LARGE SYNOPTIC SURVEY TELESCOPE

Large Synoptic Survey Telescope (LSST)
Data Management

Data Management and LSST Special Programs

M. L. Graham, M. Jurić, K.-T. Lim, and E. Bellm

DMTN-065

Latest Revision: 2019-06-25

Abstract

This document provides an in-depth description of the role of the LSST Project in preparing software and providing computational resources to process the data from Special Programs (deep drilling fields and/or mini-surveys). The plans and description in this document flow down from the requirements in LSE-61 regarding processing for Special Programs. The main target audience is the LSST Data Management (DM) team, but members of the community who are preparing white papers on science-driven observing strategies may also find this document useful. The potential diversity of data from Special Programs is summarized, including boundaries imposed by technical limitations of the LSST hardware. The capability of the planned Data Management system to processes this diversity of Special Programs data is the main focus of this document. Case studies are provided as examples of how the LSST software and/or user-generated pipelines may be combined to process the data from Special Programs.

Change Record

Version	Date	Description	Owner name
1	2017-11-14	Status: internal working document.	Melissa Graham
1	2018-06-17	Updated to finalize and issue.	Melissa Graham

Document source location: `https://github.com/lsst-dm/dmtn-065`

Contents

Data Management and LSST Special Programs

1 Introduction

The main LSST science goals will be met by the Wide-Fast-Deep (WFD) Main Survey, but this is expected to be accomplished with 85–90% of the observing time available over the 10 year survey. The remaining 10–15% of the time will be spent on Special Programs: alternative survey areas and/or observing strategies driven by a specific science goal. A call for white papers that provide scientific motivation and observing strategies for the WFD survey and Special Programs, Document-28382, released in June 2018, has a deadline of November 30 2018. It is conceivable that Special Programs might obtain imaging data that are significantly different from the WFD main survey, and/or requires special processing in order to achieve the program's science goals. This document provides an in-depth description of the role of the LSST Project in preparing to process the data from Special Programs (expanding on the general description in Section 6 of LSE-163). The main target audience is the LSST Data Management (DM) team, but members of the science community who are preparing white paper proposals may also find this document useful.

Special Programs Terminology – The WFD main survey will survey the sky with a series of **visits**, and each field will be visited > 800 times over 10 years. A **standard visit** is composed of 2×15 second exposures (commonly referred to as "snaps") and an **alternative standard visit** is composed of a single 30 second exposure. A **non-standard visit** is any other exposure time(s) or number of snaps. Special Programs are typically divided into two types: **Deep Drilling**, a single pointing for which many exposures are obtained in a relatively short amount of time (e.g., $> 2\times$ as many visits in six months as the WFD will obtain in 10 years); and **Mini-Surveys**, which refer to either new sky areas observed with a WFD-like survey, or sky areas within the WFD but observed with a specialized strategy.

The LSST Project's Role in Processing Special Programs Data – The formal requirements regarding LSST's role in processing Special Programs data are in LSE-61, and the following statements have been derived from those requirements. The LSST Project will not take formal responsibility for specialized data reduction algorithms needed to process data, including images taken in non-standard modes. The term "specialized algorithms" refers to software that is not already within scope of the LSST Data Management (DM) science pipelines, and may include, for example: difference imaging for short exposures in which the PSF is not well-formed, shift-and-stack for faint moving objects, or any software with computational needs that significantly surpass the processing budget per image (compared to the processing of a WFD image). The Project will incorporate Special Programs data into the Prompt and/or Data Release processing pipelines and data products of the WFD Main Survey, such as `Alerts`, `CoAdds`, or `Source` and `Object` catalogs (with appropriate flags; LSST data products are described in LSE-163), whenever this (1) can be accomplished with existing software, and (2) is scientifically beneficial to that data product. The Project will also reconfigure its pipelines to generate separate imaging and catalog data products for Special Programs, whenever this can be accomplished with existing software. Finally, the Project will enable user-generated processing via the Science Platform (LSE-319), which will provide software tools and computational resources for (re)processing LSST data.

Document Overview – The purpose of this document is to assess the potential diversity of imaging data that might be obtained by Special Programs (Section 2), and to explore and clarify whether – and to what extent – DM's planned pipelines and user services will be able to handle this diversity of data (Section 3). As described in the call for white papers on cadence optimization (Document-28382), Special Programs that may require specialized or computationally-intense algorithms to meet their science goals are required to describe how these processing needs will be met. Sections 2 and 3 have been designed to help white paper authors figure out whether their proposed Science Program(s) require processing that would be considered out of DM's scope, and in Section 4 we have provided a series of science-driven case studies as examples.

2 The Potential Diversity of Special Programs Data

To define the extent to which Data Management will be able to process data from Special Programs, a comprehensive understanding of the potential diversity of data from Special Programs – compared to the WFD's sky survey of standard visits – is needed.

To build up this comprehensive understanding, in Section 2.1 we considered any technical limitations that the facility and instrumentation will or might place on the data. We find that the hardware imposes few boundaries on how data can be obtained, but that a high number of filter changes and/or long slews are inefficient due to their large overheads. The minimum exposure time is 1 second (stretch goal: 0.1 seconds), but there is currently a technical boundary that limits the readout rate to 1 every 15 seconds.

Next, we considered the Special Programs that have been openly discussed so far in the Science Community (Appendix A). Based on this consideration, most of the proposed Special Programs are likely to use visits that are similar to the WFD Main Survey, but some will require exposures that are significantly shorter, or that are obtained with a bright sky background during twilight. The cadence and patterns may also differ from the WFD main survey, such as long series of exposures obtained of the same field (i.e., deep drilling), or a strategy optimized to find very fast-moving objects. In addition, images of very crowded fields in the Galactic Plane may be included (at all or more often than the WFD Main Survey). It does not appear that any of the previously proposed Special Programs violate the hardware-imposed boundaries discussed in Section 2.1.

