From FITS to SQL
Loading and Publishing the SDSS Data

Ani Thakar and Alex Szalay (JHU), Jim Gray (Microsoft Research)

Sloan Digital Sky Survey

Abstract

The Sloan Digital Sky Survey Data Release 1 (DR1) contains nearly 1 TB of catalog data published online as the Catalog Archive Server(CAS) and accessible via the **SkyServer** web interface. The DR1 CAS is the end product of a data loading pipeline that transforms the FITS file data exported by the SDSS Operational Database or OpDB, converts it to CSV (comma separated values), and loads it into a MS Windows-based relational database management system (SQL Server DBMS).

Loading the data is potentially the most time-consuming and labor-intensive part of archive operations, and it is also the most critical: it is realistically your one chance to get the data right. We have attempted to automate it as much as possible, and to make it easy to diagnose data and loading errors. We describe this pipeline, focusing on the highly automated SQL data loader framework (sqlLoader) - a **distributed workflow system** of modules that check, load, validate and publish the data to the databases. The workflow is described by a **directed acyclic graph (DAG)** whose nodes are the processing modules, and it is designed for parallel loading on a cluster of load-servers. The pipeline first reads the data from Samba-mounted CSV files on the LINUX side and stuffs it into the SQL databases on the Windows side.

The **validation** step, in particular, represents a systematic and thorough scrubbing of the data before it is deemed worthy of publishing. The **finish** step merges the different data products (imaging, spectra, tiling) into a set of linked tables that can be efficiently searched with specialized indices and pre-computed joins.

We are in the process of making the sqlLoader generic and portable enough so that other archives may adapt it to load, validate and publish their data.

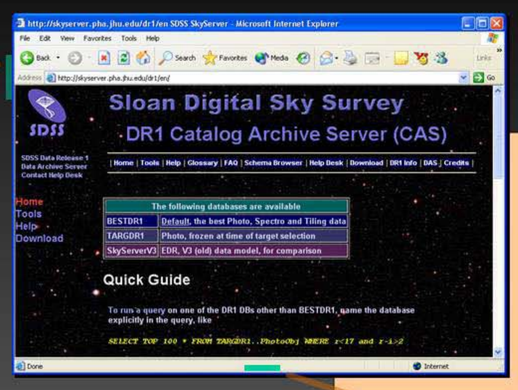

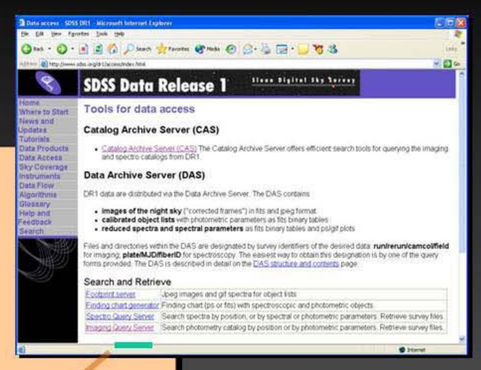

SDSS Data Release 1
http://www.sdss.org/dr1/

Coverage

2100 sq.deg (imaging), 1400 sq.deg. (spectra)
20% of total survey area
5-6 times size of EDR (Early Data Release)

Catalog data
Catalog Archive Server (SkyServer)
http://skyserver.pha.jhu.edu/dr1/

Released June 2003, 1 Tbyte DBMS
85M photometric (image) objects
160k spectroscopic objects
2 versions of data: BEST and TARGET
Relational Data Model (tables)

Image data
Data Archive Server
http://www.sdss.org/dr1/access/

Released April 2003, 2-3 Tbyte flat files
Atlas images (cutouts) – PNG, FITS
Spectra - FITS, GIF
Corrected frames - FITS
Binned images - FITS
Mask images - FITS

The Loading Process

The basic processing entity is a *task*. A task is started when a data *chunk* is exported by the OpDB. Exported chunks are converted to CSV format, and are contained in a single directory. There are several different export types: TARGET, BEST, RUNS, PLATE and TILING. Each task comes with a id number that is unique within its category.

