

# Implementing the Palomar Transient Factory Real-Time Detection Pipeline in GLADE: Results and Observations

Florin Rusu<sup>1</sup>, Peter Nugent<sup>2</sup>, Kesheng Wu<sup>2</sup>  
frusu@ucmerced.edu, penugent@lbl.gov, kwu@lbl.gov

<sup>1</sup>University of California, Merced, CA 95343, USA

<sup>2</sup>Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

**Abstract.** Palomar Transient Factory is a comprehensive detection system for the identification and classification of transient astrophysical objects. The central piece in the identification pipeline is represented by an automated classifier that distinguishes between real and bogus objects with high accuracy. Given that the classifier has to identify the most significant transients out of a large number of candidates in near real-time, the response time it provides is of critical importance. In this paper, we present an experimental study that evaluates a novel implementation of the classifier in GLADE—a parallel data processing system that combines the efficiency of a database with the extensibility of Map-Reduce. We show how each stage in the classifier – candidate identification, pruning, and contextual realbogus – maps optimally into GLADE tasks by taking advantage of the unique features of the system—range-based data partitioning, columnar storage, multi-query execution, and in-database support for complex aggregate computation. The result is an efficient classifier implementation capable to process a new set of acquired images in a matter of minutes even on a low-end server. For comparison, an optimized PostgreSQL implementation of the classifier takes hours on the same machine.

**Keywords:** parallel databases, multi-query processing, real-time classification

## 1 Introduction

The Palomar Transient Factory (PTF) project [1, 2] aims to identify and automatically classify transient astrophysical objects such as variable stars and supernovae in real-time. As a secondary objective, a catalog containing the identified transients and other celestial objects is constructed for subsequent querying and analysis. PTF is a comprehensive transient detection system including a wide-field survey camera, an automated real-time data reduction pipeline, a dedicated photometric follow-up telescope, and a full archive of all detected sources (Figure 1 [2]). The computational system supporting the project consists of two separate processing pipelines [2] fed with the images taken by the camera. Between 2000 and 4000 high-resolution ( $2048 \times 4096$  pixels) images are taken each night and fed into the two pipelines through high-speed communication links. The total amount of raw data varies between 60 and 100 GB per night. The near-real-time transient detection pipeline [3] has the goal of identifying and classifying transient objects within 30-45 minutes of images being taken. Observation of potential

transients by a network of follow-up telescopes is triggered immediately after detection in order to confirm their existence. The main objective of the time-consuming archival pipeline [4] is to create a comprehensive catalog of high-quality images and celestial objects that can be queried using a variety of criteria. It is executed on the entire set of images acquired during one night in order to achieve high accuracy. The execution of the archival pipeline typically takes 4-5 hours [4], but it can extend to several days in some cases [2].

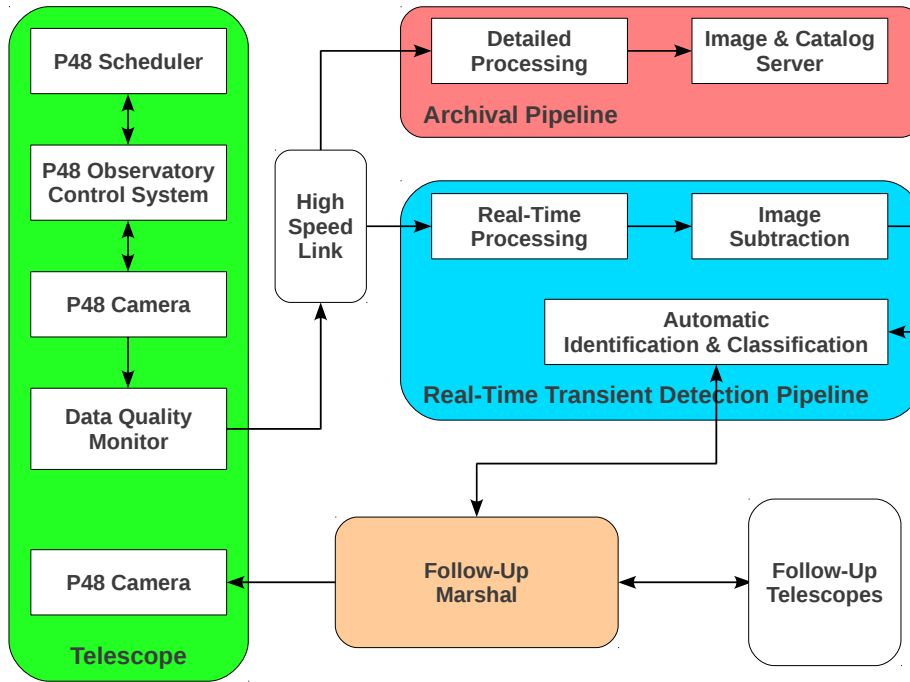


Fig. 1: PTF data flow [2].