2.1 Hardware Boundaries

Here we consider the technical boundaries on the diversity of data products that are expected to (or may) be imposed by limitations from the camera, telescope, and/or site (boundaries imposed by DM processing capabilities are considered in Section 3).

Filter Changes – The maximum time for filter change is 120 seconds: 30 seconds for the telescope to reorient the camera to its nominal zero angle position on the rotator, and 90 seconds to the camera subsystem for executing the change (OSS-REQ-0293; LSE-30). Assuming that most Special Programs would be designed to keep overheads $< 100\%$ and would be using standard 30 second visits, the filter change time indicates that it is likely that at least 4 exposures in a given filter would be obtained between filter changes – but this is not actually a technical boundary.

Filter Carousel Loads – As described in Document-28382, the filter carousel can hold five of the six LSST filters at a time, and filter loads are done in the day. The system is designed to support 3000 loads and 100000 filter changes in 15 years, which is an average of 17 changes per night (after accounting for filter changes during calibrations). Individual filters will support 30000 changes in 15 years. Based on these technical boundaries, we know that there will never be data in more than five filters in a given night.

Exposure Times – The minimum exposure time is 1 second, with a stretch goal of 0.1 seconds (OSS-REQ-0291; LSE-30). The maximum exposure time is not restricted. The readout time is 2 seconds, and would be significant overhead on short exposures. Images with exposure times < 15 seconds may still have to be separated by 15 seconds for thermal tolerance; i.e., that the minimum readout rate is one image every 15 seconds, regardless of exposure time (OSS-REQ-0291; LSE-30). We therefore consider the 15 interval between images a technical boundary on the potential diversity of data products.

Telescope Slew – As described in Document-28382, large slews would have considerable overheads, but there are no technical boundaries on the size of a single slew or the accrued slew distance.

Telescope Tracking – The requirement that the LSST system be able to perform non-sidereal tracking is set by OSS-REQ-0380 in LSE-30. This capability will include angular rates of up to 220 arcseconds per second in both azimuth and elevation.

Camera Rotation – There requirements on the rotator's capabilities do not set any limits on the per-night or total lifetime rotation (OSS-REQ-0301, -0300; LSE-30) which might put boundaries on the distance between successive visits or the ability to jump between two widely separated fields. JIRA ticket DM-12573 is currently open and asking for clarification from the camera team on this. Until then, we assume there are no technical boundaries imposed by camera rotation constraints on the potential diversity of data products.

3 Processing Special Programs Data with LSST DM Pipelines

In this section we go into greater detail regarding the Projects' role in processing data from Special Programs, as introduced in Section 1. We discuss the ability of the planned DM pipelines to: process diverse imaging data that are unlike the WFD's standard visits in Section 3.1; incorporate Special Programs data into the Prompt and Data Release pipelines and data products for the WFD main survey in Section 3.2; reconfigure the pipelines and generate unique sets of data products for each Special Program in Section 3.3; and enable user-generated pipelines and data products in Section 3.4.

3.1 Processing Diverse Imaging Data Unlike the WFD Main Survey

As described in Section 2, Special Programs will be able – and likely – to request observing modes with shorter (or longer) exposure times, long sequences of visits to the same field, and/or imaging of very crowded fields. We review the capability of DM's planned pipelines to process such each of these kinds of diverse data, keeping in mind that the processing boundaries might ultimately be defined not by what is technically possible, but by the resulting image quality parameters such as the number of stars with sufficient flux for photometric calibration or the source area density (crowding). Furthermore, that these boundaries imposed by the data quality might not be constrained until the final performance of DM's algorithms, as described in the Data Management Applications Design (LDM-151) document, is fully characterized.

3.1.1 Exposure Times

Images which deviate significantly from the 15 second duration for the WFD main survey may encounter issues in the instrument signature removal routine, in the correction for differential chromatic refraction, in the difference imaging analysis pipeline, and/or in the photometric and astrometric calibrations due to a differently sampled set of standard stars per CCD. We discuss shorter and longer exposure in turn.

Shorter Exposures. The camera constraint on the minimum supported exposure time is currently 1 second (stretch goal 0.1 seconds). The minimum exposure time for an image to be successfully reduced with Instrument Signature Removal (ISR) is under consideration (see JIRA ticket DM-12574). Assuming that 1 second exposure can be reduced and calibrated, its detected point sources will span a dynamic range of $r \approx 12.9$ – 21.0 magnitudes. A template image built on 15 second exposures will saturate at $r \approx 15.8$, but this still leaves stars between

15.8–21.0 magnitudes to be used in the PSF-matching (and all other filters have a similarly large overlap). However, in order for an image to be successfully PSF-matched to the template, the PSF must be well formed (no speckle pattern), and have a spatial variance that the pipeline is capable of modeling (be smoothly varying on some minimal scale). As a simple demonstration, Figure 1 shows that perhaps exposure times shorter than 2 seconds do not have a well-formed PSF (using the centroid of a 2D Gaussian fit as a proxy for ”well-formed”). The minimum exposure time for an image to be successfully processed by the Difference Imaging Analysis (DIA) pipeline is currently under consideration (see JIRA ticket DM-12574).