The loading process consists of *steps*. The first step is to load each chunk of data into a separate *task DB*, containing only a thin set of indices. Then we *validate* the data. This includes verifying that there are no primary key collisions and all foreign keys point to a valid record. We build several ancillary tables for spatial searches (HTM, Neighbors, etc.) After the validation step we *publish* the data: we perform a DB-to-DB copy, where the target is the final production database. After publishing it, we make a backup of the task DB.

At the very end, all the different datasets are merged together in the *finish* step, and indices are created for efficient data mining.

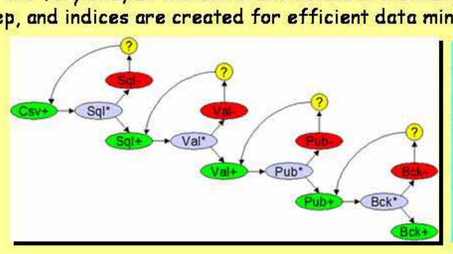

A state-machine representation of the loading process. Each step is a sequence of rather complex steps in itself. The yellow question marks represent a manual Undo step, which is performed as necessary.

DR1 Schema

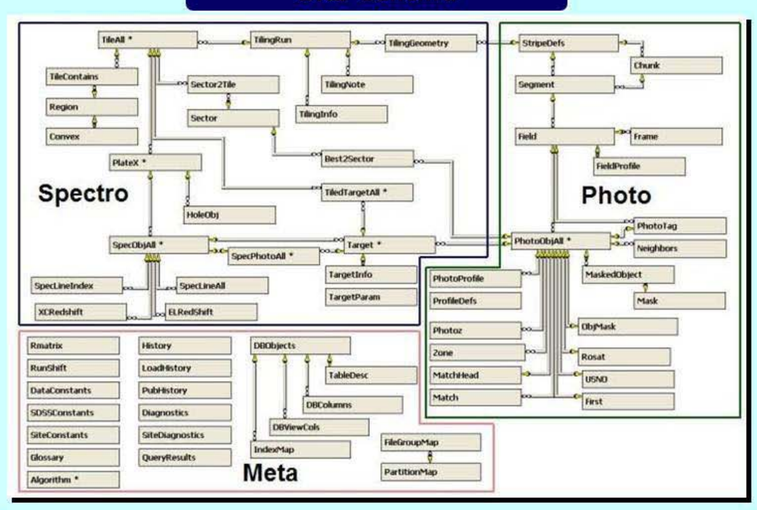

Distributed Loading

Loading a Terabyte or more of data is a time-consuming process even with fast disks, and parallelization of the loading steps is a big help, especially as we get into the multi-TB data volumes of future SDSS releases.

The *load*, *validate* and *publish* steps in the sqlLoader are fully parallelizable and can be executed in a distributed configuration with a cluster of load-servers.

Distributed loading makes use of the following SQL Server features:
- **linked servers** with
- **distributed views** and
- **distributed transactions**.

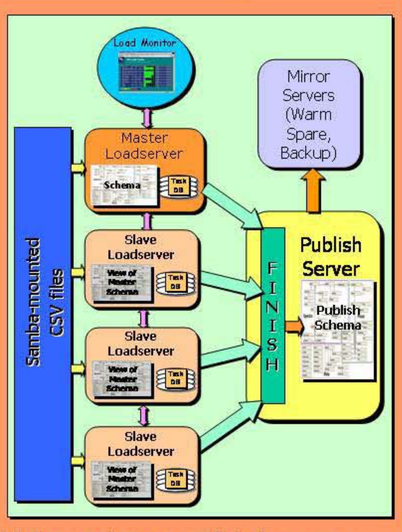

After loading, validating and publishing is done in parallel, the merging of the parallel data streams and the finish step are performed sequentially on the publish server.

Data Export Pipeline

The data is exported from the OpDB in the form of FITS files that are organized into blocks called *chunks*. Each chunk is the result of a block of data being resolved and exported by the OpDB. Four different datasets are exported: two for imaging data and one each for spectroscopic and tiling data. The imaging data is resolved at least twice – once when the spectroscopic targets are chosen, and once when the data is recalibrated with the latest, greatest calibration. These datasets are called TARGET and BEST respectively.

All the data must be converted to CSV format so it can be stuffed into the databases using **bulk insertion**. The CSV files also serve as a blood-brain barrier between the LINUX and Windows worlds. The CSV files are Samba-mounted and loaded into temporary DBs before being bulk-copied to their final destination as part of the *publish* step. All the tasks and steps are logged into a logDB that is queried by the Load Monitor to generate the various logs that it provides.