**Problem formulation.** The problem we address in this paper is the identification of real transient candidates in the near-real-time detection pipeline. Specifically, we focus on the classification phase of the real-or-bogus classifier. The goal of this classifier is to identify real transients with high accuracy. The input consists of a set of candidates extracted during image subtraction and a trained random forest classifier. In the output, the candidates are given scores, i.e., the realbogus score, quantifying the probability of them being real. Only the candidates with realbogus scores higher than a threshold and satisfying a set of additional constraints are considered real.

**Contributions.** Our objective is to design and implement a real-time classifier capable to keep-up with the continuously increasing size of the PTF repository. Our motivation is the incapacity of the existing PostgreSQL solution to identify transient candidates accurately due to the larger data volumes it has to handle. We present a novel

implementation for the real-or-bogus classification in GLADE [5]—a parallel multi-query processing system targeted specifically at analytical workloads. We show how each stage in the classification process is natively supported in GLADE – this is not true for the existing PostgreSQL [6] solution – and prove with experimental results the effect on query execution performance. Since the GLADE implementation reduces the time to investigate a set of candidates to minutes – from hours in PostgreSQL – this allows for considerably more candidates to be thoroughly evaluated, thus increasing the likelihood to find many transients that are otherwise missed by the current solution.

**Outline.** Section 2 presents in detail automatic transient identification and real-or-bogus classification. It also discusses the PTF solution deployed in production and the problems it has. GLADE is introduced in Section 3. The implementation of the real-or-bogus classifier in GLADE and the results of the experimental evaluation are given in Section 4 which also contains a comparison with two PostgreSQL solutions. We conclude the paper in Section 5.

## 2 Automatic Transient Identification & Classification

In this section, we present the details of the *automatic detection and classification stage* in the real-time transient detection pipeline (Figure 1 [2]). The input to this stage is represented by the transient candidates extracted during the image subtraction stage. There are in the order of  $10^5$  such candidates extracted every 45 minutes. Two questions have to be answered for every candidate:

1. Is the candidate real?
2. If real, what is the transient type of the candidate?

Both these questions are answered using automated machine learning classification techniques, i.e., random forest classifiers [3] in this case, that require human intervention only in the follow-up stage. This is necessary considering the number of candidates – 1 to 1.5 million – extracted every night. Since the focus of this work is identifying real candidates – the first question – we present the details and the existing solution in the following. A description of the type classifier can be found elsewhere [3].

### 2.1 Real-or-Bogus Classification

Any machine learning method consists of multiple phases. First, a series of features have to be defined for the input data. These are used as parameters for the classifier. The features used by the real-or-bogus classifier are extracted during image subtraction and stored in the candidate database. There are 28 features used by the classifier. Second, a training dataset containing labeled examples is used to compute the parameters of the classifier—the training phase. The training dataset consists of 574 candidates manually labeled with the realbogus score by multiple human scanners. At last, the trained classifier is presented with unlabeled examples and the class has to be determined—the classification phase. The output consists of 5 classes, i.e., {*bogus*, *suspect*, *unclear*, *maybe*, *realish*}. The probability of a candidate being in each of these classes is returned by the random forest classifier—probabilistic model. The score, i.e., realbogus score, for a candidate is computed as a weighted average of

the class probabilities. The final realbogus score takes into account additional information, i.e., the scores of neighboring candidates from the same subtraction. Moreover, a high-scoring candidate is deemed real if and only if it appears in at least two subtractions within 6 days. All the candidates identified as real at the end of the classification phase – 30 to 150 out of  $10^5$  – are flagged for immediate follow-up and sent to the type identification classifier. We present the exact details of the online classification phase in the following since training is a one-time offline process.

---

**Algorithm 1** *Real-or-Bogus Classification*

---

**Input:** new subtraction set (*subtraction*) with corresponding candidate set (*candidate*) and their probabilities (*rb\_classifier*) computed by the random forest classifier

**Output:** a subset of real candidates (*real*)

```

1. real  $\leftarrow$  IdentifyCandidates(subtraction, candidate, rb_classifier)
2. for all r  $\in$  real do
3.   if SingleAppearance(r, subtraction, candidate, rb_classifier) then
4.     real  $\leftarrow$  real - r
5.   end if
6. end for
7. for all r  $\in$  real do
8.   ctxt_score  $\leftarrow$  CtxtRealBogus(r, subtraction, candidate, rb_classifier)
9.   if ctxt_score < threshold then
10.    real  $\leftarrow$  real - r
11.   end if
12. end for

```

---

**Identify candidates.** The initial realbogus score – the score returned by the random forest classifier – corresponding to a candidate is computed during image subtraction. It is stored together with other candidate data in the subtraction and candidate database. Candidate identification requires a simple query that returns all the candidates with high realbogus score extracted from subtractions computed during a specified time interval.

