelectStars(matches, exposures)
    calexp.psf = FitPSF(matches, sources, calexp.{mi,wcs})
    parallel for ccd in ALL_SENSORS:
        calexp[ccd].mi = SubtractStars(calexp[ccd].{mi,psf}, sources[ccd])
    calexp.{mi,background} = SubtractBackground(calexp.mi)
    parallel for ccd in ALL_SENSORS:
        sources[ccd] = DetectSources(calexp.{mi,psf})
        calexp[ccd].mi, sources[ccd] =
            ReinsertStars(calexp[ccd].{mi,psf}, sources[ccd])
        sources[ccd] = DeblendSources(sources[ccd], calexp.{mi,psf})
        sources[ccd] = MeasureSources(sources[ccd], calexp.{mi,psf})
    calexp.psf.apcorr = FitApCorr(matches, sources)
    parallel for ccd in ALL_SENSORS:
        sources[ccd] = ApplyApCorr(sources[ccd], calexp.psf)
    return calexp, sources
```

This is essentially just another iteration of the loop in in BootstrapImChar, without the WCS-fitting or artifact-handling stages. We assume that we continue to process the wavefront sensors here (because we will use them in the FitPSF step), but it may be that previous processing may be sufficient.

Note that RefineImChar does not update the CalExp's WCS, PhotoCalib, Image, or Variance (and its Mask is only updated to indicate new detections).

5.1.4 FinalImChar

FinalImChar is responsible for producting the final PSF models and source measurements. While similar to RefineImChar, it is run after the WarpAndPsfMatch and BackgroundMatchAndReject pipelines, which provide it with the final background model and and updated mask.

The steps in FinalImChar are identical to those in RefineImChar, with just a few exceptions:

- The background is not re-estimated and subtracted.
- The suite of plugin run by Single Visit Measurement is expanded to included all algorithms indicated in the first column of Figure 5. This should provide all measurements in the DPDD Source table description.
- We also classify sources by delegating to Single Visit Classification, to fill the final Source table's *extendedness* field. It is possible this will also be run during RefineImChar and BootstrapImChar for diagnostic purposes.

5.1.5 FinalJointCal

FinalJointCal is *almost* identical to StandardJointCal, and the details of the differences when surrounding pipelines are more mature and the approach to absolute calibration is more clear. Because it is responsible for the final photometric calibration, it may be need to perform some steps that could be omitted from StandardJointCal because they have no impact on the ImChar pipelines. This could include a role in determining the absolute photometric calibration of the survey, especially if an external catalog (e.g. Gaia) is relied upon exclusively to tie different tracts together.

There is no need for FinalJointCal to produce a new or updated Reference dataset (except for its own internal use), as subsequent steps do not need one, and the DRP-generated reference catalog used by Alert Production will be derived from the Object table.

Unlike StandardJointCal, FinalJointCal is also responsible for applying its improved photometric and astrometric calibrations to raw the Source table generated by FinalImChar, yielding a Source table suitable for database ingest. However, if the final *absolute* calibration is not determined by FinalJointCal, these will need to be further adjusted at a later stage (probably after database ingest).

5.2 Coaddition and Difference Imaging

> **Coaddition, DiffIm Diagram:**
> Extract Coaddition and DiffIm pipelines from "DRP Top-Level Overview" on confluence and expand detail to show data flow and ordering of "Task/Process" boxes.

5.2.1 WarpAndPsfMatch

5.2.2 BackgroundMatchAndReject

5.2.3 WarpAndPsfCorrelate

5.2.4 CoaddTemplate

5.2.5 DiffIm

5.2.6 DecorrelateCoadds

5.3 Object Definition

> **Detection/Association/Deblending Diagram:**
> Extract process_coadds pipeline from "DRP Top-Level Overview" on confluence and expand detail to show data flow and ordering of "Task/Process" boxes.

5.3.1 DeepDetect

5.3.2 DeepAssociate

5.3.3 DeepDeblend

5.3.4 ResolvePatchOverlaps

5.3.5 ResolveTractOverlaps

5.4 Object Characterization

> **Object Characterization Diagram:**
> Extract multifit/forced_photometry pipelines from "DRP Top-Level Overview" on confluence and expand detail to show data flow and ordering of "Task/Process" boxes.

5.4.1 MeasureCoadds

5.4.2 MultiFit

5.4.3 ForcedPhotometry

5.5 Postprocessing

[**Postprocessing Diagram:**
Extract Afterburner pipelines from "DRP Top-Level Overview" on confluence and expand detail to show data flow and ordering of "Task/Process" boxes.]

5.5.1 MOPS

5.5.2 MakeSelectionMaps

5.5.3 Classification

5.5.4 GatherContributed

5.6 UNCAPTURED DEPENDENCIES

- Where does the initial reference catalog at the start of the DRP come from? This could require special observations in commissioning or before the start of the survey, as well as addition algorithms and software. If DRP always uses a reference catalog for star selection in ImChar, we need to actually do the classification for that at some point.

- How do we test all of the wavelength-dependent photometric calibration and PSF stuff on precursor data? Are we going to characterize DECam well enough to just use it directly, or do we need to mock things up or rely more on JointCal?

6 Services for Data Quality Analysis (SDQA)

6.1 Key Requirements

SQDA is a set of loosely coupled services intended to service LSST's quality assessement needs through all phases of Construction, Commissioning and Operations. Consumers of these services may include developers, facility staff, DAC (e.g., L3) users, and the general LSST science user community. Use of these services is intended for routine characterisation, fault detection and fault diagnosis.

- SDQA shall provide sevices for science data quality analysis of Level 1, 2, and Calibration Processing pipelines.
- SDQA shall provide services to support software development in Construction, Commissioning and Operations.
- SDQA shall provide for the visualization, analysis and monitoring capabilities needed for common science quality data analysis usecases. Its inputs may be gathered from SDQA services, the production pipelines, engineering data sources and non-LSST data sources.
- SDQA shall have the flexibility to support execution of ad-hoc (user-driven) tests and analyses of ad-hoc datasets (provided they are supported by the LSST stack) within a standard framework.
- SDQA shall support usecases involving interactive "drill-down" of QA data exposed through its visualisation interfaces.
- SDQA shall allow for notifications to be issued when monitoring quantities that exceed their permissible thresholds and/or have degraded over historical values.
- SDQA shall be able to collect and harvest the outputs and logs of execution of the production pipelines, and extract and expose metrics from these logs.
- SQDA shall make provision to store outputs that are not stored through other LSST data access services.

- SDQA should be deployable as high-reliability scalable services for production as well as allow for core data assessment functionality to be executed on a developer's local machine.
- SDQA shall be architected in a manner that would enable it to be deployable on standard cloud architectures outside of the LSST facilities.

6.2 Key Tasks for Each Level of QA

SDQA system will provide a framework that is capable of monitoring QA information at four different stages of capability and maturity:

QA Level 0	Testing and Validation of the DM sub-system in pre-commissioning
QA Level 1	Real-time data quality and system assesment during commissioning + operations (also, forensics)
QA Level 2	Quality assessment of Data Releases (also, forensics)
QA Level 3	Ability for the community to evaluate the data quality of their own analyses. These should made available as well-documented and deployable versions of core QA Level 0–2 services.

[**Figure summarising QA key tasks:**
Summary figure under construction]

The majority of the work described in this section falls under the 02C.10 WBS (Science Quality and Reliability Engineering). Exceptions are noted in the text as appropriate.

6.2.1 QA Level 0

The first step to good quality data is good quality software. The purpose of QA0 services is to enable testing of the DM software during pre-commissioning as well as validate software improvements during commissioning and operations, quantifying the software performance against known and/or expected outputs.