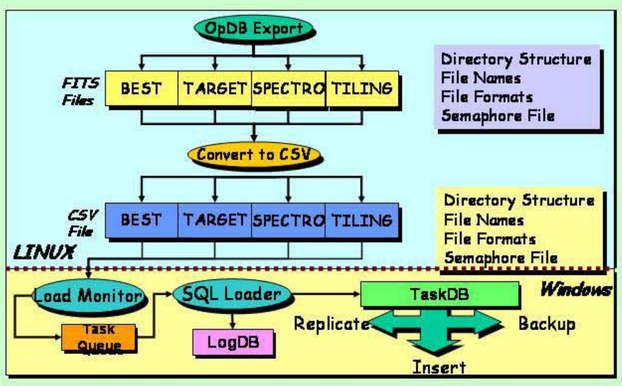

The Load Monitor

The Load Monitor is the admin web interface to the sqlLoader. It enables job submission, control and tracking via a user-friendly GUI. Loading jobs (tasks) can be submitted either a single chunk at a time or bulk-uploaded with the file upload feature. Tasks can be monitored at several levels, and information is available on the status of the individual files being loaded, detailed logs of each step in the task, and separate listings of errors and warnings encountered. Tasks are listed in a task table that shows the status of each step in the task at a glance in a color-coded chart. Sample Load Monitor screens are shown below to illustrate the features that are available.

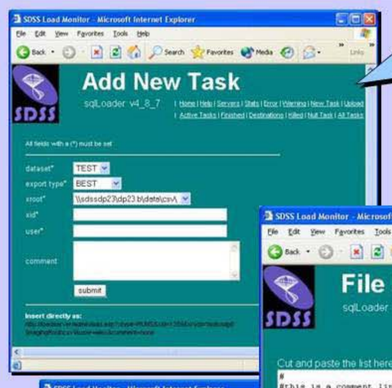

Example of an **Add New Task** page. The user must enter the target database (the databases available to publish to), the dataset this chunk goes into (BEST, TARGET, PLATES, TILING), the source path for the input CSV files, an unique chunk identifier, the name of the user and an optional comment field. All of this information can also be scripted in an upload file, as shown in the example of the **File Upload** page below.

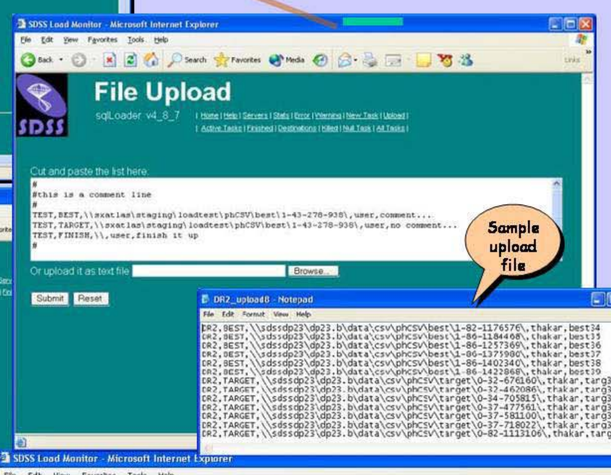

Active Tasks listing shows all the tasks that are currently executing. The task status is displayed as a color code for each step (*export*, *check*, *build*, *validate*, *backup*, *detach*, *publish*, *cleanup* and *finish*). Amber means the step is in progress, green means it is done, red means it failed and purple means that the task was killed in that step.

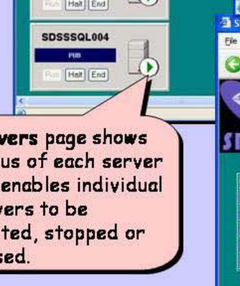

Servers page shows status of each server and enables individual servers to be started, stopped or paused.