**Prune single appearance candidates.** In order to increase the probability that a high-scoring candidate is indeed real, a candidate identified in the first query has to satisfy an additional condition. The candidate has to appear at a close spatial position in other subtractions close in time to its originating subtraction. Independent of its original realbogus score, a single appearance candidate is pruned away. Since the area and time interval are dependent on the candidate, pruning requires a separate query with different space and time bounds for every candidate. The larger the number of candidates identified in the first query, the more queries have to be executed for pruning.

**Compute contextual realbogus score.** The final realbogus score of a candidate takes into consideration the score returned by the random forest classifier for the closest  $k$  other candidates extracted from the same subtraction. This is called the contextual realbogus score. It is based on this score that the final classification decision is made. Computing the contextual realbogus score is a rather complicated process that involves a nearest-neighbor query followed by a complex aggregate computation. Unlike pruning, the nearest-neighbors are computed only along the spatial dimension, i.e., in the

same subtraction. Nonetheless, a separate query has to be executed for every candidate that survives pruning—a considerable problem when the number of candidates is large. Algorithm 1 summarizes formally the stages of the real-or-bogus classification.

## 2.2 Existing Solution

The current solution implemented in the PTF pipeline is a standard Python [7] application with a PostgreSQL [6] database backend. The database contains 3 tables – `subtraction`, `candidate`, and `rb_classifier` – storing subtractions, candidates, and the scores returned by the random forest classifier. `subtraction` and `candidate` are wide tables having 51 and 46 columns, respectively. Many of the columns are never used in queries. The number of rows in these tables increases continuously as more observations are taken daily. `candidate` and `rb_classifier` already contain a few billion tuples each. Data corresponding to a new set of images are added to the tables during the image subtraction stage. Transient identification and classification execute as database queries. Possible candidates in a given time window are identified with a complex query over the 3 tables. For each such candidate, a time and space nearest-neighbor query is executed from the application to find additional appearances of the candidate. These queries are executed iteratively. For the remaining candidates, another spatial nearest-neighbor query is executed to find close candidates in the same subtraction. The contextual realbogus score is computed in the application by combining the score of the candidate with those of its neighbors. This is another iterative process that goes over the non-pruned candidates one at a time.

**Problems with the existing solution.** The existing approach suffers from a series of problems that make real-time candidate identification and classification in the limited time interval between two subtraction sets – approximately 45 minutes [3] – difficult. Experimental results over a relatively small snapshot of the database from the early stages of the project confirm this problem (Section 5). As the size of the repository increases with the acquisition of new images, the situation will become only worse. As a result, many of the transients are missed simply because there is not enough time to carefully investigate them.

The fundamental limitation is raised by the need to evaluate each candidate sequentially even though the same query template is used across all the candidates. Essentially, two passes over the candidates extracted based on the realbogus score are required to take a decision. And in each pass, a complicated nearest-neighbor query is executed for each candidate. Since the time taken to process one query over the increasingly larger `candidate` table grows continuously with the size of the table, the number of candidates that can be inspected between two subtractions decreases.

The PostgreSQL row-based storage format affects query performance negatively considering the width of `subtraction` and `candidate` tables and the number of attributes used in the query—a small fraction out of the overall number of attributes. While query execution speed can be improved with appropriate indexes, this results in data ingestion time increase due to index maintenance, thus limiting the time available for querying. Essentially, indexing moves the bottleneck from querying to data ingestion. With the increase in repository size, index maintenance under data ingestion only becomes worse.

Data transfer between the database and the application is another limitation that is a direct consequence of the large number of queries that have to be executed. The reason for this is the lack of support for complex computations inside the database. While user-defined functions (UDF) and user-defined aggregates (UDA) provide extensibility to in-database complex computations, the exclusive SQL invocation limits their applicability. As a result, these complex computations are executed in the application in the current PTF solution.

### 3 GLADE

Given the aforementioned problems of the existing solution and the incapacity to accurately identify some highly-probable candidates, novel solutions have to be explored. The approach we take in this paper is a novel implementation of the real-or-bogus classifier in GLADE—a parallel data processing prototype we have developed from scratch over the past few years. Although we showed the considerable performance gains GLADE provides over PostgreSQL and Hadoop on a limited set of tasks [5], we have not evaluated GLADE on a complex real-life application yet. Moreover, the characteristics of the PTF real-or-bogus classifier map perfectly on the GLADE architectural features. While these are compelling reasons to carry out such an investigation, the experimental results in Section 5 prove that GLADE is indeed a suitable solution that outperforms the existing PostgreSQL implementation.