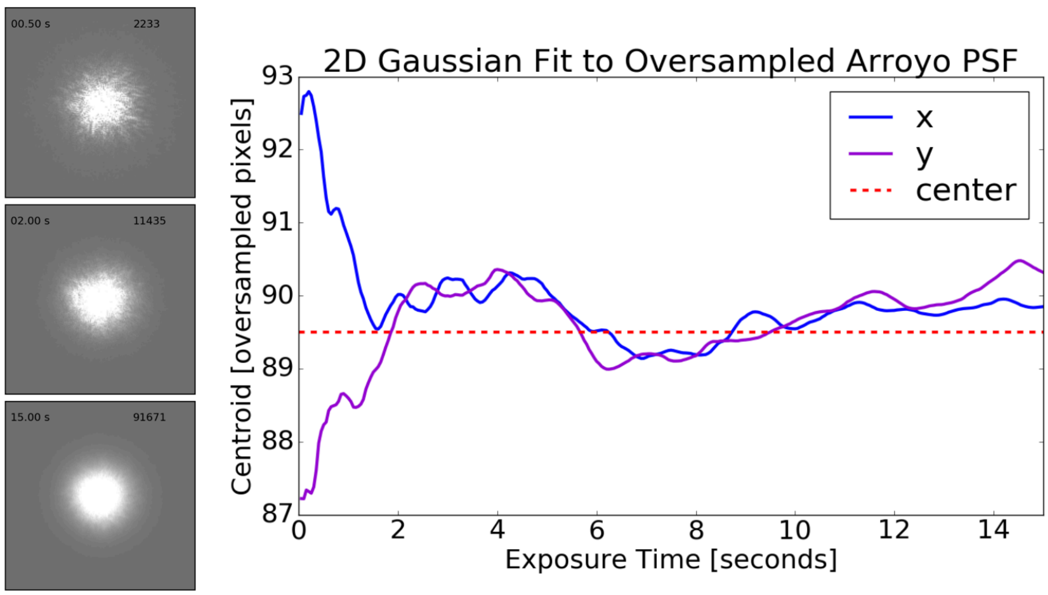

FIGURE 1: At left, Arroyo atmosphere-only simulated PSF for LSST (with oversampled pixels) with exposure times of 0.5, 2, and 15 seconds (top to bottom), courtesy of Bo Xin. At right, blue and purple lines show the location of the centroid derived from a 2D Gaussian fit to the PSF as a function of exposure time, with the red dashed line showing the true center. We can see that for exposure times greater than 2 seconds, the centroid converges near its true value.

Longer Exposures. There is no maximum exposure time specified for an LSST image. Given that the template image will be a stack of at least a year or two of data, processing a 5–10 times deeper single image through the difference imaging pipeline should be fine. However, a 2×150 second exposure would saturate at $r \approx 18.3$, and cosmic-ray rejection completeness might suffer (unknown), which could impact the quality of a difference image and the detected sources. Additionally, any system qualities that vary on short (but > 30 second) timescales could inhibit photometric calibration (e.g., tracking). A potential maximum exposure time from a processing perspective is currently under consideration (see JIRA ticket DM-12574).

3.1.2 Number of Exposures per Visit (Long Sequences of a Single Field)

There is no processing constraint on the number of consecutive exposures that could be obtained of a single field. From a DM perspective, it would be best if these exposures were packaged into visits of no more than 2 exposures per visit, to minimize the need to reconfigure of the pipelines, and because the camera only “clears” between visits.

3.1.3 Images in Very Crowded Fields

The LSST pipelines’ performance in crowded fields is documented in DMTN-077, which finds that, e.g., in Galactic Plane regions with a source density of 500000 sources per square degree, the completeness drops to 50% at 20.2 magnitudes. The slide deck at Document-27962 also describes DM’s plans for processing crowded fields. These may or may not be appropriate for Special Programs data, depending on the science goals.

3.1.4 Twilight Images with a Bright Background

Images obtained during twilight for scientific purposes are also likely to have shorter exposure times, and so the issues described in Section 3.1.1 also apply here. Whether or not bright-background images can (or shall) be fully processed – reduced, calibrated, background-subtracted, and delivered with astrometric and photometric solutions – or whether this will require a User-Generated pipeline, is TBD (see also the example in Section 4.3). This may depend on the exposure time and the number of stars available in the image.

3.1.5 Images Obtained with Non-Sidereal Tracking

Non-sidereal tracking leads to images in which stars are streaked, but the moving object appears as a point source. Full processing – providing reduced, calibrated, background-subtracted images that are delivered with astrometric and photometric solutions – of these images is beyond the scope of the DM pipelines as it would require the development of new algorithms, and will need to be done as a User-Generated pipeline. The first steps of such a pipeline, such as Instrument Signature Removal, will probably be possible to achieve by reconfiguring the relevant DM software tasks.

3.2 Including Special Programs Data in the WFD Main Survey's Data Products

As described in Section 1, the Project may incorporate Special Programs data into the WFD main survey's pipelines and data products whenever this is (1) possible and (2) scientifically beneficial. This will likely be at the discretion of the data quality assessment team during Operations. In the following three sections we project when and how Special Programs data might be incorporated into the pipelines and data products of the Prompt pipeline and the Alert stream (Section 3.2.1) and the annual Data Release pipeline (Section 3.2.2).

3.2.1 The Prompt Pipeline and Alert Generation

It would be beneficial to transient science to include as many LSST images into the Difference Imaging Analysis (DIA) pipeline and the Alert Stream as possible. Only images that can be processed with the Prompt DIA pipeline with the 60-second timeline can contribute to the Alert Stream. As discussed in Section 2, this might prohibit exposures shorter than < 15 seconds and/or visit cadences shorter than 1 per ~ 30 seconds. This might also prohibit the inclusion of very crowded fields that require more computational resources. A field must have an LSST template image to be processed by the DIA pipeline, which would prohibit immediate Alerts from new survey fields. It should be possible to load and use an alternative template image (than what would be used for that field if and when it is covered by the WFD main survey) in the Prompt pipeline for fields from a Special Program.