The core capabilities of QA 0 services are:

6.2.1.1 Continuous Integration Services

- Continuous integration services compile code to uncover syntax errors.
- Builds of references (tags, branches) can happen on a schedule, on developer request or on development events (eg merge to master)
- SDQA provides CI services on multiple reference platforms and uses OS portability testing as a way to ensure the codebase is well engineered for future use.

6.2.1.2 Test execution harness

- A test execution harness runs tests (such as data analysis unit tests) on a regular cadence (eg nightly/weekly/monthly) to allow basic functional checkout of the code. Tests can be added directly by developers and be caused to execute without manual intervention.
- Results from such tests are exposed in such a way to allow summary reports and meaningful failure notifications.

6.2.1.3 Validation Metrics Code

- During Construction, progress towards meeting DM subsystem requirements revolve around the Key Performance Metrics (KPMs) outlined in LDM-240. SDQA implements code to calculate these KPMs. Consult *reference to KPM Verification document* for a list of those metrics and how (and by whom and on what) they will be calculated.
- Additional metrics must be calculated to be met in order for the DM subsystem to demonstrate its operational readiness. The list of those metrics and how (and by whom) they will be calculated will be in *reference to DM Verification Plan CoDR document.* In terms of QA infrastructure, these metrics will not require different capabilities than the KPMs.
- Validation code will be implemented in such a way that it can run inline with normal pipeline processing on developer's laptops.

- Additional metrics may be devised during construction that are helpful to development or algorithm characterisation. SDQA will provide ways of executing that code in a similar way to KPMs, but apps developers may need to contribute the code (or at least document the algorithmic approach) to calculate those metrics.

6.2.1.4 Computational Metrics

- While the scope of this document is the scientific aspects of the pipelines, SDQA must also service non-scientific KPMs such as computational performance characterisation.
- SDQA will provide a capability to instrument the production pipelines to calculate computational performance metrics
- The computational performance metrics that SQDA calculates will be in practice surrogates for the actual computational performance in production since those will depend on the DAC architecture. The purpose of calculating those as part of SDQA is to continuously monitor relative performance to alert the developers that a regression has occured.
- SDQA can calculate modeled system performance from the surrogate computational metrics if a model is provided to it (eg from Architecture).
- A library of those instrumentations will be provided so that they can be mixed and matched to pipelines depending on the performance metric of interest.

6.2.1.5 Curated Datasets

- Part of the process for validating the software and its performance is selecting rich but targeted standardized data sets to generate directly comparable metrics between different versions of the software.
- SDQA will select and curate a combination of simulated and precursor datasets that are modest enough for “canary” test runs but rich enough to characterise the envelope of algorithmic performance.
- SDQA will “publish” (make available) these datasets so developers can run the validation tests directly against them in their own environments.

6.2.1.6 SQUASH - Science Quality Analysis Harness SQUASH is a QA-as-a-service architecture that comprises of the following elements:

1. The execution of simple pipeline workflows for the purposes of QA
2. The construction of those QA workflows with an emphasis on usability (as opposed to performance)
3. The collection and exposure of the results of those runs for further retrieval and analysis
4. A monitoring system to detect threshold trespass and excursions from past trends
5. Exposure of the data for retrieval and to interactive analysis tooling

Notes:

- As construction progresses, first-party DM systems to underwrite the functions of SQUASH will become production ready. In the meantime, basic implementations of minimum viable functionality may be done with boostrap or off-the-shelf solutions either as an interim measure or, in some cases, a more lightweight solution.
- A simple example of a "factory" analysis based on SQUASH is "Calculate the astrometric repeatability on this dataset; display the trend; drill down to to show the historgram of the points that went into calculating this trend".
- An advanced example of a bespoke analysis based on SQUASH is "Display a three-color diagram of the sources in this run; compute the width of point sources in the selected – e.g., blue – part of the locus"
- SQUASH will likely expose results to the LSST Science User Interface for advanced interaction scenarios (both because of the SUI team's front-end expertise but also because they are likely to be similar to science-driven interactions in intent and in execution)

6.2.2 QA Level 1

QA-1 designates the capability to assess data quality in real-time observing modes such as integration, commissioning and operations; if the role of QA-0 is to validate the software, the role of QA-1 is to validate the performance of the facility.

There are two distinct aspects to this capability:

1. Some metric products and services serve standalone user-driven use cases as in QA-0 but with additional data sources, such as the Engineering Facilities Database (EFD), and with real LSST data as opposed to simulated data or pre-cursor data sets. An example use case is "Show the width of point sources on data taken this week in windy conditions with all vents closed versus only the vents in the wind direction as a function of wind speed".
2. Some metric products are produced as part of the routine opreational processing for Level 1 and Calibration pipelines. These will predominantly use the production DM architecture at the DAC and produce metric products either through QA-specific steps in the processing or via the outputs of task instrumentation. An example use case is "show the running **XYZ**"

In the first case the architecture is based on components re-used from QA-0 (with modifications made if made necessary by more stringent performance concerns). Additional out-of-scope (for DM) work may be funded by the Commissioning WBS to support "quick-look" or "comfort display" scenarios where some facility health data is gathered directly from Telescope & Site systems, in which case a component will be added to the QA-0 architecture to support this.

In the second case, the Level 1 DM system software and processing infrastructure at the DAC is used. The Data Access framework (DAX) is used to access all data including values from the EFD and Calibration products.

Note that the EFD is specified to hold all telemetry generated by any observatory system.

All QA-0 components will be involved in QA-1 workflows. The following additional components originate from QA-1 requirements:

6.2.2.1 Alert QA There are two QA components developed for Alert Production:

- A static analysis component that can check, for example, whether the alerts conform to a valid format. This kind of component can be incorporated in the normal Alert Production pipeline.
- A component to receive alerts (akin to a mini-broker) and collect statistics on received events. This would run as a canary node outside LSST facilities to test the alert system is functioning correctly.

6.2.2.2 Validation Metrics Performance As noted, the components of QA-0 to devise key metrics are qualititatively suitable for QA-1. However:

- We expect to make some optimizations to allow them not to consume a significant portion of the 60-second alert time budget.
- In the area of computational performance metrics, additional metrics or instrumentations could be needed due to specific elements of the data center architecture, which at this point is still under design. These will be provided under the Processing Control and Site Infrastructure WBS (02C.07)

6.2.2.3 Dome / Operator Displays Some QA displays may be useful as "comfort displays" (or "facility heartbeats") to staff on site at the telescope, or remote operators. These may require interfaces to data that within the DM system is handled by Data Access Services. If that is the case, this work will be provided from a non-DM (Commissioning) WBS.

6.2.2.4 Telescope Systems Outputs of the SQQA system may be required by the Observatory Control System in order to take some automated action (e.g., reschedule a field). An API will be provided if there is a requirement not already covered by Data Access Services, or DAX may need to be extended to support that access (in the latter case, out of the (02C.06 WBS))

6.2.2.5 Camera Calibration The SDAQ system will also provide QA of Calibration images and products.

- Images taken from the Camera will require "prompt QA" that will run in the quasi-real-time image processing system. Camera is interested in the monitoring infrastructure of SDQA for tracking parameters such as read noice, cross-talk, linearity etc.

- QA of Calibration Products Production data products (i.e., master calibration images and calibration database entries). These are similar in architecture and implementation to other DRP-related tests

- The one exception to the above is the daily daytime/twilight calibration operations prior to night-time observing. QA done for this calibration sequence needs to run under Obsetvatory Control System. There is therefore an explicit or perhaps implicit interface to the OCS that is yet to be specified.