**GLADE.** GLADE [5, 8] is a parallel data processing system specifically designed for the execution of analytical tasks specified in SQL enhanced with Generalized Linear Aggregates (GLA). This allows for the execution of a much larger class of analytical computations beyond the standard SQL aggregates. Essentially, GLADE provides an infrastructure abstraction for parallel processing that decouples the algorithm from the runtime execution. The algorithm has to be specified in terms of a clean interface – SQL + GLA – while the runtime takes care of all the execution details including data management, memory management, and scheduling. Contrary to existent parallel data processing systems designed for a target architecture, typically shared-nothing, GLADE is architecture-independent. It runs optimally both on shared-disk servers as well as on shared-nothing clusters. The reason for this is the exclusive use of thread-level parallelism inside a processing node while process-level parallelism is used only across nodes. There is no difference between these two in the GLADE infrastructure.

**Architecture.** GLADE consists of two types of entities—a coordinator and one or more executor processes (Figure 2). The coordinator is the interface between the user and the system. Since it does not manage any data except the catalog metadata, the coordinator does not execute any data processing task. These are the responsibility of the executors, typically one for each physical processing node. It is important to notice that the executors act as completely independent entities, in charge of their data and of the physical resources. Each executor runs an instance of the DataPath [9] relational execution engine enhanced with a GLA metaoperator for the execution of arbitrary user code specified using the GLA interface.

*Communication Manager* is in charge of transmitting data across process boundaries, between the coordinator and the executors, and between individual executors. Different

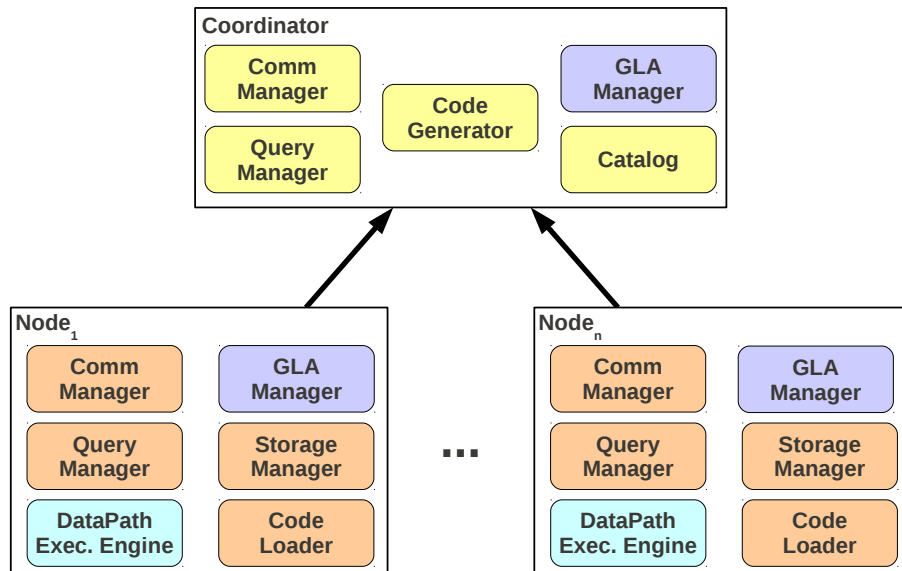


Fig. 2: GLADE architecture.

inter-process communication strategies are used in a centralized environment with the coordinator and the executor residing on the same physical node and for a distributed shared-nothing system. The communication manager at the coordinator is also responsible for maintaining the list of active executors. This is realized through a heartbeat mechanism in which the executors send alive messages at fixed time intervals.

*Query Manager* is responsible for admission, setup, and query processing coordination across executors and queries. This is a particularly important task since processing is asynchronous both with respect to executors as well as to queries.

*Code Generator* fills pre-defined M4 templates with macros specific to the actual processing requested by the user generating highly-efficient C++ code similar to direct hard-coding of the processing for the current data. The resulting C++ code is subsequently compiled together with the system code into a dynamic library. This mechanism allows for the execution of arbitrary user code inside the execution engine through direct invocation of the GLA interface methods.

*Code Loader* links the dynamic library to the core of the system allowing the execution engine and the GLA manager to directly invoke user-defined methods. While having the code generator at the coordinator is suitable for homogeneous systems, in the case of heterogeneous systems both the code generator and the code loader can reside at the executors.