There may be a couple of issues encountered with Alerts from many consecutive visits of Deep Drilling fields. One is that, since the Alert contains the full record of all associated `DIASources` from the past 12 months (LSE-163), for a Deep Drilling Field with significantly more visits over the year, the size of the Alert might become prohibitively large (TBD). Another is that the self-consistency of the `DIAObjects` catalog may suffer during consecutive visits of a single field. For example, the processing for image 2 of a sequence would begin when the processing for image 1 is only halfway complete. Any new `DIASource` in image 1 that cannot be associated (by coordinate) with an existing `DIAObject` becomes a new `DIAObject`. When this source is again detected in image 2, another new `DIAObject` would be created if the catalog has not yet been updated. The best solution might be to flag `Alerts` which may suffer from incomplete `DIASource`–`DIAObject` associations (see JIRA ticket DM-12574).

The Moving Object Processing System (MOPS) – Since MOPS takes `DIASources` as input, any Special Programs images that can be run through the Alert Pipeline can be ingested by MOPS. As discussed under "Solar System Objects (SSO)" in Appendix A, most of the Special Programs data associated with SSO science will obtain standard visit images anyway.

3.2.2 The Data Release Pipeline (DRP)

This document is not the place for a full consideration of whether or not it would be “scientifically beneficial” to include any Special Programs data in the DRP data products – namely, the deep image CoAdds and their corresponding `Source` and `Object` catalogs (LSE-163) – and we leave that decision for the data quality assessment team in LSST Operations. One example might be when Special Programs data brings additional area up to the same level of depth and cadence as the rest of the WFD main survey. Another may be if including some or all of the shallower Galactic Plane coverage suppresses edge effects or low-order modes in the all-sky photometric solutions.

3.3 Reconfiguring DM Pipeline Components for Special Programs

Whether or not Special Programs images are incorporated into the WFD main survey’s data products, it is anticipated that most of the Special Program’s science goals will require (or benefit from) separate data products (i.e., CoAdds and/or catalogs). For this reason, LSST intends to reconfigure the DM’s pipelines in order to generate unique and separate – but joinable – imaging and catalog products for Special Programs data, whenever possible. In this context, “possible” means that no new algorithms need to be written and that an intensive amount of additional computational resources is not required for the processing.

In a “possible” scenario, DM would assemble a pipeline from existing DM codes in order to process data associated with a given Special Program and build image and catalog products that meet the science needs of that particular program. For example, for a DDF SN survey (see Section 4.2), existing DM codes would be used to: (1) make a deep template image from a certain time window, (2) process standard single visit images, (3) create a nightly CoAdd, (4) run difference imaging analysis, (5) run source detection on the difference images, and (6) create `DIASource` and `DIAObject` catalog equivalents (this example is also given in Section 6 of the DPDD, [LSE-163]). This type of reconfiguration would also be possible to create as a user-generated pipeline (Section 3.4), but having these products provided by the Project ensures a consistent and verified level of quality.

The above statements of intent are derived from the Data Management Subsystem Requirements document, LSE-61, which contains several requirements related to the processing of data from Special Programs (DMS-REQ-0069, 0320, 0321, 0322, and 0344). To ensure that the work-hours needed to reconfigure and test the pipelines, and run them and verify the data products for public release (which may potentially be needed on intermediate timescales that do not coincide with the Prompt/Yearly timescales, e.g., monthly stacks of deep drilling fields), have been included in the personnel budget, JIRA ticket DM-12575 is currently under consideration.

DMS-REQ-0320 states that ”it shall be possible for special programs to trigger their own data processing recipes”. A header keyword identifying an image as related to a Special Program would be sufficient to send it to a dedicated processing pipeline, and would satisfy this requirement. JIRA ticket DM-12576 is currently open to make sure that this happens.

3.4 Support for User-Driven Processing of Special Programs Data

In cases where the science goals of a Special Program require specialized algorithms and cannot be achieved by reconfiguring DM's software, then user-generated pipelines will be needed. Towards this end, LSST DM is making all of its software open-source, and preparing the Science Platform (LSE-319), through which users can access the tools and computational resources to assemble data processing pipelines to achieve their science goals (whether related to Special Programs data or not). During Operations, there will be a method for the system to allocation processing resources in the case of over-subscription. JIRA ticket DM-12577 is currently open to inspire an investigation of whether additional Science Platform capabilities are needed to enable user-driven processing of Special Programs data.

If the user-generated processing pipeline for Special Programs data requires requires significantly more computational resources than have been allocated – where that allocation has been sized approximately, based on image processing for WFD main survey data (i.e., difference imaging, source detection, and/or stacking) – then external computational resources may be necessary. To support such external processing DM intends to make the data and its code base accessible to and exportable by users in the science community.

No user-generated pipeline may contribute Alerts to the Alert Stream, although a separate stream should be possible if the packet and transport formats are adopted (see also Section 3.4).

It is furthermore expected that, over time, some user-designed pipelines might become "adopted", installed and operated (and change controlled) by the LSST Operations team. For both adopted and user-run code, whether they are for Special Programs or WFD survey data, the LSST DM team will encourage and facilitate data product databases that are built with the same schema as – and can easily be joined with – the tables of the Prompt and DRP data products. An alternative option to "adopted" code is "adopted" data products: situations in which user-generated code is run externally, a data catalog is returned to LSST to be ingested, verified, and made public. JIRA ticket DM-12578 is currently open to consider the staffing needs and process for adopting user-generated code or products; this issue is not necessarily limited to Special Programs.

4 Special Programs Processing Case Studies

For further insight to the DM-related needs of potential Special Programs, we can write out all of the data acquisition and processing steps, in order, that some of the proposed Special Programs might use. This kind of thought experiment of describing the reductions and processing could also be a required section of all future white paper proposals. Note that we are not including any analysis in these descriptions, only processing and products. These are not necessarily complete and may even be incorrect in some places, as we are not experts in the science needs of these potential Special Programs; they could use some more thought and input.

