

Incremental Read-Aheads

A. Soydan Bilgin*

Department of Computer Science,
North Carolina State University,
Raleigh, NC 27695, USA
asbilgin@unity.ncsu.edu

Abstract. In spite of the advances in caching, query optimization, and object persistence techniques in the past few years, the cost of interactions of large-scale data-intensive applications with a relational database where the persistent objects are implemented remains a performance bottleneck. To reduce the cost of such interactions, we present a read-ahead scheme, which allows the application to reduce the number of database roundtrips by retrieving the data before it is actually needed by the transactions in the applications.

We focus on designing generic rules for determining the efficient sequences of SQL statements for read-ahead queries on relational databases, such that the rules would be useful across application domains and data-access patterns. This paper explains our research methodology for generating generic access patterns and studying the parameters that influence the costs of various combinations of read-ahead SQL statements that implement the generic access patterns of applications.

1 Introduction

The objective of this project is to develop new methods for improving a range of performance metrics in modern relational databases. The main direction of the project is to design efficient scalable techniques for increasing the throughput of database accesses and minimizing the cost of database interactions, by reading ahead of time the data before it is actually needed by the transactions in the data-intensive applications.

Relational databases are the most popular and commonly used DBMSs in the enterprise, so minimizing the cost of interactions between large-scale data-intensive applications and relational databases will result in considerable gains in performance of applications. Each data access operation results in one or several relatively expensive physical disk accesses and network overhead that reduce the total throughput in terms of business transactions performed in a unit of time.

Traditional caching and prefetching techniques are heavily used to efficiently handle database queries, which process large number of data objects or Web documents [1, 8]. The technique of prefetching refers to the process of guessing an application's future requests for data and getting those data into the cache before they are actually referenced. Caching is used to store the actively referenced data objects, thereby avoids unnecessary requests to the database. However, even with caching, if an access to a non-cached data

* Doctoral student under the guidance of Rada Chirkova and Munindar P. Singh.

item (e.g., uncacheable data, compulsory misses) occurs that entails a round-trip to the database, performance will suffer. [2] provides a feeling for the performance penalty of relational databases with and without the technique of prefetching. According to the experiments described in [2], the retrieval time of data is up to 5.5 times faster to get rows in a batch of 100 rows than a row at a time. Simple prefetch mechanisms, which create read-ahead threads for certain queries that return large quantities of data sequentially from a single table, are already used in commercial data servers or application servers.

The prefetching technique of reading ahead of time, which we refer to as read-aheads in the rest of the paper, is used to reduce the number of database roundtrips either by augmenting the current query to include the answer to the next queries or to generate additional queries that are likely to come next. In current application servers, some read-ahead techniques are already used by object managers or containers. However, these techniques should be used systematically, because incorrect guessing reduces the throughput of the server, and the overhead examining the query for possible concatenations increases the latency. Current object managers that are responsible for accessing the database cannot efficiently benefit from the application's context descriptions. In this respect, we want to help such object managers to maximize the probability of correct guessing and to have efficient database interactions by finding the rules of thumb regarding what strategies to use to determine efficient sequences of SQL statements. We want to provide a systematic vision for state-of-the-art read-ahead techniques.

2 Motivation and Problem Formulation

Our goal is to bring the potentially accessible data to application's environment in the most efficient way in terms of maximal throughput and minimal database interaction cost. We use relational structure of data such as table relationships and data dependencies, and access patterns of the applications to predict the potentially accessible data that has the lowest retrieval cost. We separate this original problem into two parts.

1. Design generic rules for determining the efficient sequences of SQL statements for read-ahead queries on relational databases, such that the rules would be useful across application domains and data-access patterns.
2. Develop efficient and scalable algorithms for fine-tuning the read-ahead access rules in each particular application.

We seek rules of thumb that would be applicable across application domains and data-access patterns. To derive such *generic domain-independent rules*, we need to study the parameters that influence the costs of various combinations of read-ahead SQL statements that implement the generic access patterns. These generic domain-independent rules are then fine-tuned according to the data-access patterns of particular applications.