*DataPath Execution Engine* implements a series of relational operators – SELECT, PROJECT, JOIN, AGGREGATE – and a special GLA metaoperator for the execution of arbitrary user code specified using the GLA interface. They are all configured at runtime with the actual code to execute based on the requested processing. The execution engine has two main tasks—manage the thread pool of available processing resources

and route data chunks generated by the storage manager to the operators in the query execution plan. Parallelism is obtained by processing multiple data partitions – chunks – simultaneously and by pipelining data from one operator to another.

*GLA Manager* executes `Merge` at executors and `Terminate` at coordinator, respectively. These functions from the GLA interface [5] are dynamically configured with the code to be executed at runtime based on the actual processing requested by the user. Notice that the GLA manager merges only GLAs from different executors, with the local GLAs being merged inside the execution engine.

*Catalog* maintains metadata on all the objects in the system such as table names and attribute names and their partitioning scheme. These data are used during code generation, query optimization, and execution scheduling. In addition to the global catalog, each executor has a local catalog with metadata on how its corresponding data partition is organized on disk.

*Storage Manager* is responsible for organizing data on disk, reading, and delivering the data to the execution engine for processing. The storage manager operates as an independent component that reads data asynchronously from disk and pushes it for processing. It is the storage manager rather than the execution engine in control of the processing through the speed at which data are read from disk. In order to support a highly-parallel execution engine consisting of multiple execution threads, the storage manager itself uses parallelism for simultaneously reading multiple data partitions.

**Range-based data partitioning.** Parallel execution is supported in GLADE through data partitioning, i.e., multiple partitions are processed simultaneously by different executors. The tuples of a relation are divided horizontally into chunks containing thousands to a few million tuples. Chunks are stored continuously on disk. The larger the size of the chunk, the longer the size of sequential scans, thus the smaller the number of disk seeks. While tuples can be assigned to chunks in arbitrary order, it is particularly useful for many workloads to have tuples with close values along some attributes grouped together in the same chunk. This corresponds to range-based data partitioning. Generating range-partitioned chunks is typically more complicated since data have to be partially ordered along the partitioning attributes. The benefit is faster execution for range queries since a reduced number of chunks have to be processed.

**Column-oriented storage.** Inside a chunk the columns of a relation are further partitioned vertically, with a disk page storing only values from the same column. Attribute values corresponding to the same tuple are stored at the same relative position inside each column. This allows for immediate tuple reconstruction in memory. The benefit of column-oriented storage is evident in the case of wide relations containing a large number of attributes with only a few of them accessed by every query. When range-based partitioning is combined with columnar storage – the case in GLADE – the amount of data read from disk is minimized since only the chunks and the columns required by the query are retrieved.

**Multi-query processing.** GLADE supports concurrent execution of multiple queries by sharing data access across the entire hierarchy—from disk to CPU registers through memory and cache. All the queries reading data from the same relation are connected to a single circular shared scan operator that reads a chunk only once and distributes it to all the queries that require it. While this is standard practice in any multi-query process-



ing system, chunk sharing in GLADE is taken considerably further. Essentially, chunks are shared across all the common operators in the query execution trees corresponding to two queries. This requires merging separate operators with similar functionality into a single mega-operator that combines the operations corresponding to each query. For example, instead of having two selection operators with different predicates – one for each query – a single operator containing both predicates is created in GLADE. The new combined operator is responsible for identifying what queries a chunk is valid for and for setting the correct tuple validity based on the query predicates. The same logic can be applied to other relational operators, including JOIN, GROUP BY, AGGREGATE, and the GLA metaoperator. The code executed by each operator is dynamically generated at runtime based on the running queries.

**Complex aggregates.** The GLA metaoperator supports the execution of arbitrary user code specified using the abstract GLA interface [10]. This allows for the execution of complex computations far beyond standard SQL aggregates inside the execution engine without the need to extract data into an application with more powerful computational capabilities. This paradigm shift – bring the code near the data instead of moving data to the code – results in considerable gains especially in the cases where a large amount of data have to be moved.

## 4 Experimental Evaluation

In this section, we present the GLADE implementation for real-or-bogus classification. We show how data are mapped onto the GLADE storage model, how each task in the classification is expressed as GLADE computations, and how native GLADE features – range-based data partitioning, columnar storage, multi-query processing, and complex aggregate computation – are used in this workload. We provide measurement results that prove a significant improvement over the existing solution and we analyze the reasons for this.

**Data.** The data we use in the experiments are a snapshot of the subtraction and candidate database. The 3 tables referenced in real-or-bogus classification and their characteristics are given in Table 1. The overall size of the 3 tables when loaded in GLADE is 161 GB. Notice that approximately 5,000 candidates are not classified by the random forest classifier—`rb_classifier` contains less tuples than `candidate`. There are 647 candidates per subtraction on average.