Basic steps that we use to describe a processing case study:
Step 1. Data Acquisition.
Step 2. Inclusion in the Prompt Pipeline and Alert Generation.
Step 3. Delivery of LSST Processed Images.
Step 4. Reconfigured Processing Pipelines and Separate Data Products.
Step 5. Inclusion in the DRP Data Products for the WFD Main Survey.
Step 6. User-Generated Pipelines and Products.

4.1 Searching for TNOs with Shift-and-Stack

This Special Programs processing summary is based on Becker et al. (2011) white paper to find TNOs with shift-and stack (SAS) [Document-11013].

Step 1. Data Acquisition.
The observational sequence is triggered. In a single night, the 9 adjacent fields in a 3x3 grid are observed with 336×15 second r-band exposures. This sequence is always repeated 2-3 nights later. This re-visit sequence is repeated 3 more times: 1.5 months, 3 months, and 13.5 months later. Data obtained in the g-band filter is also acceptable. [Document-11013]

Step 2. Inclusion in the Prompt Pipeline and Alert Generation.
Each 2×15 second visit is processed in the Prompt pipeline and Alerts are released within 60 seconds.

Step 3. Delivery of LSST Processed Images.
The raw, reduced, and calibrated exposures and difference images from the Prompt pipeline

are made available within `L1PublicT` (currently 24 hours; LSR-REQ-0104), but this is not very relevant for this program, which requires a year of dispersed observations before the processing pipelines for SAS can be run.

Step 4. Reconfigured Processing Pipelines and Separate Data Products.
Shift-and-stack processing is beyond the scope of DM's algorithms.

Step 5. Inclusion in the DRP Data Products for the WFD Main Survey.
As with all Special Programs data, they might be included in the products of the WFD main survey if DM decides it is beneficial. However, since these images are much deeper than stacks made from the WFD survey, and the strict timing of the observations might lead to their acquisition in sub-optimal conditions, it is unlikely that they would *all* be incorporated.

Step 6. User-Generated Pipelines and Products.
The user-generated pipeline running the shift-and-stack processing will be set up and submitted for batch processing by the user through the Science Platform or on an external processor. Pipeline inputs will be the 336 processed exposures per field per re-visit sequence. The DRP difference imaging routine will be used with the same template tract/patch for all. Custom, user-generated algorithms will shift the exposures and create difference images, then DRP routines can stack and do source detection and characterization and generate an object database. Custom code will derive orbital parameters for the detections and add them to a `SSObjects`-like database.

4.2 Searching for Supernovae in Deep Drilling Fields

Step 1. Data Acquisition.
On a single deep drilling field, the scheduler obtains e.g., 5, 10, 10, 9, and 10 visits with 2×15 second exposures in $grizy$ (or similar for the night's filter set) and a small dither pattern between visits.

Step 2. Inclusion in the Prompt Pipeline and Alert Generation.
Each 2×15 second visit is processed by the Prompt pipeline's DIA, and Alerts are released within 60 seconds. They are flagged to denote the image source is a DDF and that source association might be compromised.

Step 3. Delivery of LSST Processed Images.
The raw, reduced, and calibrated exposures and difference images from the Prompt pipeline

are made available within L1PublicT (currently 24 hours; LSR-REQ-0104).

Step 4. Reconfigured Processing Pipelines and Separate Data Products.
The required data products for this science goal can be met by reconfiguring the DM pipelines. First, a template image for the field will be made using DM stacking algorithms. On nights when this DDF is observed, at the end of the sequence of observations, DM algorithms are used to create a nightly deep stack, PSF-match it with the template, create a deep difference image, run source detection on the differences, and create separate databases of DIAObject, DIASource, and Object that are unique to this DDF. The LSST codes for alert packet and transport could be used to distribute the detected objects e.g., to the same brokers that receive the Alert Stream, or alternative destinations. However, these packets would not be distributed via the LSST Alert Stream, and would need to be identified as, e.g., DDF Alerts. Note that JIRA ticket DM-12585 is currently open to investigate whether or not the internal real/bogus routine be able to run on a nightly CoAdd of deep drilling difference images.

Step 5. Inclusion in the DRP Data Products for the WFD Main Survey.
As with all Special Programs data, they might be included in the products of the WFD main survey if DM decides it is beneficial.

Step 6. User-Generated Pipelines and Products.
For the science goal of searching for supernovae in nightly stacked DDF images, no separate user-generated software appears necessary.

4.3 A Twilight Survey with Short Exposures

Several kinds of twilight surveys with short exposures have been or might be proposed: to put brighter stars (or transients such as supernovae) that saturate in a 15 second image onto the LSST photometric system and/or to observe the Sweetspot, 60 degrees from the sun, for near-Earth objects. The processing case study for these is currently limited by unknowns about the first step: the reduction of processed single visit images.

Step 1. Data Acquisition.
At a specified time (or e.g., 6 degree twilight), the scheduler begins dither pattern of short exposures. Location and exposure times are set by the sky brightness and desired saturation limits.

Step 2. Inclusion in the Prompt Pipeline and Alert Generation.

Pending studies of short-exposure suitability for DIA (see Section 3.1) and scalable processing capabilities to incorporate a faster image-input rate than 1 every 30 seconds, these data could *potentially* be incorporated and spawn Alerts.

Step 3. Delivery of LSST Processed Images.
Pending the issues mentioned above, the raw, reduced, and calibrated exposures and difference images from the Prompt pipeline are made available within `L1PublicT` (currently 24 hours; LSR-REQ-0104).

Step 4. Reconfigured Processing Pipelines and Separate Data Products.
This is officially not determined, but so long as the short-exposure images can be processed and have enough stars for photometric and astrometric calibration, reconfigured DM pipelines will probably be sufficient for creating image and catalog products from this kind of data.