In the first phase of our project, we look for break-even points for efficient sequences of SQL statements by discovering new parameters and by exploring known parameters such as the size of data in the tables (number of rows and columns), the complexity of SQL calls (e.g., number of joins, presence of subqueries), presence of indexed columns, shape of the graph we are trying to retrieve (in a depth-first manner or breadth-first manner),

and network latency (in later stages of the project). Break-even points correspond to parameter values that are used in constructing efficient sequences of SQL statements, which give maximal throughput with minimal database interaction cost including data-retrieval time. As an outcome of this phase, we will have rules of thumb applicable across applications regarding what strategies to use for efficient sequences of SQL statements. At present, we don't consider the caching aspects of the problem to answer subsequent queries from the results that were prefetched. In addition to this, we don't consider read-aheads for batch queries. For example, instead of focusing on queries such as *find the name of all customers*, we focus on queries such as *find the name of the customer with given ID*.

Example 2.1 We consider a simple financial services data model for our examples [10]. According to this model, there can be many banks, and each customer can have multiple accounts in different banks. Relations *Person* and *Organization* store detailed information about customers of the banks. The relation *Customer* stores the common attributes of persons and organizations. The relation *Account* stores account information for customers in various banks. In the first and second transactions, we consider a customer representative who issues three and two data-access requests, respectively for the same customer from the database (The data-access requests are abbreviated with uppercase letters).

Transaction 1:

- A** Find the city and SSN of (person) customer '0011111'.
- B** Find the total amount of money that customer '0011111' has in all his accounts.
- C** Find the routing numbers of the banks where customer '0011111' has accounts.

Transaction 2:

- D** Find the name of customer '0011111'.
- C** Find the routing numbers of the banks where customer '0011111' has accounts.

For Transaction 1, request **A** requires two joins on tables for Customer, Person and Address without any read-aheads as in Figure 1. With read-aheads, it will be a smart choice to also bring data from Account table as in Figure 2, because request **B** that comes after **A**, requires data in the Account table. So using just one SQL statement with three joins as in Figure 2, we answer the first two requests **A** and **B**. For request **C**, we need to access the Bank_Customers table, so we need another SQL statement to fetch the data. As a result, with two SQL statements as in Figure 2 (three joins for the first statement and no joins for the second statement), we are able to answer the requests in Transaction 1. As another option, we can bring all the graph data in just one SQL statement by joining all the required tables at once; but in this case we will have a more complex query shown in Figure 3. This example shows the trade off between some of the parameters, such as the cost of the join operation versus the number of database roundtrips. Although the SQL statement in Figure 3 requires one database roundtrip, we may choose to use two SQL statements as in Figure 2 because these statements need fewer join operations in total and thus may result in lower total roundtrip time for large databases.

```
select Address.city, Person.ssn
from Customer, Person, Address
where Person.customerId=Customer.customerId
and Customer.addressId=Address.addressId
and Customer.customerId= '0011111'
```

Fig. 1. SQL statement for Transaction 1: A

AB:

```
select Person.ssn, Address.city, sum(amount)
from Customer, Account, Person, Address
where Person.customerId=Customer.customerId
and Customer.addressId=Address.addressId
and Customer.customerId=Account.customerId
and Customer.customerId= '0011111'
group by Person.ssn, Address.city
```

C:

```
select routingNumber
from Bank_Customers
where customerId='0011111'
```

Fig. 2. SQL statements for Transaction 1: AB and C

```
select Address.city, Person.ssn, Bank_Customers.routingNumber, sum(amount)
from Customer, Account, Person, Address, Bank_Customers
where Person.customerId=Customer.customerId and Customer.addressId=Address.addressId
and Customer.customerId=Account.customerId and Customer.customerId= '0011111'
and Customer.customerId=Bank_Customers.customerId
group by Person.ssn, Address.city, Bank_Customers .routingNumber
```

Fig. 3. SQL statement for Transaction 1: ABC

As the above example illustrates, our goal is not to find the minimal number of data-access statements, but to find the efficient number of simplest data-access statements. For example, one very complex data-access statement can bring all the data by reading ahead of time at one roundtrip, but this statement may not result in the maximal throughput due to the cost of the join operations. Also this data-access statement can bring the data that may never be needed by the application.

In subsequent phases of the project, we will explore the effect of object-to-relational mapping technique on our generic domain-independent rules and will develop and test learning algorithms to increase the efficiency of our generic rules for the data-access patterns of particular applications.

