Astronomy and Computing 23 (2018) 27–39

Contents lists available at ScienceDirect

Astronomy and Computing

journal homepage: www.elsevier.com/locate/ascom

Full length article

Fast in-database cross-matching of high-cadence, high-density source lists with an up-to-date sky model

B. Scheers [a,*], S. Bloemen [b,c], H. Mühleisen [a], P. Schellart [d,b], A. van Elteren [e], M. Kersten [a], P.J. Groot [b]

[a] *CWI—Centrum Wiskunde & Informatica, PO Box 94079, 1090 GB Amsterdam, The Netherlands*
[b] *Department of Astrophysics, IMAPP, Radboud University, 6500 GL Nijmegen, The Netherlands*
[c] *NOVA Optical InfraRed Instrumentation Group, Oude Hoogeveensedijk 4, 7991 PD Dwingeloo, The Netherlands*
[d] *Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA*
[e] *Leiden Observatory, Leiden University, PO Box 9513, 2300 RA Leiden, The Netherlands*

ARTICLE INFO

Article history:
Received 9 June 2017
Accepted 28 February 2018
Available online 12 March 2018

Keywords:
Telescopes
Time-domain astrophysics
Astronomical databases
Surveys
Catalogues
Database query processing

ABSTRACT

Coming high-cadence wide-field optical telescopes will image hundreds of thousands of sources per minute. Besides inspecting the near real-time data streams for transient and variability events, the accumulated data archive is a wealthy laboratory for making complementary scientific discoveries.

The goal of this work is to optimise column-oriented database techniques to enable the construction of a full-source and light-curve database for large-scale surveys, that is accessible by the astronomical community.

We adopted LOFAR's Transients Pipeline as the baseline and modified it to enable the processing of optical images that have much higher source densities. The pipeline adds new source lists to the archive database, while cross-matching them with the known catalogued sources in order to build a full light-curve archive. We investigated several techniques of indexing and partitioning the largest tables, allowing for faster positional source look-ups in the cross matching algorithms. We monitored all query run times in long-term pipeline runs where we processed a subset of IPHAS data that have image source density peaks over 170,000 per field of view (500,000 deg^{-2}).

Our analysis demonstrates that horizontal table partitions of declination widths of one-degree control the query run times. Usage of an index strategy where the partitions are densely sorted according to source declination yields another improvement. Most queries run in sublinear time and a few ($<$ 20%) run in linear time, because of dependencies on input source-list and result-set size. We observed that for this logical database partitioning schema the limiting cadence the pipeline achieved with processing IPHAS data is 25 s.

1. Introduction

High-cadence astronomy is a relatively new field in observational astronomy. Advances in hardware and software technology have made it possible to stream large volumes of observational data over fast links to clusters of computers that, in general, process the data in one or more automated pipelines for scientific analysis. The time available to do real-time analysis is limited by the cadence of the instrument. Therefore, additional and complementary scientific data analyses are forced to shift to non-real time environments. Here, all data accumulates over time and the growth may vary in the range of 0.1–100 PB/yr (Becla and Wang, 2014). These volumes clearly challenge many aspects of contemporary data management systems, which is also recognised by Ivezić et al. (2017).

Several instruments have shown impressive demonstrations of charting the sky down to a timescale of seconds, e.g., the international LOFAR telescope (van Haarlem et al., 2012), the Murchison Wide-field Array (MWA; Tingay et al., 2013), the Australian Square Kilometre Array Pathfinder (ASKAP; Murphy et al., 2013). High-cadence observations in image-domain astronomy, where sky regions are revisited many times in relatively short periods, produce overwhelmingly large amounts of data. Optical and radio telescopes planned for the next decade will generate even larger continuous data streams, e.g., the Large Synoptic Survey Telescope (LSST; Lazio et al., 2014; Juric and Tyson, 2015), the additional Ground-based Wide Angle optical Camera (GWAC; Cordier et al., 2015) of the Space-based multiband astronomical Variable

* Corresponding author.
E-mail address: bscheers@cwi.nl (B. Scheers).

https://doi.org/10.1016/j.ascom.2018.02.006

Objects Monitor (SVOM), BlackGEM (Bloemen et al., 2015), the Square Kilometre Array (SKA; Broekema et al., 2012).

Although high-cadence instruments are specifically designed to carry out their own unique science, they share similar observational strategies. The main ones being: high-speed, wide- or all-sky surveys, searching for transient and variable sources on a variety of time scales and gradually archiving full-source light-curve catalogues. In this respect, the archive is considered the new Big Data laboratory, equipped for making scientific discoveries in complex structured data. However, such discoveries are only possible when the infrastructure and software tools allow continuous and simultaneous data mining, statistical modelling, machine learning and ad-hoc querying.

The optical Sloan Digital Sky Survey (SDSS; York et al., 2000; Alam et al., 2015) was the first instrument to seriously integrate a database system into its survey design. It uses a database-centric computing approach for their large-scale scientific datasets. SDSS data are cumulatively released to the public in roughly annual cycles. In this respect, SDSS is a low-cadence instrument since the yearly updates of the full-source catalogue make the database essentially static.