Step 5. Inclusion in the DRP Data Products for the WFD Main Survey.
These short-exposure, high sky background images would not contribute to the DRP data products created for the WFD survey.

Step 6. User-Generated Pipelines and Products.
If short-exposure images cannot be processed with the existing DM algorithms, a user-generated processing pipeline might be needed to reduce the raw data.

Side note: A short-exposure survey of the bright stars of M67, described in Chapter 10.4 of the Observing Strategy White Paper [16], suggests using the stretch goal of 0.1 second exposures or, if that is not possible, *"custom pixel masks to accurately perform photometry on stars as much as 6 magnitudes brighter than the saturation level"*. This would be considered a user-generated algorithm.

4.4 The Galactic Plane Survey for Variable Stars and/or Exoplanets

Step 1. Data Acquisition.
The schedule incorporates fields in the Galactic Plane, and executes 2×15 second visits in these fields (or shorter, for a shallower depth than the WFD main survey).

Step 2. Inclusion in the Prompt Pipeline and Alert Generation.
Each 2×15 second visit is processed in the Prompt pipeline and Alerts are released within 60 seconds. Extremely crowded fields might have to be skipped if they take longer to process and violate the 60 second latency for Alerts.

Step 3. Delivery of LSST Processed Images.
The raw, reduced, and calibrated exposures and difference images from the Prompt pipeline are made available within `L1PublicT` (currently 24 hours; LSR-REQ-0104).

Step 4. Reconfigured Processing Pipelines and Separate Data Products.
The image and catalog products needed for science with the Galactic Plane are very similar to the products of the Prompt and DRP pipelines, so it seems that not much reconfiguration would be needed. The biggest difference might be the incorporation of a user-supplied deblender algorithm optimized for very crowded fields.

Step 5. Inclusion in the DRP Data Products for the WFD Main Survey.
It is quite likely that images from the Galactic Plane will be included into the products of the WFD main survey, as they could e.g., reduce edge effects and help with global photometric classification, but this will depend on deblender performance, and left to the discretion of DM.

Step 6. User-Generated Pipelines and Products.
It seems likely that science users will want to deploy their alternative deblending algorithms on this data set and create their own catalogs.

References

[1] **[Document-11013]**, Becker, A., et al., 2011, *Science White Paper for LSST Deep-Drilling Field Observations Opportunities for Solar System Science*, Document-11013, URL https://ls.st/Document-11013

[2] **[Document-27962]**, Bosch, J., 2018, *DM Plans for Crowded Star Fields*, Document-27962, URL https://ls.st/Document-27962

[3] **[LSE-30]**, Claver, C.F., The LSST Systems Engineering Integrated Project Team, 2018, *Observatory System Specifications (OSS)*, LSE-30, URL https://ls.st/LSE-30

[4] **[Publication-141]**, Dhital, S., et al., 2011, *Science White Paper for LSST Deep-Drilling Field Observations Mapping the Milky Way's Ultracool Dwarfs, Subdwarfs, and White Dwarfs*, Publication-141, URL https://ls.st/Publication-141

[5] Drout, M.R., Chornock, R., Soderberg, A.M., et al., 2014, ApJ, 794, 23 (arXiv:1405.3668), doi:10.1088/0004-637X/794/1/23, ADS Link

[6] **[LSE-61]**, Dubois-Felsmann, G., Jenness, T., 2018, *LSST Data Management Subsystem Requirements*, LSE-61, URL https://ls.st/LSE-61

[7] **[Publication-142]**, Ferguson, H.C., 2011, *Science White Paper for LSST Deep-Drilling Field Observations: LSST Deep Drilling for Galaxies*, Publication-142, URL https://ls.st/Publication-142

[8] **[Publication-143]**, Gawiser, E., et al., 2011, *Science White Paper for LSST Deep-Drilling Field Observations: Ultra-deep $ugrizy$ Imaging to Reduce Main Survey Photo-z Systematics and to Probe Faint Galaxy Clustering, AGN, and Strong Lenses*, Publication-143, URL https://ls.st/Publication-143

[9] Gould, A., 2013, ArXiv e-prints (arXiv:1304.3455), ADS Link

[10] **[Document-28382]**, Željko Ivezić, Jones, L., Ribeiro, T., 2018, *Call for White Papers on LSST Cadence Optimization*, Document-28382, URL https://ls.st/Document-28382

[11] **[LPM-17]**, Ivezić, Ž., The LSST Science Collaboration, 2018, *LSST Science Requirements Document*, LPM-17, URL https://ls.st/LPM-17

[12] Ivezic, Z., et al., 2008, ArXiv e-prints (arXiv:0805.2366), ADS Link

[13] **[LSE-319]**, Jurić, M., Ciardi, D., Dubois-Felsmann, G., 2017, *LSST Science Platform Vision Document*, LSE-319, URL https://ls.st/LSE-319

[14] **[LSE-163]**, Jurić, M., et al., 2017, *LSST Data Products Definition Document*, LSE-163, URL https://ls.st/LSE-163

[15] **[Publication-144]**, Kessler, R., et al., 2011, *Science White Paper for LSST Deep-Drilling Field Observations: Supernova Light Curves*, Publication-144, URL https://ls.st/Publication-144

[16] LSST Science Collaboration, Marshall, P., Anguita, T., et al., 2017, ArXiv e-prints (arXiv:1708.04058), ADS Link

[17] **[DMTN-077]**, Suberlak, K., Slater, C., Ž. Ivezić, 2018, *LSST Fall 2017 Crowded Fields Testing*, DMTN-077, URL https:/wdmtn-077.lsst.io,
LSST Data Management Technical Note