  [I am still unclear as to whether these are all the calibrations we are talking about - I get the dome flats, what about things produced by the spectrometer, the data reconstituted from the lasers etc? - FE]

6.2.2.6 Engineering and Commissioning Some data that is taken specifically for engineering or commissioning purposes will require custom treatment (e.g., an image that is taken with deliberately defocussed optics should not trigger QA alarm and instead should have the noted characteristics of the defocussed sources analysed). While architecturally these are the same as other QA tests, the scope and work for this will be defined as part of the Systems Engineering WBS.

6.2.3 QA2

QA-2 designates the capability to assess the periodic Data Release Products that will be published by LSST. The key aspects that will add on to QA-1 capabilities are (1) the ability to quickly analyze and inspect large data sets; (2) identify failure modes (excursions from expectation or specification) that are rare in QA-0 analysis or real-time QA-1 processing, but represent an identifiable and systematic population or effect on the scale of a full Data Release; and (3) closely interface with calibration efforts in support of the stringent relative color calibration requirements.

In brief, the main focus of QA2 will be to (1) assess the quality of data releases (including the co-added image data products); (2) perform quality assessment for astrometric and photometric calibration and derived products; (3) and look for problems with the image processing pipelines and systematic problems with the instrument.

In addition to the components provided in QA-0, and QA-1, the new components for QA-2 are:

6.2.3.1 DRP-specific dataset

- The scale of a DRP will impose additional performance requirements on the calculation of key performance metrics and associated quality metrics

- The need to drill down with random access to the entire DRP data set will fully exercise the SUIT capabilities.

6.2.3.2 Release data product editing tools (including provenance tracking) Understanding, assuring, and investigating data quality issues will require tight tracking of provenance tracking, particularly as different post-pixel-processing modules may be swapped for each other: e.g., photometric calibration calculations, references catalogs, etc.

6.2.3.3 Interfaces to Workflow and Provenance System(s) If the SDQA system determines that a data (whether science or calibration) is defective, it provides all the information required for the worflow system to take action on this information.

A simple example of this is that a calibration is bad, and it needs to be marked as such so that it is not used in further DRM processing - or a data frame is bad and the compute time should not be wasted processing it further for AP.

A more complex implementation is that a data product previously thought to be good is on futher processing or new tests determined to be bad. In this case will be combined with provenance information to mark *all* data polluted with the bad frame as bad, and provide sufficient information to the workflow system to allow it to trigger the necessary reprocessing with that data excluded.

These are implemented in a manner that is agnostic as to the implementation of the Workflow (e.g., they are values in a database table or API methods that different workflow systems can utilize).

In order to support the interface to the provenance system it would be useful to have some provenance analysis tools, that will allow an operator to query speficially what data went into a particular data product or used a specific data product. These would be very useful to QA but will be provided by the Data Access Services WBS (02C.06).

6.2.3.4 Output Interface to Science Pipelines

- Output interface to Science Pipelines, including from QA database

QA results may provide key feedback to model and parameter choices in the Science Pipelines. The result of the QA system should be made available to the Science Pipelines processing in clearly-tracked analysis and provenance.

6.2.3.5 Comparison tools for overlap areas due to satellite processing Data Release Processing may be distributed across multiple geographic data centers. It is important to verify consistency of the results across these data centers by analyzing both subsets of the overall data processing that are analyzed redundantly by each data center. It will be of particularly importance to test the overlap regions. A framework to define the splits and overlap region and a coherent dashboard and QA configuration to analyze these overlap regions will be key in building confidence in the merged Data Release.

6.2.3.6 Metrics/products for science users to understand quality of science data products (depth mask/selection function, etc.) The Data Release Processing should generate statistics of depth, typical seeing, etc. for regions of the sky; as well as selection functions for the sensitivity to various types of objects. These data products will need to be validated by processing of well-understood data.

6.2.3.7 Characterization report for Data Release

- Each Data Release will be accompanied by a detailed description of its key data statistics, coverage, and quality metrics.
- In addition to static summary numbers and plots, this summary may involve and interactive components. E.g., 10,000 individual is not particularly useful, but an interface to generate plots of interest based

on informed ideas of things to check will be very useful. These interactive components would be the same as those used in the validation of the data release.

6.2.4 QA3

Data quality based on science analysis performed by the LSST Science Collaborations and the community. Level 0-2 visualization and data exploration tools will be made available to the community. Make all results from the above available. Make all of the above components available to some part of the community (could be just affiliated data centers or could be individual scientists) as a supported product. Ingest external science analysis data as Level 3 data products; ingest useful external science analysis tools.

6.2.5 Interactive Visualisation

For QA to happen effectively, before it can be captured to be performed by systems it must be done by humans; the requirements of a QA system do not include all the requirements of a QA Analyst.

Interactive visualisation and free-form data exploration are critical parts of scientific and engineering insight, and for a system the size of LSST it cannot be effectively done on a developer's laptop and/or using traditional tooling. It follows that for the QA process to happen effectively, custom tooling will be necessary to support discovery workflows.

The design of these workflows is out of scope for the this document, which is focused on pipelines generating the products defined in the Data Products Definition Document and the design is described in a document under preparation. But briefly, they fall into three categories:

1. Capabilities that involve structure pre-defined high-semantics displays (e.g., dashboards) with fixed drill-down workflows. These will be serviced by the QA system, specifically the Science Quality Analysis Harness interactive dashboards.

2. Capabilities that are similar to science-user workflows in that they involve generic free-form exploration of the dataset. These will be serviced through the Science User Interface through the Science User Interface Data Analysis and Visualization Tools WBS (02C.05.02), with the Data Access services acting as interface between the SUI and SDAQ.

This is partly to leverage the superior features of the SUI system, and partly to encourage early in-house testing of the SUI features.

3. A more complex case is the situation where curated pre-defined display is desired, but free-form generic exploration of the results is required. In this situation, SDQA will have an API or facility for exporting the former into a tool suitable for the latter. One example of this would be a QA report on, say, a standardised KPM measurement that is produced as a Jupyter Notebook; the user can inspect it, or take it and further interact with the results. Further design is underway in this area.

4. In some cases specific algorithms need to be implemented to drive required visualisation scenarios; these are provided as part of the Alert Production (02C.03) Or Data Release Production (02C.04) as appropriate. An example of this is N-way matching across multiple visits (**??**)

6.2.6 Who validates the validator?

It should be apparent from the above that the QA sevices, while not a critical component in operational terms, nevertheless comprise a system of high semantic value to multiple audiences - dome operators, software developers, science operations staff, data release production engineers and science consumers. Therefore care must be taken to design into the system sanity self-checks to ensure the reliability of its own resuls as well as its upstream pipelines. This section outlines some of the planned features in this area:

6.2.6.1 Intrinsic design features Many of the features described so far provide an alert path for misbehaviours of the QA system. For example a trending excursion for a specific key performance metric could either be due to an algorithmic error or a validation code error. Either way, detection will be a necessary first step to investigation.

6.2.6.2 Known Truth While it may be a matter of debate as to how accurate construction-era simulations are compared to the eventual on-sky data, they are extremely valuable as a fixed source of "known truth" which allow for algorithmically simple QA tests that result in quantifiable performance.

6.2.6.3 Reference Truth [Someone like Z should sign off on this] Comm Cam may allow us to early on develop a small library of representative "reference fields" (eg at different galactic latitudes or ecliptic planes) to provide a minimal standard dataset against which competing algorithmic approaches can be compared (this is similar to the approach taken in Construction with percursor datasets). There would be made available outside the project too alloweing groups working on alternative algoritms and/or implementations to compare their results with the "factory" reductions. Finally, the possibility exists that these reference fields could be unencumbered by proprietary periods so that scientific groups without data rights (and perhaps not even interested in LSST per se) could also utilise them for algorithmic and/or software development.