3 Related Work

In [11], various types of prefetching are characterized according to its short-term and long-term benefits. In short-term prefetching, future accesses to data are predicted according to the cache's recent access history. In long-term prefetching, global object access patterns are used to identify valuable objects that are worth prefetching (i.e., if an object is accessed by one client, it is likely that it will be accessed by other clients) [11].

Both prefetching techniques are mainly used in reducing the latency used in loading web pages [4]. In our work, we use long-term prefetching that also uses data-access costs to identify cheap and valuable prefetching sequences.

Prefetching is used to either pre-load the data needed for the subsequent queries for the given workload or load a specific collection of objects related to the requested object. The former case requires a more complex mechanism to track the cache contents and an analysis of which parts of the previous query is contained in the subsequent query, is required. [5] addresses the former case as an optimization for computing overlapping queries that generate Web pages (e.g., online shopping, where users narrow down their search space as they navigate through a sequence of pages). Predicate-based caching also serves the same idea where current cache content is used to answer future queries [7]. The approach in [12] uses the transition probabilities for each query to find the most probable query that can appear after the current query. So while executing the current query, they also execute the most probable subsequent query by using probable parameters and query pattern. In our approach, we take into account the cost of the data-access statements and parameters that affect this cost, to find the efficient sequence SQL statements.

[9] proposes to use a predictive cache to recognize and exploit access patterns for applications by incorporating prefetching mechanism with cache replacement mechanism to eliminate erroneous or least-likely prefetches. [6] also aims to ‘pre-cache’ the objects that are likely to be subsequently accessed by the application. Haas et al. propose a heuristic approach to cache and prefetch the objects whose object identifiers were returned as part of the query, so they make prefetching decisions according to the object identifiers found in the result set of the query. However, instead of focusing on finding the most beneficial prefetches, their goal is to find the cost of caching the prefetched objects. They incorporate the prefetching process into query processing to find the best execution plan by considering the cost of caching the prefetched tuples.

[2] specifically addresses the prefetching technique on relational databases where persistent objects are implemented. They use the context of an object as a predictor for future accesses in navigational applications. This context describes the structure in which the object was fetched. Main prefetching methods are listed (e.g., prefetching all the attributes of the requested object(s)). The results are applicable across application domains, because they use generic access patterns that are applicable across a wide range of applications. However, they only make one-level prefetching for referenced objects, so they don’t actually answer the ‘how deep’ question for read-aheads. Their overall goal is to minimize database latency for future data-access statements, so they don’t explore the cost of efficient sequence of data-access statements.

We explore the cost of data-access statements to find efficient generic data-access patterns. None of the previous work consider the prefetching problem both with query optimization parameters and navigational access patterns, such as following a relationship, at the same time. Also by providing a mechanism for merging simple data-access statements to find an efficient sequence of data-access statement, we actually use a different aspect for multi-query optimization [3] where dependencies or common subexpressions between the queries in a sequence are explored and computed.

4 Proposed Approach

Applications use objects, but these objects are mapped to tuples of the appropriate tables. In relational databases, the objects accessed by the applications are always associated with each other. Most of the time, the data in the database is accessed according to these associations, and again most of the time these associations are intuitive and work just as you would expect. This is an important observation for the first phase of our project in which we don't specifically use data-access patterns of particular applications. Instead, for the first phase of the project, our goal is to come up with generic read-ahead rules that are applicable across applications. An important property of applications for systems such as health-care, financial, or human resources is repetitive usage of the same query templates. For example, an application can request the due date of a credit card payment after requesting the balance, or it can request the transactions of the same card in the last billing period. The similar associations and dependencies that can be found in different domains form a basis for guessing the useful generic access patterns in our project.

Generic access patterns aren't enough to determine the most beneficial and efficient read-ahead scheme configuration. These patterns are helpful to determine the useful sequence of SQL statements. On the other hand, we need to find efficient sequences of (merged) SQL statements. Finding such efficient sequences includes answering the following questions:

1. How much to read ahead? This question requires figuring out which tables and table columns may be subsequently accessed, and how the structure of data (e.g., existence of an indexed attribute) and relational constraints affect the structure of a read-ahead query.
2. How deep to read ahead? This question requires figuring out the levels of the object hierarchy that may be subsequently accessed. How does the number of joins affect the efficiency of sequence of SQL statements?
3. In what direction to read ahead? In each level of the object hierarchy, each object can be associated with many different objects. This question requires figuring out the trade-off of typical directions of the traversal on the object hierarchy that is stored in the relational database.