The maximum chunk size is fixed at  $2^{20} \approx 1$  million tuples across all the tables. This generates a single `subtraction` chunk and 642 chunks for the other two tables. The size of a full chunk is different though across tables since they contain a different number of columns. Notice that only the columns required in query processing are read for a chunk – not the entire chunk – due to the columnar storage. Thus, even the `subtraction` table is never read in full unless all the columns are requested by the query. `candidate` and `rb_classifier` are range-based partitioned along the `subtraction_id` attribute. This guarantees that all the candidates extracted from the same subtraction end up in the same chunk. Moreover, candidates from subtractions close in time are also co-located in the same chunk with high probability. This partitioning has two benefits. It minimizes the number of chunks read from disk. And it isolates

processing to the chunk level, thus increasing the amount of parallelism achieved by processing multiple chunks simultaneously.

Table name	Columns	Rows	Chunks
subtraction	51	1,039,758	1
candidate	46	672,912,156	642
rb_classifier	9	672,906,737	642

Table 1: Tables used in real-or-bogus classification.

**Setup.** The machine used in the experiments is a low-end server with an Intel Core2 Quad CPU running at 2.66 GHz, 4 GB of memory, and a single 1 TB disk with sequential I/O throughput of 100 MB/s. Ubuntu SMP 10.10 64-bit is the operating system. There is a single GLADE executor in this configuration. It is co-located with the coordinator. Only thread-level parallelism is employed. The DataPath execution engine is configured to use 4 worker threads – one for each core – while the storage manager corresponding to every table assembles 4 chunks simultaneously. The reader might be surprised by our modest system choice given that the PTF pipeline is running in production on a powerful NERSC supercomputer. Nonetheless, our results confirm that even on such a low-end machine GLADE manages to load and classify the candidates in a set of subtractions in less than 20 minutes.

**Data ingestion.** The time it takes to ingest the entire dataset depicted in Table 1 in GLADE is 8,222 seconds ( $\approx$  2 hours 15 minutes) out of which 7,056 are spent loading the `candidate` table. This dataset corresponds though to many nights of observations. To determine how long it takes to ingest a set of subtractions generated at one instance in time, we chose a random night in the dataset, i.e., the night of October 10, 2011, compute its corresponding statistics, and then extrapolate the loading time based on these statistics. There are 2,997 subtractions taken during this night and 1,939,059 candidates at an average rate of 647 candidates per subtraction. The time taken to ingest these data into GLADE is only 24 seconds. Considering that the ingestion is distributed over 10 periods of 45 minutes each, the average ingestion time for a set of subtractions is less than 3 seconds.

**Candidate identification.** The first stage in real-or-bogus classification is to identify candidates with high realbogus score assigned by the random forest classifier. The corresponding query (1) contains a 3-way join between `subtraction`, `candidate`, and `rb_classifier` and a series of selection predicates on each of the tables. The most important predicate is a range selection on `subtraction` that limits the search to the images acquired most recently. Due to range-based partitioning and columnar storage, in the GLADE implementation this query reads only the chunks and columns that generate results. In the optimal situation, a single chunk is processed from each of the 3 tables. Out of the almost 2 million candidates detected during the night of October 10, 2011, only 40,087 are classified as real by the random forest classifier. This is only

2%. The number can be easily increased by relaxing the conditions in the query. It takes GLADE only 9.2 seconds to find the real candidates, i.e., 0.92 seconds per subtraction.

```

SELECT s.ujd, c.sub_id, c.id, c.ra, c.dec,
       c.xint_new, c.yint_new, c.pos_sub
FROM
  subtraction s JOIN candidate c ON (c.sub_id = s.id)
  JOIN rb_classifier rbc ON
    (rbc.sub_id = c.sub_id AND rbc.candidate_id = c.id) (1)
WHERE s.ujd > 2455844 AND s.ujd < 2455845 AND
      rbc.realbogus > 0.17 AND rbc.bogus < 0.35 AND
      c.b_image > 0.7 AND c.pos_sub = 'True' AND
      (c.a_image < 3.0 OR c.mag < 15.0)