[18] **[LDM-151]**, Swinbank, J.D., et al., 2017, *Data Management Science Pipelines Design*, LDM-151, URL https://ls.st/LDM-151

[19] **[Publication-145]**, Szkody, P., et al., 2011, *Science White Paper for LSST Deep-Drilling Field Observations High Cadence Observations of the Magellanic Clouds and Select Galactic Cluster Fields*, Publication-145, URL https://ls.st/Publication-145

A Previously Proposed Special Programs

In this section we compile information about the science goals and observational methods for Special Programs that have been previously proposed or discussed in the Science Community. We use these to infer the potential deviations from standard visit images, and to get a basic idea of the DM processing needs that would be required to enable the science. The main resources from which we have collected information about the Community's Special Program are: [12]; [11]; the LSST Deep Drilling Field white papers from 2011[1]; presentations by Niel Brandt and Stephen Ridgway at the LSST Project and Community Workshop in August 2016[2]; [9]; and Chapter 10 of [16].

So far, only one aspect of the LSST Special Programs are set: the locations of the four chosen deep drilling fields[3]. There are three mini-survey areas that have been discussed extensively by the Science Community: the North Ecliptic Spur (NES), the South Celestial Pole, and the Galactic Plane (see Figure 8 of [12]). In Table 1 we list the four extragalactic deep drilling fields have already been specified, along with an *incomplete* list of potential mini-surveys that have been openly discussed in the Science Community. In Section 4, we create detailed DM Processing Case Studies for several of these Special Programs in order to identify any potential issues with reconfiguring the DM pipelines to create specific data products for these programs.

TABLE 1: Approved DDF and Incomplete List of Potential Special Programs.

Name	Coordinates	Description
DDF Elias S1	00:37:48, -44:00:00	approved, cadence TBD
DDF XMM-LSS	02:22:50, -04:45:00	approved, cadence TBD
DDF Extended Chandra Deep Field-South	03:32:30, -28:06:00	approved, cadence TBD
DDF COSMOS	10:00:24, +02:10:55	approved, cadence TBD
North Ecliptic Spur		solar system objects (find and characterize)
Galactic Plane		more intensive stellar surveying
South Equatorial Cap		S/LMC and more Galactic science
Twilight		short exposures (0.1s) for bright stars
Mini-Moons		finding mini-moons
Sweetspot		60 deg from Sun for NEOs on Earth-like orbits
Meter-Sized Impactors		detection a week before impact
GW Optical Counterparts		search and recovery
Old Open Cluster M67	dec +12	compact survey above Galactic plane

[1] https://project.lsst.org/content/whitepapers32012

[2] https://project.lsst.org/meetings/lsst2016/sites/lsst.org.meetings.lsst2016/files/Brandt-DDF-MiniSurveys-01.pdf and https://project.lsst.org/meetings/lsst2016/sites/lsst.org.meetings.lsst2016/files/Ridgway-SimulationsMetrics_1.pdf

[3] https://www.lsst.org/scientists/survey-design/ddf

Here we consider a variety of scientific fields in turn, the Special Programs that have been discussed in that Science Community so far, and the implications of these Programs for the diversity of data and data products. Generally, the types of LSST Special Programs that are open for proposals include: (i) additional deep drilling fields; (ii) refined observing strategies for deep drilling fields; (iii) optimized survey areas for the NES, South Pole, and Galactic Plane; (iv) refined observing strategies for the NES, South Pole, and Galactic Plane; and (v) additional mini-surveys (areas and observing strategies).

A Nominal DDF Observing Strategy – Ivezić et al. (2008, [12]; Section 3.1.2) describes a nominal DDF data set as ~ 50 consecutive 15 second exposures in each of four filters, repeated every two nights for four months. Each exposure would have a 5σ limit of $r \sim 24$; the nightly stack would have a limit of $r \sim 26.5$; and the final deep stack of all exposures would have a limit of $r \sim 28$. This description does not comment on the processing mode, but, depending on the science goals the exposures could be done as either a series of 50 non-standard visits (1×15 seconds) or 25 standard visits (2×15 seconds).

Solar System Objects (SSO) – Four of the mini-surveys in Table 1 have science goals related to studies of SSO. Observations of the North Ecliptic Spur area could yield more ≥ 140 m near-earth objects (NEOs) for the final LSST sample (reference: Brandt's talk). The Mini-Moons Mini-Survey aims to find and study the temporarily captured satellites of the Earth (Section 10.2, [16]). The Sweetspot Survey would use twilight fields to find NEOs in Earth-like orbits (i.e., these objects are never in opposition fields, but overhead at sunrise/sunset; Section 10.2, [16]). The Meter-Sized Impactors program would find and track meter-sized impactors < 2 weeks before impact (Section 10.2, [16]). **Summary:** most of these science goals do not seem to require non-standard visits or exposure times, with the exception of the Sweetspot survey which occurs during twilight and thus may require shorter exposures. The cadence and patterns of these mini-surveys may differ from the WFD main survey, especially when very fast-moving objects are sought. From a processing perspective, it seems that many of these science goals will be achievable by using the products of the Moving Object Processing System (MOPS), which runs on the Prompt Pipeline's `DIASource` catalogs after they are updated each night. The exception is finding faint SSOs (e.g., Trans-Neptunian Objects Trojans, asteroids, long-period comets, dwarf planets) through shift-and-stack (SAS) processing [Document-11013], because SAS is not a capability being built within the DM system and cannot be done solely by reconfiguring DM pipelines. An example of user-generated pipeline for SAS is described in Section 4.