7 Science User Interface and Toolkit

7.1 Science Pipeline Toolkit (WBS 02C.01.02.03)

7.1.1 Key Requirements

The Science Pipeline Toolkit shall provide the software components, services, and documentation required to construct Level 3 science pipelines out of components built for Level 1 and 2 pipelines. These pipelines shall be executable on LSST computing resources or elsewhere.

7.1.2 Baseline Design

The baseline design assumes that Level 3 pipelines will use the same `Tasks` infrastructure (see the Data Management Middleware Design document; DMMD) as Level 1 and 2 pipelines[7]. Therefore, Level 3 pipelines will largely be automatically constructible as a byproduct of the overall design.

The additional features unique to Level 3 involve the services to upload/-download data to/from the LSST Data Access Center. The baseline for these is to build them on community standards (VOSpace).

7.1.3 Constituent Use Cases and Diagrams

Configure Pipeline Execution; Execute Pipeline; Incorporate User Code into Pipeline; Monitor Pipeline Execution; Science Pipeline Toolkit; Select Data to be Processed; Select Data to be Stored;

7.1.4 Prototype Implementation

While no explicit prototype implementation exists at this time, the majority of LSST pipeline prototypes have successfully been designed in modular and portable fashion. This has allowed a diverse set of users to customize and run the pipelines on platforms ranging from OS X laptops, to 10,000+ core clusters (e.g., BlueWaters), and to implement plugin algorithms (e.g., Kron photometry).

---

[7] Another way of looking at this is that, functionally, there will be no fundamental difference between Level 2 and 3 pipelines, except for the level of privileges and access to software or hardware resources.

8 Algorithmic Components

8.1 Instrument Signature Removal

AUTHOR: Merlin

- Mask defects and saturation
- Assembly
- Overscan
- Linearity
- Crosstalk
- Full frame corrections: Dark, Flats (includes fringing)
- Pixel level corrections: Brighter fatter, static pixel size effects
- Interpolation of defects and saturation
- CR rejection
- Generate snap difference
- Snap combination

8.1.1 AP: just skip some steps?

AUTHOR: Simon

- Indicate steps to be done by camera
- call out other steps that are omitted/modified relative to the DRP version

8.1.2 DRP: do all the steps

AUTHOR: Merlin

8.2 Artifact Detection

8.2.1 Single-Exposure Morphology

AUTHOR: Simon

- Find CRs via morphology.
- Find some optical ghosts (etc?) from bright star catalog and optics predictions.
- Needs to work without PSF (amybe using placeholder PSF), but also make use of PSF if available.

8.2.2 Single-Exposure Aggregation

AUTHOR: Simon

- Find satellites via Hough transform.

8.2.3 Snap Subtraction

AUTHOR: Simon

- All of the above, but improve by looking at both snaps.

8.2.4 Warped Image Comparison

AUTHOR: Jim

- Find more optical artifacts by looking at differences between warped images (this is run during background matching).
- Find transient astronomical sources we don't want to include in coadds.

8.3 Artifact Interpolation

AUTHOR: Jim

- Set mask planes for all artifacts.
- Eliminate small artifacts by interpolating them.
- Uses PSF model as interpolant.

8.4 Source Detection

AUTHOR: Jim

- Detect above-threshold regions and peaks in direct or difference images.
- Needs to work on preconvolved and unconvolved images.
- May need multi-pass variants: detect bright objects first, then faint; detect with approximate PSF, then improved.
- Need to work on wavefront sensors (with out-of-focus PSFs)

8.5 Deblending

AUTHOR: Jim

For templates, try:

- symmetry ansatz with additional regularization
- simultaenous fit of galaxy models
- spline-based models with regularization?
- (multi-coadd only) optimize color uniformity

Will be especially challenging in crowded fields, but it needs to work in that regime as well.

8.5.1 Single Visit Deblending

- Generate HeavyFootprint deblends using only a single image.
- May need to be able to work with approximate/guess PSF, even in crowded fields, if we need to deblend before PSF estimation in DRP.
- May need to work on wavefront sensors (with out-of-focus PSFs)

8.5.2 Multi-Coadd Deblending

- Generate consistent HeavyFootprint deblends from coadds over multiple bands and possibly epoch ranges.

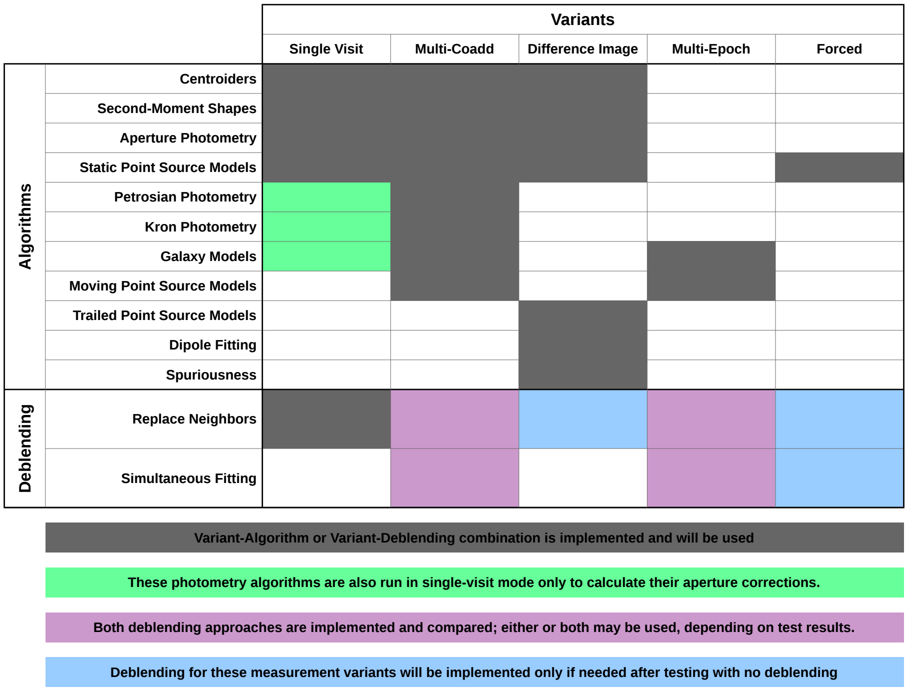

		Variants				
		Single Visit	Multi-Coadd	Difference Image	Multi-Epoch	Forced
Algorithms	Centroiders	⬛	⬛	⬛		
	Second-Moment Shapes	⬛	⬛	⬛		
	Aperture Photometry	⬛	⬛	⬛		
	Static Point Source Models	⬛	⬛	⬛		⬛
	Petrosian Photometry	🟩	⬛			
	Kron Photometry	🟩	⬛			
	Galaxy Models	🟩	⬛		⬛	
	Moving Point Source Models		⬛		⬛	
	Trailed Point Source Models			⬛		
	Dipole Fitting			⬛		
	Spuriousness			⬛		
Deblending	Replace Neighbors	⬛	🟪	🟦	🟪	🟦
	Simultaneous Fitting		🟪		🟪	🟦

⬛ Variant-Algorithm or Variant-Deblending combination is implemented and will be used

🟩 These photometry algorithms are also run in single-visit mode only to calculate their aperture corrections.

🟪 Both deblending approaches are implemented and compared; either or both may be used, depending on test results.

🟦 Deblending for these measurement variants will be implemented only if needed after testing with no deblending

Figure 5: Matrix showing combinations of measurement variants, algorithms, and deblending approaches that will be implemented.

8.6 Measurement

AUTHOR: Jim

8.6.1 Variants

Measurement is run in several contexts, but always consists of running an ordered list of algorithm plugins on either individual objects or families thereof. Each context corresponds to different variant of the measurement driver code, and has a different set of plugin algorithms and approaches to measuring blended objects.