On account of Gray's law to ship computations to the data instead of data to the computations (Szalay and Blakeley, 2009) many algorithms are designed to run *inside* the database engine. Another design rule includes knowledge of the 50 most frequent and intensive queries. Since astronomical pipelines process the data in a structured way, this allows one to optimise execution plans for known queries.

In the radio regime, the automated Transients Pipeline (TraP) of the international LOFAR telescope adopted many database techniques from SDSS (Swinbank et al., 2015). The TraP applies source finding and fitting to calibrated radio images after which, per image, all image and source properties, i.e. the *source list*, are handed over to the database. Note that the images themselves are not stored in the database. The loop of tasks of the TraP database consists in total of about 50 queries which can be divided into four successive steps, all executed in bulk mode:

1. load source list
2. cross-match source list with catalogue of known sources
3. update catalogue: maintain up-to-date statistical sky model
4. find/identify transient and variable sources or other significant deviations from the sky model

Typical source lists for LOFAR survey-mode observations do not exceed 10^3 entries, whereas averages are less than 10^2 for cadence modes as high as 10 s (Swinbank et al., 2015). The total number of unique sources in the LOFAR radio catalogue is of the order 10^6. Long-term monitoring of the database tasks is essential to predict pipeline performance and understand the instrument as a whole. Queries with poor scaling (e.g. exponential) will eventually jam the processing. Significant increases of cadence and/or source density determine the critical limits of the system and permitted types of observations. Swinbank et al. (2015) show that the TraP run times increase linearly with input size within the LOFAR observation constraints.

Source lists produced by optical instruments are in general much larger, primarily due to the intrinsic higher resolution in combination with the increased sensitivity. Also the catalogues that represent the optical sky models hold orders of magnitude more sources than their radio counterparts. Therefore, one avoids naive implementations of the TraP for optical instruments, since the extrapolation of the source counts into the optical spectrum will most probably break linear performance or even in a best-case long-term linear performance scenario, the processing time will pass the cadence time at some point.

Table 1
Characteristics of the MeerLICHT telescope, a single BlackGEM prototype telescope. DB source data size is the storage size that all properties of single source would need when stored in a database. We assume observation nights of 10 h.

Mirror diameter	65 cm	
FoV	2.7 deg^2	
CCD size	10,536 × 10,536	
Resolution	0.57″/px	
bits per px	16	
Image size	222 MB	
Calib. images per night	2 × (10 bias + 5 × 5 flats) = 70	
DB source data size	402 B	
Observation mode	Nominal	Fast
Integration time	5 min	1 min
Sensitivity	23 mag	21 mag
Science images per night	120	600
Data rate per night	42 GB	148 GB

The planned wide-field optical telescope array BlackGEM is dedicated to measure optical emission from pairs of merging neutron stars and black holes (Bloemen et al., 2015). BlackGEM will start with 3 telescopes, all of which will be located at ESO La Silla, Chile. MeerLICHT,[1] a single BlackGEM telescope acting as a prototype, is coupled to the MeerKAT radio array (a precursor to SKA; Brederode et al., 2016) to operate in concert and allowing to study the optical–radio sky simultaneously as a true multi-wavelength instrument. Table 1 shows the characteristics of a single BlackGEM telescope. From a database–pipeline perspective, the most influential properties are the source data size, the source density and the integration time, where the latter determines the cadence.

In this paper we use the TraP-like queries as a baseline and investigate its scalability to the MeerLICHT environment. We need to know to what extent the processing of images with source densities of 500,000 deg^{-2} or source lists with hundreds of thousands of sources is still feasible.

The paper is outlined as follows. Section 2 gives the rationale behind the choice of a column-oriented relational database management system (RDBMS). Section 3 describes the experimental set up of the performance tests, the data and the TraP queries that were adjusted and optimised. The results are presented in Section 4 and concluded in Section 5. Although all source codes are publicly available, the Appendices A–C show the relevant query snippets for readability.

2. Rationale for MonetDB, a column-oriented relational database management system

Non-relational databases, e.g., key–value and NoSQL stores, lack solid support for many data manipulations that are required for this type of astronomical application. We cannot afford redundant or duplicate storage, and therefore have to distribute the data over a minimal set of related tables. The absence of fast join and cross-match functionality, schema free storage formats, no transactional support and own unique query languages makes non-relational databases hard to perform on such applications. Furthermore, the bulk processing requires fast data aggregation, ordering and indexing to avoid large table scans. All these functionalities are well-known and implemented in relational database systems.

Relational database storage models follow either the row-oriented or column-oriented principles. The row-oriented storage model partitions tabular data horizontally. Such record layouts consist of rows that store all their columns contiguously, with the adverse effect that a data block contains multiple data types.

[1] Check the current status at www.meerlicht.org.