Stars in the Milky Way and Magellanic Clouds – As described in [Publication-141], mini-surveys of the Galactic Plane can better distinguish faints stars from faint red galaxies by including at least 3 filters of coverage (e.g., izy; similar to WFD), and could mitigate losses from proper motion and increase the detection rate of stellar flares by obtaining all the images in short time span (i.e., a more concentrated cadence than the WFD). As described in [Publication-145], applying the nominal DDF observing strategy over the full area of the Large and Small Magellanic Clouds can characterize stellar variability to $M_V < 6.5$ on timescales from 15 seconds to 3 days. For this, special co-adds may be required, e.g., *"to reach variability levels of 0.1 to 0.005 mag will require co-adds depending on the timescale of the particular variables"* [Publication-145]. The Twilight survey in Table 1 proposes short exposures to enable bright stars to be put on the same photometric system as the deeper LSST WFD main survey catalog, and enable science that is based on their long monitoring baselines from historical observations. In Chapter 10.4 of [16], a proposed short-exposure survey of M67 would use the camera's stretch goal of 0.1 second exposures or, if that is not possible, *"custom pixel masks to accurately perform photometry on stars as much as 6 magnitudes brighter than the saturation level"*. **Summary:** while some of these science goals can be accomplished with standard visits, MW & L/SMC science goals are likely to request shorter exposure times, perhaps down to 0.1 seconds. These science goals are also likely to propose cadence and filter distributions that are significantly different from the WFD main survey. From a processing perspective, the science goals depending on shorter exposures will only be able to be met by reconfiguring the DM pipelines if the short exposures can be shown to successfully be processed (with, e.g., instrument signature removal); the science goals can likely be met with data products in the same format as the Prompt or DR Pipeline (i.e., `Source` and `Object` catalogs, single visits and deep CoAdds). Although it is not mentioned in the above paragraph, the MW & L/SMC science community is also most likely to require special processing to extract information from saturated stars, which is outside the scope of DM. See Section 4.4 for more detailed DM processing case studies.

Exoplanets – As described in Section 3.1.2 of [12], transiting exoplanets could be detected with the nominal DDF plan, which would allow for 1% variability to be detected over hour-long timescales; a DDF field at Galactic latitude 30 degrees would yield 10^6 stars at $r < 21$ that would have $\mathrm{SNR} > 100$ in each single exposure of the sequence. [9] describes how transits can be extract from a wider-area survey of the Galactic Plane, and how microlensing candidates can be found with ~ 22 mag imaging over the Galactic Plane region every 3-4 days (since microlensing events are slower; these would then require follow-up with external facilities). Dealing with the more crowded fields would be mitigated by the shallower images, in this case. One of the main points of [9] is that the Galactic Plane can yield a lot of science despite the fact that its eventual deep co-adds would be uselessly confusion limited, and therefore should not be skipped. **Summary.** Some of these science goals appear possible with standard visit images, and some might request shorter exposures to avoid confusion in crowded fields when the science can be done with brighter stars. From a processing perspective, the science goals are likely to be achievable with reconfigured DM pipelines, but this depends heavily on performance in crowded fields. See Section 4.4 for a more detailed DM processing case study for Galactic Plane regions.

Supernovae – The nominal DDF plan described in [12], which builds nightly stacks with a limit of $r \sim 26.5$ out of standard visit images, would extend the SN sample to $z \sim 1.2$ and provide more densely sampled light curves for cosmological analyses. The optimal exposure time distribution might be 6, 5, 10, 10, 9, 10 in $ugrizy$ [Publication-144]. High-cadence observations of DDF would be the only way to detect fast transients, particularly extragalactic novae, some tidal disruption events, optical counterparts to gamma-ray bursts, and peculiar SNe [5]. Generating the best-possible individual SN light curves for cosmological analyses requires building special, deep-as-possible, SN-free host galaxy images and using them as a template. This will also be necessary for studying SNe that appear in the template image; i.e., that last > 1000 days. These are mostly Type IIn, probably explosions of massive stars into dense circumstellar material, which are not used for cosmology but rather to study late-stage stellar evolution and mass loss. SN-free images will also be needed to measure correlated properties for cosmology and to do host-galaxy science. The latter, specifically the ”characterization of ultra-faint SN host galaxies”, is also mentioned in the Galaxies DDF WP [Publication-142]. Short-exposure observations of bright, nearby SNe may also be useful to include near-peak photometry in the LSST magnitude system, and enable full light-curve analyses. **Summary.** All of these science goals appear possible with standard visit images (with the exception of a target-of-opportunity short-exposure program to observe bright SNe). From a processing perspecitve, the science goals appear to be accessible with reconfigured DM pipelines to stack and difference the data. In particular, the DRP codes to create ”transient-free CoAdds” will be suitable for generating the SN-free templates for DDF, as they will do for the Main Survey images. See also Section 4.2 for a DM processing case study to find SNe in a DDF.

Galaxies – The additional depth of a DDF may provide access to a larger collection of low-μ objects. [Publication-142] mentions ”identification of nearby isolated low-redshift dwarf galaxies via surface-brightness fluctuations” and ”characterization of low-surface-brightness extended features around both nearby and distant galaxies”. The DDF stacks could also be used to characterize of high-z clusters, although this ability might depend on deblending extended objects.

Also, the DDF observations, when combined with the WFD, allow for AGN monitoring on a variety of timescales in well-characterized galaxies [Publication-142, Publication-143]. **Summary.** As with the SN science goals, these use standard visit images and reconfigured DM pipelines to make deep CoAdds and extract sources. In addition, it seems likely that user-generated algorithms that are optimized to detect and characterize particular types of faint extended sources will be needed, and these are beyond the scope of DM.

Weak Lensing – The deeper imaging from DDFs can help with shear systematics and the effects of magnification in the analysis of WFD data (community forum, Jim Bosch). **Summary.** As with the SN and Galaxies DDF-related science goals, these use standard visit images and reconfigured DM pipelines can be used to make deep CoAdds and extract sources, as Jim notes.