8.6.1.1 Single Visit Measurement: Measure a direct single-visit CCD image, assuming deblend information already exists and can be used to replace neighbors with noise (see 8.6.3.2).

Single Visit Measurement is run in both AP's Single Frame Processing pipeline) and DRP's BootstrapImChar, RefineImChar, and FinalImChar. It must be capable of running on wavefront sensor images, though this may require different plugin algorithms.

The driver for Single Visit Measurement is passed an input/output SourceCatalog and an Exposure to measure. Plugins take an input/output SourceRecord and an Exposure containing only the object to be measured.

8.6.1.2 Multi-Coadd Measurement: Simultaneously measure a suite of coadds representing different bandpasses, epoch ranges, and flavors. This is run only in DRP's MeasureCoadds pipeline.

The driver for Multi-Coadd Measurement is passed an input/output ObjectCatalog and a dict of Exposures to be measured. Plugins take an input/output ObjectRecord and a dict of Exposures, each containing only the object to be measured. Some plugins will also support simultanous measurement of multiple objects, which requires they be provided the subset of the ObjectCatalog to be measured and a dict of Exposures containing just those objects.

8.6.1.3 Difference Image Measurement: Measure a difference image, potentially using the associated direct image as well. Difference image measurement is run in AP's Alert Detection pipeline and DRP's DiffIm pipeline.

The signatures of difference image measurement's drivers and algorithms are at least somewhat TBD; they will take at least a difference image Exposures and a SourceCatalog/SourceRecord, but some plugins such as dipole measurement may require access to a direct image as well. Because difference imaging dramatically reduces blending, difference image measurement may require any approach to blended measurement (though any use of the associated direct image would require deblending).

8.6.1.4 Multi-Epoch Measurement: Measure multiple direct images simultaneously by fitting the same WCS-transformed, PSF-convolved model to them. Blended objects in Multi-Epoch Measurement will be handled by *at least* fitting them simutaneously (8.6.3.3), which may in turn require hybrid galaxy/star models (8.6.3.4). These models may then be used as templates for deblending and replace-with-noise (8.6.3.2) measurement if this improves the results.

Because the memory and I/O requirements for multi-epoch measurement of a single object or blend family are substantial, we will not provide a driver that accepts an ObjectCatalog and measures all objects within it; instead, the peline will submit individual family-level jobs directly to the orchestration layer. The multi-epoch measurement driver will thus just operate on one blend family at a time, and manage blending while executing its plugin algorithms.

Multi-epoch measurement for DRP only includes two plugin algorithms, so it is tempting to simply hard-code these into the driver itself, but this driver will also need to support new plugins in Level 3.

Multi-epoch measurement will also be responsible for actually performing forced photometry on direct images, which it can do by holding non-amplitude parameters for moving point-source models fixed and adding a new amplitude parameter for each observation.

8.6.1.5 Forced Measurement: Measure photometry on an image using positions and shapes from an existing catalog.

In the baseline plan, we assume that forced measurement will only be run on difference images; while forced photometry on direct images will also be performed in DRP, this will be done by multi-epoch measurement.

Because difference imaging reducing blending substantially, forced measurement may not require any special handling of blends. If it does, simultaneous fitting (with point-source models) should be sufficient.

The driver for Forced Measurement is passed an input/output SourceCatalog, an additional input ReferenceCatalog, and an Exposure to measure. Plugins take an input/output SourceRecord, an input ReferenceRecord and an Exposure. If simultaneous fitting is needed to measure blends, plugins will instead receive subsets of the catalogs passed to the driver instead of individual records.

Forced measurement is used by the DRP ForcedPhotometry pipeline and numerous pipelines in AP.

TODO:
Add references to specific AP pipelines that will use forced measurement.

8.6.2 Algorithms

8.6.2.1 Centroids

- should be equivalent to PSF model fit for stars
- use larger weight function (TBD) for extended objects
- need variant that doesn't require a PSF model (or can work with a poor guess) to run before PSF estimation.
- need to have a version (possibly the main version) that works on wavefront sensors

8.6.2.2 Pixel Flag Aggregation

- Compute summary statistics of masked pixels in the neighborhood of the source/object.

8.6.2.3 Second-Moment Shapes

- probably adaptive elliptical Gaussian weights, with fall back to unweightd, PSF-weighted, or some fixed Gaussian
- add regularization for unresolved objects - avoid crazy ellipticities for objects much smaller than PSF
- Should also compute moments of PSF model.
- Need to have a version (possibly the main version) that works on wavefront sensors to characterize the donut-like out-of-focus sources.

8.6.2.4 Aperture Photometry

- Aperture fluxes are computed by summing the total flux within an elliptical region defined on the image.
- Aperture fluxes are computed at a series of logarithmically spaced aperture sizes. Per the DPDD, the total number of apertures will vary depending on the size of the source.

- When computing fluxes for small apertures—for configurable values of "small"–we use sinc interpolation [5]. For large apertures, we use a naive summation of pixel values.
- May need to change ellipticity as a function of aperture radius.
- If run before PSF estimation, will need a variant that does not rely on the PSF model to choose aperture size/ellipticity.

8.6.2.5 Static Point Source Models

- Fit PSF model for flux only (hold center fixed at centroid or reference value)
- Doesn't use per-pixel variances for flux measurement, but might also provide measurement with per-pixel variances (for diagnostics?)

8.6.2.6 Kron Photometry

- Compute Kron radius (hard to make this robust)
- Compute flux in elliptical aperture at Kron radius.

8.6.2.7 Petrosian Photometry

- Compute Petrosian radius. Harder than it seems due to need for improvements to splines? (ask RHL)
- Compute flux in elliptical aperture at Petrosian radius.

8.6.2.8 Galaxy Models

- Some sort of bulge+disk model. Lots of need for experimentation.
- Will Monte Carlo sample in MultiFit (and maybe on coadds, too, if that helps).
- May also fit to PSF-matched coadds for consistent colors.
- Will need to support simultaneous fitting (and sampling).
- Hybrid model candidate

8.6.2.9 Moving Point Source Models

- Fit point source with flux, centroid, parallax, and proper motion parameters.
- May need to support simultaneous fitting.
- Might want to sample this too, at least if we fit it simultaneously with sampled galaxy models.
- Hybrid model candidate

8.6.2.10 Trailed Point Source Models

- Fit PSF convolved with line segment to individual images

8.6.2.11 Dipole Models

- Fit PSF dipole for separation and flux to a combination of difference image and direct image.
- Deblending on direct image very problematic.

8.6.2.12 Spuriousness

- Some per-source measure of likelhood the detection is junk (in a difference image).
- May use machine learning on other measurements or pixels.
- May be augmented by spuriouness measures that aren't purely per-source.

8.6.3 Blended Measurement

- Integrate text from blended-measurement doc here.

8.6.3.1 Deblend Template Projection

8.6.3.2 Neighbor Noise Replacement

8.6.3.3 Simultaneous Fitting

8.6.3.4 Hybrid Models

8.7 Background Estimation

AUTHOR: Simon

- Fit or interpolate large-scale variations while masking out detections.
- Needs to work in crowded fields.
- Needs to work on both difference images and direct images.
- Need to be able to compose backgrounds measured in different coordinate systems on different scales.
- Needs to work on single CCDs for AP even if we use full FoV in DRP.

8.8 Build Background Reference

AUTHOR: Simon

- Given multiple overlapping visit images (already warped to a common coordinate system), synthesize a continuous single-epoch image that can be used as a reference for background matching.

8.9 PSF Estimation

8.9.1 Single CCD PSF Estimation

AUTHOR: Simon

- Fit simple empirical PSF model to stars from a single exposure.