```

**Candidate pruning.** Each of the candidates identified by the random forest classifier is further checked before deemed real. The first condition a candidate has to pass is that it appears in more than one subtraction at a position close to the original position where it was first spotted. This is expressed as a complex nearest-neighbor query (2) along the space and time attributes. In the current implementation of the PTF pipeline, one such query is executed sequentially for every candidate. This is 40,087 queries for our example night or approximately 4,000 queries for every subtraction set. Since all these queries have to be executed in less than 45 minutes, this step is by far the bottleneck of the entire process. GLADE multi-query processing kicks in perfectly in this situation thus allowing for multiple candidates to be checked at the same time. Most importantly though, the time to check many candidates – up to 64 candidates in the current GLADE implementation – is the same as checking a single candidate. The reason for this is that the queries have identical execution plans – more or less some constants – which allows for maximum data access sharing. When coupled with range-based partitioning and columnar storage, it takes only 18 minutes to check the 4,000 candidates identified in a subtraction set—18 seconds for a group of 64 candidates. 560 candidates survive pruning on average for each subtraction set.

```

SELECT COUNT(*)
FROM
  subtraction s JOIN candidate c ON (c.sub_id = s.id)
  JOIN rb_classifier rbc ON (rbc.candidate_id = c.id)
WHERE c.ra BETWEEN (%1f,%2f) AND (2)
      c.dec BETWEEN (%3f,%4f) AND
      (s.ujd BETWEEN (%5f,%6f) OR s.ujd BETWEEN (%7f,%8f))
      AND (rbc.realbogus > 0.07 OR c.pos_sub <> 'True') AND
      c.b_image > 0.7 AND (c.a_image < 3.0 OR c.mag < 15.0)

```

**Contextual realbogus computation.** For the surviving candidates, the contextual realbogus score is computed based on the probability of being real of their nearest-neighbor candidates in the subtraction. This requires another iterative process in which

each surviving candidate is examined independently. The difference from pruning is that the contextual realbogus score cannot be computed inside the database. Instead it is computed in a Python script that extracts the necessary data from the database using query (3). This is not required in GLADE though since complex aggregates can be expressed as GLAs and executed inside the system without moving data between processes. In addition to the savings in execution time, the GLA mechanism allows for all the computation to be confined to the database engine—a cleaner and easier to understand solution. In GLADE, the contextual realbogus score for the 560 candidates surviving pruning in a subtraction set is computed in 88 seconds—it takes 10 seconds on average to compute the score for a group of 64 candidates.

Table 2 summarizes the results we obtained for processing the October 10, 2011 data in GLADE. These are average results for processing a subtraction set. The overall time to classify the candidates is less than 20 minutes. This is less than half the length of the interval between two sets of images are ingested, i.e., 45 minutes. The remaining time can be used either to increase the rate at which images are ingested or to analyze more candidates—some of the conditions based on which candidates are pruned are arbitrary and they are targeted at reducing the overall classification time. This is not a problem in the GLADE implementation though.

```

SELECT c.id,
        $\sqrt{(c.xint\_new-\%1f)^2 + (c.yint\_new-\%2f)^2}$  AS dist
FROM
  subtraction s JOIN candidate c ON (c.sub_id = s.id)
  JOIN rb_classifier rbc ON
    (rbc.sub_id = c.sub_id AND rbc.candidate_id = c.id)
WHERE s.id = %3i AND c.pos_sub = '%4s'
ORDER BY dist

```

**PostgreSQL solutions.** In order to compare the proposed GLADE approach with the existent solution, we devise two alternative PostgreSQL databases. The first database does not contain any optimizations. There are no indexes or any other structures for enhancing query performance. The second database defines indexes for all the attributes used in selection predicates or join conditions across the three workload queries. This is the solution implemented in the PTF production pipeline. We deploy these two databases on a PostgreSQL 8.4 server running on the same test machine. We modify the server configuration in order to maximize usage of the available memory resources in the system, e.g., we set `shared_buffers` to 3 GB.

Table 2 contains the results for the two PostgreSQL database implementations. The indexed implementation outperforms the non-indexed version considerably at query processing. The gap is as much as 6 orders of magnitude for the contextual realbogus computation. The reason for this is that the non-indexed database has to read all data from all the tables in order to perform any query. Since no indexes are available, sequential scan is the only feasible path access strategy. Indexes reduce dramatically query execution time since the tuples satisfying the highly-selective predicates can be identified with as little as a single disk access. Due to the large buffer pool, disk access

is not even required at all in many situations. Indexes also play an important role in the selection of the join algorithms used in query execution plans. The situation is radically different though for data ingestion. While it takes less than a minute to load a new set of candidates in the non-indexed database, it takes 68 minutes to do so in the indexed version. This 61 factor difference is entirely due to index maintenance. Adding tuples to the `candidate` table requires insertions in each of the 9 indexes defined over it. Although this might not seem such a difficult problem at first, in the case of a batch of 200,000 insertions the probability to encounter some time-consuming index reorganizations is quite high.