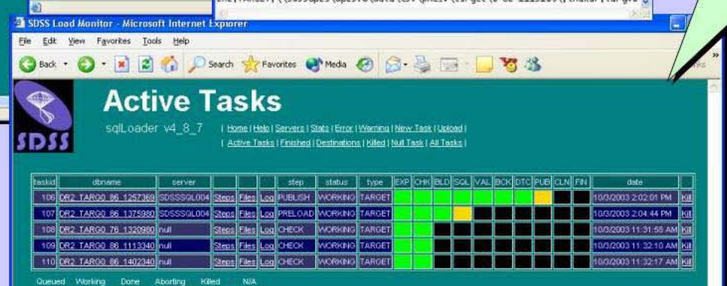

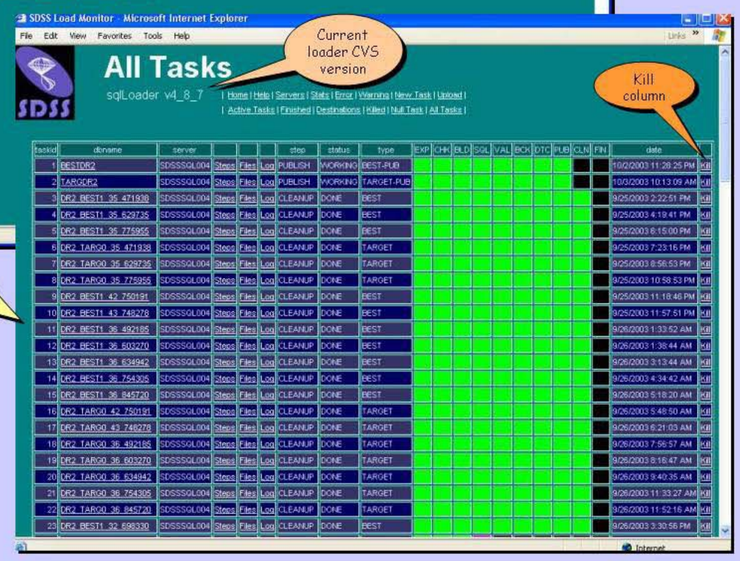

All Tasks listings shows the status of every task ever submitted to the loader. Tasks are numbered in the order they are submitted. The user can kill a given task by clicking on the **Kill** column for that task. All kill commands are checked for confirmation before being executed.

Data Validation

Validation is perhaps the most important step in the loading process. The speed, integrity and convenience that databases offer come at a price: data once published cannot be retracted or corrected easily. This is not only because the loading process itself is difficult and arduous, but also because the data must always be available once science has been done with it. Hence it is crucial to get the data right the first time.

The validate step in sqlLoader represents a systematic scrubbing and sanity-check of the data, from a scientific as well as data integrity point of view. The figure on the right shows the various operations that are performed on the data.

The primary and foreign key tests are run on all the tables. The photo (imaging) and spectro tables are tested for HTM IDs, which are 64-bit IDs that provide fast spatial indexing according to the Hierarchical Triangular Mesh indexing scheme. The image data also has parent-child relationships defined for deblended objects. The referential integrity of these is also checked as part of the validation. Finally, the consistency of counts of various quantities is checked.

This validation process has proven invaluable in finding numerous inconsistencies and errors in the data and catching them early, during the testing of DR1 rather than after the data is published.

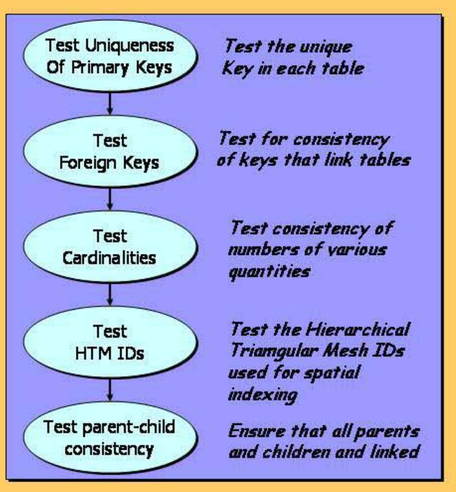

Help and Documentation

Online help is bundled with the sqlLoader product, and includes a user guide that describes how to set up and run the sqlLoader. The troubleshooting section lists solutions to typical problems.

There is also a document that details each step in the loader processing. When compared with the log for a loader task in the Load Monitor, this enables the user to troubleshoot a problem that causes the loader to get stuck.

The sqlLoader has enabled the entire loading for DR1 to be completed largely as a turnkey operation with very little human intervention.