- No chromaticity.
- May use external star catalog, but doesn't rely on one.
- Used only in Alert Production.

8.10 Wavefront Sensor PSF Estimation

AUTHOR: Jim

- Build an approximate PSF model using only the very brightest stars in the wavefront sensors. Because WF sensors are out-of-focus, these stars may be saturated on science CCDs.
- Model can have very few degrees of freedom (very simple optical model + elliptical Moffat/Double-Gaussian?)
- Only needs to be good enough to bootstrap PSF model well enough to bootstrap processing of science images (but it needs to work in crowded fields, too).
- Being able to go to brighter magnitudes may be important in crowded fields because the shape of the luminosity function may make it easier to find stars with (relatively) significant neighbors.

8.10.1 Full Visit PSF Estimation

AUTHOR: Jim

- Decompose PSF into optical + atmosphere.
- May also use wavefront sensors.
- Constrain model with stars, telemetry, and wavefront data.
- Wavelength-dependent.
- Used in RefineImChar in DRP.
- Must include some approach to dealing with wings of bright stars.

8.11 Model Spatial Variation of PSF

8.11.1 Within a CCD

- Estimate PSF at discrete locations using a set of basis functions
- Fit interpolation functions to fit coefficients to enable interpolation

8.11.2 Over a focal plane – Do we need this?

8.12 Aperture Correction

AUTHOR: Jim

- Measure curves of growth from bright stars (visit-level, at least in DRP)
- Correct various flux measurements to infinite (CCD-level)
- Propagate uncertainty in aperture correction to corrected fluxes; covariance is tricky.

8.13 Astrometric Fitting

AUTHOR: Simon

8.13.1 Single CCD

Used by AP, probably (RHL worries we might need full-visit)

- If this uses DRP's internal reference catalog, this does all we need. THIS IS A NEW DEPENDENCY BETWEEN DRP AND AP.

8.13.2 Single Visit

- Fit multi-component WCS to all CCDs in a single visit simultaneously after matching to reference catalog.

8.13.3 Joint Multi-Visit

- Fit multi-component WCS to all CCDs from multiple visits simultaneously after matching to reference catalog.

8.14 Photometric Fitting

AUTHOR: Simon (and Merlin?)

8.14.1 Single CCD (for AP)

- Match to photometric calibration reference catalog
- Calculate single zeropoint using available color terms

8.14.2 Single Visit

- Fit zeropoint (and some small spatial variation?) to all CCDs simultaneously after matching to reference catalog.
- Need for chromatic dependence unclear; probably driven by AP.

8.14.3 Joint Multi-Visit

- Derive SEDs for calibration stars from colors and reference catalog classifications.
- Utilize additional information from wavelenth dependent photometric calibration built by calibration products production.
- Fit zeropoint and possibly perturbations to all CCDs on multiple visits simultaneously after matching to reference catalog.

8.15 Retrieve Diffim Template for a Visit

AUTHOR: Simon

- Determine appropriate template to use
- Generate template for observation (may include DCR correction)

8.16 PSF Matching

AUTHOR: Simon

8.16.1 Image Subtraction

- Match template image to science image, as in Alert Production and DRP Difference Image processing.
- Includes identifying sources to use to determine matching kernel, fitting the kernel, and convolving by it.

8.16.2 PSF Homogenization for Coaddition

- Match science image to predetermined analytic PSF, as in PSF-matched coaddition.

8.17 Image Warping

AUTHOR: Jim

8.17.1 Oversampled Images

Oversampled images are warped to a new WCS and resampled using a two dimensional Lancsoz kernel of configurable order. The baselined default order is 3.

The one dimensional Lancsoz kernel of order a is defined as

$$L(x) = \begin{cases} \mathrm{sinc}(x)\ \mathrm{sinc}(x/a) & \text{if } \ -a < x < a \\ 0 & \text{otherwise.} \end{cases}$$

The two dimensional Lancsoz kernel is $L(x, y) = L(x) \cdot L(y)$.

For each integer pixel position in the remapped image, the associated pixel position in the source image is determined using the source and destination WCS. The warping kernel is then applied to the source image to compute the remapped pixel value. A flux conservation factor is applied based on the relative sizes of the pixel in the source and destination WCS.

For performance reasons, it is desirable to reduce the total number of WCS calculations. It is therefore acceptable to perform the mapping between source and destination images over a regular grid and linearly interpolate between grid points, rather than mapping every pixel independently.

Since chromaticity is accounted for in the PSF rather than the WCS, no special account is taken of color when warping.

Note:
The above describes the current warping implementation in afw. We should identify deficiencies with the current implementation to establish resource requirements.

8.17.2 Undersampled Images

- Can use PSF model as interpolant if we also want to convolve with PSF (as in likelihood coadds). Otherwise impossible?

8.17.3 Irregularly-Sampled Images

- Approximate procedure for fixing small-scale distortions in pixel grid.

8.18 Image Coaddition

AUTHOR: Jim

- Can do outlier rejection (but usually doesn't).
- Needs to propagate full uncertainty somehow.
- May need to propagate larger-scale per-exposure masks to get right PSF model or other coadded quantities.
- Should be capable of combining coadds from different bands and/or epoch ranges ranges as well as combining individual exposures.
- Also needs to support combining snaps

8.19 DCR-Corrected Template Generation

AUTHOR: Simon

- Somwewhat like coaddition, but may need to add dimensions for wavelength or airmass, and may involve solving an inverse problem instead of just compute means.

8.20 Image Decorrelation

8.20.1 Difference Image Decorrelation

AUTHOR: Simon

- Fourier-space (?) deconvolution of preconvolved difference images before measurement - ZOGY as reinterpreted by Lupton (could apply correction in real space, too)

- Need to test with small-scale research before committing to this approach.

8.20.2 Coadd Decorrelation

AUTHOR: Jim

- Fourier-space/iterative deconvolution of likelihood coadds, as in DMTN-15.
- Need to test with small-scale research before committing to this approach.

8.21 Star/Galaxy Classification

AUTHOR: Jim

8.21.1 Single Visit S/G

- Extendedness or trace radius difference that classifies sources based on single frame measurements that can utilize the PSF model. Used to select single-frame calibration stars, and probably aperture correction stars.

8.21.2 Multi-Source S/G

- Aggregate of single-visit S/G post-PSF numbers in jointcal.

8.21.3 Object Classification

- Best classification derived from multifit and possibly variability.

8.22 Variability Characterization

Following the DPDD, lightcurve variability is characterized by providing a series of numeric summary 'features' derived from the lightcurve. The DPDD baselines an approach based on Richards et al. [24], with the caveat that ongoing work in time domain astronomy may change the definition, but not the number or type, of features being provided.

Richards et al. define two classes of features: those designed to characterize variability which is periodic, and those for which the period, if any, is not important. We address both below.

All of these metrics are calculated for both Objects (DPDD table 4, `lcPeriodic` and `lcNonPeriodic`) and DIAObjects (DPDD table 2, `lcPeriodic` and `lcNonPeriodic`). They are calculated and recorded separately in each band. Calculations for Objects are performed based on forced point source model fits (DPDD table 5, `psFlux`). Calculations for DIAObjects are performed based on point source model fits to DIASources (DPDD table 1, `psFlux`). In each case, calculation requires the fluxes and errors for all of the sources in the lightcurve to be available in memory simultaneously.