5 Research Methodology

We construct a testbed to experiment with the parameters that affect the cost of SQL statements used while reading ahead of time. The focus of our case study is to discover and experiment with the data-access patterns of financial applications with parameters that are related to the structure of the data in the database. For this case study, we use a slightly modified version of the standard data model for financial services (e.g., banking and investment services) as described in [10]. Some of the main entities of this model are shown in Figure 4. We use Oracle 9i as our data server to implement the entire data model which has 55 tables and maximum fan-out of 7 relationships. After creating a sample database for this model, we list possible query templates for the model. Here, we list some query templates.

- Find the attributes of a given customer
- Find the accounts for a given customer
- Find the agreements to which a given financial product(s) is related
- Find the assets of a given customer that are used for loan agreements
- Find the account transactions of a given account
- Find the names of the customers that have more than a given number of accounts

The above list can be easily extended. By using these potential query templates, we can come up with meaningful generic access sequences for the database.

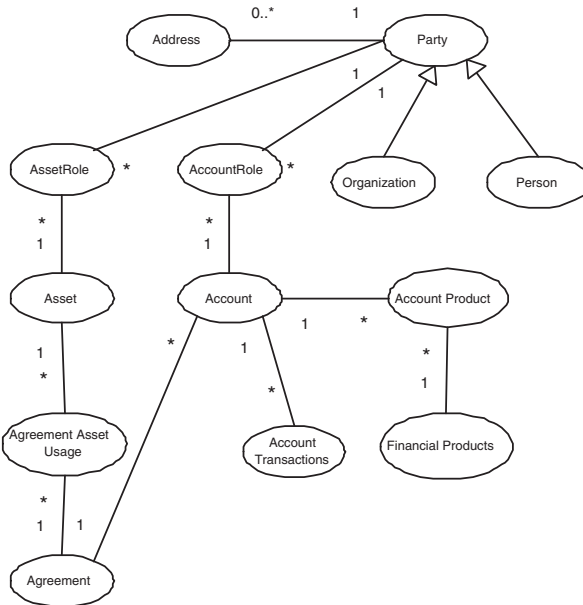


Fig. 4. Some entities in the financial services model

We generate SQL statements that implement access patterns, which are part of the given access sequence. To generate various combinations of SQL statements for the given access sequence with read-ahead functionality, we can use the following prefetch methods:

- Prefetch only primary keys of the associated objects
- Prefetch primary and non-primary foreign keys
- Prefetch key and non-key attributes, or only non-key attributes
- Prefetch via traversing the inheritance-extension, one-to-many association (i.e., aggregation), or many-to-many type of relationships

We model the data accesses as a directed graph. In this directed graph, vertices correspond to simple data-access statements and edges correspond to transition probabilities between the statements. These probabilities can be calculated and updated using access

log files of the database server. For example, if we consider each data-access request used in Example 2.1 as simple data-access statement, we obtain the simple *data-access statement transition graph* in Figure 5. Based on this simple hypothetical transition graph, whenever A is accessed there is a 40% chance that B will be accessed next. We assume an acyclic graph that can have multiple sources and sinks for the first phase of the project.

To initialize large transition graphs for financial services applications, we generate vertices by using above query templates and generate edges with random probabilities by above prefetch methods that use generic associations and dependencies among relations.

We need to merge simple data-access statements to find optimal sequences of SQL statements for the given access sequence. There are two principle ways to combine two or more SQL statements. Either the statements are executed *at the same time* or they are executed *one after the other*. We call the first combination *merged execution* and denote it by operator $*$; the second combination represents *sequential execution* and is denoted by operator $+$. We can apply $*$ repeatedly to describe access patterns, and apply $+$ repeatedly to combine patterns to form the access sequence. By definition, $*$ is commutative, while $+$ is not, and $*$ has precedence over $+$. For example, if we have an access sequence $ADEF$ for a larger graph in Figure 6, we can come up with compound statements such as $A*D*E+F$, $A*D*E*F$, $A*C*D+E*F$, $A*B*D+E*F$, $A+D*E*F$ and so on. By using large acyclic graphs, we can consider more complex access patterns. To find efficient compound statements, we derive a benefit formula, which takes into account the processing cost and the probabilities of each simple and compound statements that can generate the given access sequence. This formula also includes a risk factor, which aims to balance the cost and usefulness of the generated sequence. According to the query workload or even type of applications, the risk factor can change.