Overall, none of the PostgreSQL solutions meets the requirement to ingest and identify a new set of candidates in less than 45 minutes. The index-based solution deployed in production takes 75 minutes on our test machine—out of which 60 minutes are spent for data ingestion. The non-indexed version is far from this requirement. A possible solution to decrease the loading time for the indexed database is to reduce the number of indexes. The expectation is that the decrease in loading time offsets the increase in query execution time and for some combination of indexes the overall time drops below 45 minutes. Finding the optimal index combination is a hard problem that requires the investigation of an exponential number of alternatives.

Phase	GLADE	PostgreSQL	PostgreSQL + indexes
Data ingestion	3 sec	59 sec	1 hour 8 sec
Identification	0.92 sec	45 sec	4.67 sec
Pruning	18 min	607 hours	15 min 39 sec
Contextual realbogus	1 min 30 sec	68 hours	0.79 sec

Table 2: Average results for processing a subtraction set on October 10, 2011.

**Observations.** When we compare the proposed GLADE solution to the PostgreSQL indexed database, we remark some interesting aspects. Overall, GLADE outperforms PostgreSQL by a factor of 3.88. This is entirely due to the efficient GLADE data loading mechanism which is faster by 3 orders of magnitude. Since GLADE does not employ any secondary data structures to enhance query performance, it is not as efficient as indexed PostgreSQL in answering queries. The difference between the two systems – only 24% – is considerably smaller when compared to the basic PostgreSQL implementation. The GLADE architecture specifically targeted at analytical processing and optimized for read-mostly workloads is responsible for providing similar query performance to indexed PostgreSQL but without the associated increase in database size – the indexed PostgreSQL database is twice as large as GLADE – and ingestion time—GLADE loads new candidates a factor of 60 faster. Range-based partitioning and columnar storage combine together in order to minimize the amount of data read from disk across all types of range queries. Dedicated support for the execution of any user code inside the system eliminates data movement almost entirely and allows for complex computations to be executed right near the data. For the PTF workload though,

the most significant gains are due to multi-query processing. Instead of verifying each candidate one at a time, GLADE allows for up to 64 candidates to be evaluated simultaneously in the same amount of time. This is because all the queries have identical execution plans and GLADE is capable to combine them into a single plan in which the operators share access along the entire data path—from disk to CPU registers through main memory and cache.

## 5 Conclusions

In this paper, we present a novel implementation for the real-or-bogus classification in GLADE—a parallel multi-query processing system targeted specifically at analytical workloads. We show how each stage in the classifier – candidate identification, pruning, and contextual realbogus – maps optimally into GLADE tasks by taking advantage of the unique features of the system—range-based data partitioning, columnar storage, multi-query execution, and in-database support for complex aggregate computation. The result is an efficient classifier implementation capable to process a new set of acquired images in a matter of minutes even on a low-end server. For comparison, the existing optimized PostgreSQL implementation of the classifier is a factor of 3.88 slower. Due to this reduction in the time to investigate a set of new candidates, considerably more candidates can be thoroughly evaluated, thus increasing the likelihood to find many transients that are otherwise missed by the current solution.

## References

1. Palomar Transient Factory. [www.astro.caltech.edu/ptf/](http://www.astro.caltech.edu/ptf/) [November 2013].
2. Law, N.M., al.: The Palomar Transient Factory: System Overview, Performance and First Results. CoRR **abs/0906.5350** (2009)
3. Bloom, J.S., al.: Automating Discovery and Classification of Transients and Variable Stars in the Synoptic Survey Era. CoRR **abs/1106.5491** (2011)
4. Grillmair, C.J., al.: An Overview of the Palomar Transient Factory Pipeline and Archive at the Infrared Processing and Analysis Center. In: ASP Conf. Ser. 434, Astronomical Data Analysis Software and Systems XIX. (2010) 28–36
5. Cheng, Y., Qin, C., Rusu, F.: GLADE: Big Data Analytics Made Easy. In: Proceedings of 2012 ACM SIGMOD International Conference on Management of Data. (2012) 697–700
6. PostgreSQL: <http://www.postgresql.org/> [November 2013].
7. Python Programming Language. <http://www.python.org/> [November 2013].
8. Cheng, Y., Rusu, F.: Astronomical Data Processing in EXTASCID. In: Proceedings of 2013 SSDBM Conf. on Sci. and Stat. Database Management. (2013) 387–390
9. Arumugam, S., Dobra, A., Jermaine, C., Pansare, N., Perez, L.: The DataPath System: A Data-Centric Analytic Processing Engine for Large Data Warehouses. In: Proceedings of 2010 ACM SIGMOD International Conference on Management of Data. (2010) 519–530
10. Rusu, F., Dobra, A.: GLADE: A Scalable Framework for Efficient Analytics. Operating Systems Review **46**(1) (2012) 12–18