8.22.1 Characterization of periodic variability

- Characterize lightcurve as the sum of a linear term plus sinusoids at three fundamental frequencies plus four harmonics:

  $$y(t) = ct + \sum_{i=1}^{3}\sum_{j=1}^{4} y_i(t|jf_i) \tag{1}$$

  $$y_i(t|jf_i) = a_{i,j}\sin(2\pi j f_i t) + b_{i,j}\cos(2\pi j f_i t) + b_{i,j,0} \tag{2}$$

  where i sums over fundamentals and j over harmonics.

- Use iterative application of the generalized Lomb-Scargle periodogram, as described in [24], to establish the fundamental frequencies, f_1, f_2, f_3:
  - Search a configurable (minimum, maximum, step) linear frequency grid with the periodogram, applying a $\log f/f_N$ penalty for frequencies above $f_N = 0.5\langle 1/\Delta T\rangle$, identifying the frequency f_1 with highest power;
  - Fit and subtract that frequency and its harmonics from the lightcurve;
  - Repeat the periodogram search to identify f_2 and f_3.

- We report a total of 32 floats:
  - The linear coefficient, c (1 float)
  - The values of f_1, f_2, f_3. (3 floats)
  - The amplitude, $\mathrm{A}_{i,j} = \sqrt{a_{i,j}^2 + b_{i,j}^2}$, for each i, j pair. (12 floats)

- The phase, $\mathrm{PH}_{i,j} = \arctan(b_{i,j}, a_{i,j}) - \frac{jf_i}{f_1}\arctan(b_{1,1}, a_{1,1})$, for each i, j pair, setting $\mathrm{PH}_{1,1} = 0$. (12 floats)
- The significance of f_1 vs. the null hypothesis of white noise with no periodic signal. (1 float)
- The ratio of the significance of each of f_2 and f_3 to the significance of f_1. (2 floats)
- The ratio of the variance of the lightcurve before subtraction of the f_1 component to its variance after subtraction. (1 float)

NB the DPDD baselines providing 32 floats, but, since $\mathrm{PH}_{1,1}$ is 0 by construction, in practice only 31 need to be stored.

8.22.2 Characterization of aperiodic variability

In addition to the periodic variability described above, we follow [24] in providing a series of statistics computed from the lightcurve which do not assume peridoicity. They define 20 floating point quantities in four groups which we describe here, again with the caveat that future revisions to the DPDD may require changes to this list.

Basic quantities:

- The maximum value of delta-magnitude over delta-time between successive points in the lightcurve.
- The difference between the maximum and minimum magnitudes.
- The median absolute deviation.
- The fraction of measurements falling within 1/10 amplitudes of the median.
- The "slope trend": the fraction of increasing minus the fraction of decreasing delta-magnitude values between successive pairs of the last 30 points in the lightcurve.

Moment calculations:

- Skewness.

- Small sample kurtosis, i.e.

  $$\text{Kurtosis} = \frac{n(n+1)}{(n-1)(n-2)(n-3)} \sum_{i=1}^{n} \left(\frac{x_i - \overline{x}}{s} \right)^4 - \frac{3(n-1)^2}{(n-2)(n-3)} \tag{3}$$

  $$s = \sqrt{\frac{1}{n-1} \sum_{i=1}^{n} (x_i - \overline{x})^2} \tag{4}$$

- Standard deviation.

- The fraction of magnitudes which lie more than one standard deviation from the weighted mean.

- Welch-Stetson variability index J [26], defined as

  $$J = \frac{\sum_k \operatorname{sgn}(P_k)\sqrt{|P_k|}}{K},$$

  where the sum runs over all K pairs of observations of the object, where sgn returns the sign of its argument, and where

  $$P_k = \delta_i \delta_j \tag{5}$$

  $$\delta_i = \sqrt{\frac{n}{n-1}} \frac{\nu_i - \overline{\nu}}{\sigma_\nu}, \tag{6}$$

  where n is the number of observations of the object, and ν_i its flux in observation i. Following the procedure described in Stetson [26], the mean is not the simple weighted algebraic mean, but is rather reweighted to account for outliers.

- Welch-Stetson variability index K [26], defined as

  $$K = \frac{1/n \sum_{i=1} N|\delta_i|}{\sqrt{1/n \sum_{i=1} N|\delta_i^2|}},$$

  where N is the total number of observations of the object and δ_i is defined as above.

Percentiles. Taking, for example, $F_{5,95}$ to be the difference between the 95% and 5% flux values, we report:

- All of $F_{40,60}/F_{5,95}$, $F_{32.5,67.5}/F_{5,95}$, $F_{25,75}/F_{5,95}$, $F_{17.5,82.5}/F_{5,95}$, $F_{10,90}/F_{5,95}$
- The largest absolute departure from the median flux, divided by the median.
- The radio of $F_{5,95}$ to the median.

QSO similarity metrics, as defined by Butler & Bloom [10]:

- $\chi^2_{\mathrm{QSO}}/\nu$.
- $\chi^2_{\mathrm{False}}/\nu$.

8.23 Proper Motion and Parallax from DIASources

AUTHOR: Simon

- Fit proper motion and parallax models to all potisions of DIASources belonging to a DIAObject taking into account errors.

8.24 Association and Matching

8.24.1 Single CCD to Reference Catalog, Semi-Blind

AUTHOR: Simon

- Want to match in image coordinates, so also needs to transform reference catalog.
- Run prior to single-visit WCS fitting, with only telescope's best guess as a starting WCS.
- Single CCD form needed by AP.

8.24.2 Single Visit to Reference Catalog, Semi-Blind

AUTHOR: Simon

- Want to match in focal plane coordinates, so also needs to transform reference catalog.
- Run prior to single-visit WCS fitting, with only telescope's best guess as a starting WCS.

8.24.3 Multiple Visits to Reference Catalog

AUTHOR: Jim

- Match sources from multiple visits to a single reference catalog, assuming good WCSs.

8.24.4 DIAObject Generation

AUTHOR: Simon

- Match all DIASources to predicted Solar System object positions and existing DIAObjects and generate new DIAObjects. Definitely run in AP, maybe run in DRP.

8.24.5 Object Generation

AUTHOR: Jim

- Match coadd detections from different bands/SEDs/epoch-ranges, merging Footprints and associating peaks.
- Also merge in DIASources or (if already self-associated) DIAObjects.

8.24.6 Cross-Patch Merging

AUTHOR: Jim

- Resolve duplicates in patch overlap regions by flagging "primary" objects. Difficult due to blending.

8.24.7 Cross-Tract Merging

AUTHOR: Jim

- Resolve duplicates in tract overlap regions by flagging "primary" objects. Difficult due to blending.

8.25 Ephemeris Calculation

AUTHOR: Simon

- Calculate positions for all solar system objects in a region at a given time.

8.26 Make Tracklets

AUTHOR: Simon

- Make all tracklet pairs
- Merge multiple chained observation into single longer tracklets
- Purge any tracklets inconsistent with the merged tracklets

8.27 Attribution and precovery

AUTHOR: Simon

- Predict locations of known Solar System objects
- Match tracklet observation to predicted ephimerides taking into account velocity
- Update SSObjects
- Possibly iterate

8.28 Orbit Fitting

AUTHOR: Simon

- Merge unassociated tracklets into tracks.
- Fit orbits to all tracks.
- Purge unphysical tracks.
- Update SSObjects
- Possibly iterate

9 Software Primitives

9.1 Images

9.1.1 Exposure

Image A 2-d array of calibrated, background-subtracted pixel values in counts.

Mask A boolean representation of artifacts, detections, saturation, and other image. This may include (but is not limited to) a 2-d integer arrays with bits interpreted as different "mask planes"; it may also include using Footprints to describe labeled regions.

Variance A representation of the uncertainty in the image. This includes at least a 2-d array capturing the variance in each pixel, and it may involve some other scheme to capture the variance.