Fig. 5. Query transition graph for Example 2.1




Fig. 6. A sample query transition graph

By using the results of the benefit formula and the computed cost of an access sequence, e.g., processing cost for $A+D+E+F$, we will determine rules of thumb regarding what strategies to use for efficient sequences of SQL statements for the financial services domain. Then, we will test and fine-tune the rules on other application domains.

6 Discussion

The results of our work can be used to effectively preload associated dataset of a requested object that may be subsequently accessed. Our work can also be integrated with query optimizers to find the efficient execution plans for compound SQL statements. This integration provides us another aspect for multi-query optimization.

This work can also be integrated with Container Managed Persistence containers or Java Data Object drivers as a performance tuning technique. Using this technique, we will be able to get ahead of time the working set of objects for the current transaction in very concurrent environments. Also this work can be helpful in determining the right cache size for systems that use prefetching.

One of the important implementation challenges is to merge simple SQL statements on-the-fly to experiment with the compound SQL statements. This isn't a trivial task, because it requires the detailed knowledge of the relational structure of data. Even for the statements that are syntactically similar, we can have different merged SQL statements.

We use applications' common behaviours and the cost of database interactions to generate read-ahead rules that can provide a significant performance gain for systems where many concurrent data-intensive read-only applications access huge databases.

Acknowledgements

I would like to thank Rada Chirkova and Munindar P. Singh for their guidance and suggestions. I would also like to thank Timo Salo for his helpful comments. This research is supported under NCSU CACC Grant 11019.

References

1. Sibel Adali, K. Selçuk Candan, Yannis Papakonstantinou, and V. S. Subrahmanian. Query caching and optimization in distributed mediator systems. In *Proc. 1996 ACM SIGMOD Conf. on Management of Data*, pages 137–148, 1996.
2. Philip A. Bernstein, Shankar Pal, and David Shutt. Context-based prefetch – An optimization for implementing objects on relations. *VLDB Journal*, 9(3):177–189, 2000.
3. Sunil Choenni, Martin Kersten, Amani Saad, and Johan van den Akker. A framework for multi-query optimization. In *Proc. 8th Int. Conf. on Management of Data (COMAD'97)*, pages 165–82, 1997.
4. Brian D. Davison. *The Design And Evaluation Of Web Prefetching and Caching Techniques*. PhD thesis, Department of Computer Science, Rutgers University, October 2002.
5. Daniela Florescu, Alon Levy, Dan Suciu, and Khaled Yagoub. Optimization of run-time management of data intensive Web sites. In *Proc. 25th VLDB Conf.*, pages 627–638, September 1999.
6. Laura M. Haas, Donald Kossmann, and Ioana Ursu. Loading a cache with query results. In *Proc. 25th VLDB Conf.*, pages 351–362, 1999.
7. Arthur M. Keller and Julie Basu. A predicate-based caching scheme for client-server database architectures. *VLDB Journal*, 5(1):35–47, 1996.
8. Tom M. Kroeger, Darrell D. E. Long, and Jeffrey C. Mogul. Exploring the bounds of Web latency reduction from caching and prefetching. In *USENIX Symposium on Internet Technologies and Systems*, 1997.

9. Mark Palmer and Stanley B. Zdonik. FIDO: A cache that learns to fetch. In *Proc. 17th VLDB Conf.*, pages 255–264, Barcelona, Spain, 1991.
10. Len Silverston. *The Data Model Resource Book*, volume 2. John Wiley and Sons, New York, 2001.
11. Arun Venkataramani, Praveen Yalagandula, Ravindranath Kokku, Sadia Sharif, and Mike Dahlin. The potential costs and benefits of long term prefetching for content distribution. In *Proc. of Web Content Caching and Distribution Workshop*, 2001.
12. Dazhi Wang and Junyi Xie. An approach toward Web caching and prefetching for database management systems. Technical report, Department of Computer Science, Duke University, 2001. <http://www.cs.duke.edu/~junyi/cps216/report.pdf>.