Background An object describing the background model that was subtracted from the image; the original unsubtracted image can be obtained by adding an image of this model to the Exposure's image plane. Backgrounds are more complex than merely an image or even an interpolated binned image; background estimation will proceed in several stages, and these stages (which may happen in different coordinate systems) must be combined to form the full background model.

PSF A model of the PSF; see PSF. This includes a model for aperture corrections.

WCS The astrometric solution that related the image's pixel coordinate system to coordinates on the sky; see WCS.

PhotoCalib The photometric solution that relates the image's pixel values to magnitudes as a function of source wavelength or SED. Some PhotoCalibs may represent global calibration and some may represent relative calibration.

9.2 Tables

9.2.1 Source

9.2.2 Reference

9.2.3 Object

9.2.4 Reference

9.3 Footprints

9.4 Convolution Kernels

Must support correlation as well.

9.5 Basic Statistics

9.6 Point-Spread Functions

9.7 Coordinate Transformations

10 Glossary

API Applications Programming Interface

CBP Collimated Beam Projector

CCOB Camera Calibration Optical Bench

CTE Charge Transfer Efficiency

DAC Data Access Center

DAQ Data Acquisition

DMS Data Management System

DR Data Release.

EPO Education and Public Outreach

Footprint The set of pixels that contains flux from an object. Footprints of multiple objects may have pixels in common.

FRS Functional Requirements Specification

MOPS Moving Object Pipeline System

OCS Observatory Control System

Production A coordinated set of pipelines

PFS Prime Focus Spectrograph. An instrument under development for the Subaru Telescope.

PSF Point Spread Function

QE Quantum Efficiency

RGB Red-Green-Blue image, suitable for color display.

SDS Science Array DAQ Subsystem. The system on the mountain which reads out the data from the camera, buffers it as necessary, and supplies it to data clients, including the DMS.

SDQA Science Data Quality Assessment.

SNR Signal-to-Noise Ratio

SQL Structured Query Language, the common language for querying relational databases.

TBD To Be Determined

Visit A pair of exposures of the same area of the sky taken in immediate succession. A Visit for LSST consists of a 15 second exposure, a 2 second readout time, and a second 15 second exposure.

VO Virtual Observatory

VOEvent A VO standard for disseminating information about transient events.

WCS World Coordinate System. A bidirectional mapping between pixel- and sky-coordinates.

References

[1] M. Ankerst, M. M. Breunig, H.-P. Kriegel and J. Sander, **OPTICS: Ordering Points To Identify the Clustering Structure**, Proc ACM SIGMOD (1999).

[2] P. Antilogus, P. Astier, P .Doherty, A. Guyonnet and N. Regnault **The brighter-fatter effect and pixel correlations in CCD sensors** Journal of Instrumentation, Volume 9, Issue 3, article id. C03048 (2014).

[3] A. Becker et al, **Report on Late Winter 2013 Production: Image Differencing** `http://ls.st/x9f`.

[4] A. Becker, **Report on Summer 2014 Production: Analysis of DCR**, `https://github.com/lsst-dm/S14DCR/blob/master/report/S14report_V0-00.pdf`.

[5] **An algorithm for precise aperture photometry of critically sampled images**, MNRAS 431, 1275–1285, 2013

[6] J. S. Bloom et al, **Automating discovery and classification of transients and variable stars in the synoptic survey era**, PASP 124, 1175–1196 (2012).

[7] J. F. Bosch, **Modeling Techniques for Measuring Galaxy Properties in Multi-Epoch Surveys**, PhD Thesis, University of California, Davis (2011). `http://adsabs.harvard.edu/abs/2011PhDT.......226B`

[8] J. Bosch, P. Gee, R. Owen, M. Juric and the LSST DM team, **LSST DM S13 Report: Shape measurement plans and prototypes**, `https://docushare.lsstcorp.org/docushare/dsweb/ImageStoreViewer/Document-15298`

[9] J. Bosch, **Measurement of Blended Objects in LSST**.

[10] **Optimal Time-Series Selection of Quasars**, ApJ 141, 93 (2011).

[11] **Automated Supervised Classification of Variable Stars I. Methodology**, A&A 475, 1159–1183 (2007).

[12] E. M. Huff et al, **Seeing in the dark – I. Multi-epoch alchemy**, `http://arxiv.org/abs/1111.6958`.

[13] H. Furusawa et al, **Hyper Suprime-Cam Survey Pipeline Description**, `http://hsca.ipmu.jp/pipeline_outputs.pdf`.

[14] M. J. Jee and J. A. Tyson, **Toward Precision LSST Weak-Lensing Measurement. I. Impacts of Atmospheric Turbulence and Optical Aberration**, PASP 123, 596(2011).

[15] M. J. Jee, J. A. Tyson, M. D. Schneider, D. Wittman, S. Schmidt and S. Hilbert, **Cosmic shear results from the Deep Lens Survey. I. Joint constraints on Ω_M and σ_8 with a two-dimensional analysis**, ApJ 765 74 (2013).

[16] J. Kubica et al, **Efficiently Tracking Moving Sources in the LSST**, Bulletin of the American Astronomical Society, 37, 1207 (2005).

[17] D. Lang, D. Hogg, S. Jester and H.-W. Rix, **Measuring the undetectable: Proper motions and parallaxes of very faint sources**, AJ 137 4400–4111 (2009).

[18] R. H. Lupton et al, **SDSS Image Processing II: The *Photo* Pipelines**. `http://www.astro.princeton.edu/~rhl/photo-lite.pdf`

[19] R. Lupton and Ž. Ivezić, **Experience with SDSS: the Promise and Perils of Large Surveys**, Astrometry in the Age of the Next Generation of Large Telescopes, ASP Conferences Series, Vol 338 (2005). `http://adsabs.harvard.edu/abs/2005ASPC..338..151L`

[20] R. Lupton, M. Jurić and C. Stubbs, **LSST's Plans for Calibration Photometry**, July 2015.

[21] L. Denneau, J. Kubica and R. Jedicke, **The Pan-STARRS Moving Object Pipeline**, Astronomical Data Analysis Software and Systems XVI ASP Conference Series, Vol. 376, proceedings of the conference held 15-18 October 2006 in Tucson, Arizona, USA. Edited by Richard A. Shaw, Frank Hill and David J. Bell., p.257.

[22] N. Padmanabhan et al, **An Improved Photometric Calibration of the Sloan Digital Sky Survey Imaging Data**, ApJ 674 1217–1233 (2008).

[23] A. Rasmussen, **Sensor Modeling for the LSST Camera Focal Plane: Current Status of SLAC Originated Code** July 2015. `https://docushare.lsstcorp.org/docushare/dsweb/Get/Document-8590`.

[24] J. Richards et al, **On Machine-learned Classification of Variable Stars with Sparse and Noisy Time-series Data**, ApJ 733 10 (2011).

[25] E. F. Schlafly et al, **Photometric Calibration of the First 1.5 Years of the Pan-STARRS1 Survey**, ApJ 756 158 (2012).

[26] Stetson, P.B., **On the Automatic Determination of Light-Curve Parameters for Cepheid Variables**, PASP 108 851–876 (1996).

[27] C. W. Stubbs, **Precision Astronomy with Imperfect Fully Depleted CCDs – An Introduction and a Suggested Lexicon**, Journal of Instrumentation, Volume 9, Issue 3, article id. C03032 (2014).

[28] A. S. Szalay, A. J. Connolly and G. P. Szokoly, **Simultaneous Multicolor Detection of Faint Galaxies in the Hubble Deep Field**, AJ 117 68–74 (1999).