

DISTRIBUTED QUERY PROCESSING ON THE GRID

Jim Smith¹
Anastasios Gounaris²
Paul Watson¹
Norman W. Paton²
Alvaro A.A. Fernandes²
Rizos Sakellariou²

Abstract

Distributed query processing (DQP) has been widely used in data intensive applications where data of relevance to users are stored at multiple locations. This paper argues: (i) that DQP can be important in the Grid, as a means of providing high-level, declarative languages for integrating data access and analysis; and (ii) that the Grid provides resource management facilities that are useful to developers of DQP systems. As well as discussing and illustrating how DQP technologies can be deployed within the Grid, the paper describes Polar*, a prototype implementation of a DQP system running over Globus. Polar* can handle complex data by adopting the ODMG object model and its query language OQL, which supports the invocation of user-defined operations. The Globus components are accessed through the MPICH-G interface rather than in a lower level way. A case study from bioinformatics is used throughout the paper, to show the benefits of the approach.

Key Words: grid computing, distributed query processing, parallel query processing, user-defined operation, globus, MPI, ODMG

1 Introduction

To date, most work on data storage, access and transfer on the Grid has focused on files. We do not take issue with this – files are clearly central to many applications, and it is reasonable for Grid middleware developers to seek to put in place effective facilities for file management and archiving. However, database management systems provide many facilities that are recognized as being important to Grid environments, both for managing Grid metadata (e.g. Dinda and Plale, 2001) and for supporting the storage and analysis of application data (e.g. Szalay et al., 2000).

In any distributed environment there are inevitably multiple related data resources, which, for example, provide complementary or alternative capabilities. Where there is more than one database supported within a distributed environment, it is straightforward to envisage higher-level services that assist users in making use of several databases within a single application. For example, in bioinformatics, it is commonly the case that different kinds of data (e.g., DNA sequence, protein sequence, protein structure, transcriptome) are stored in different, specialist repositories, even though they are often inter-related in analyses.

There are perhaps two principal functionalities associated with distributed database access and use: distributed transaction management and distributed query processing (DQP; Kossmann 2000). This paper is concerned with DQP on the Grid, and both (i) discusses the role that DQP might play within the Grid and (ii) describes a prototype infrastructure for supporting distributed query optimization and evaluation within a Grid setting.

There is no universally accepted classification of DQP systems. However, with a view to categorizing previous work, we note that DQP is found in several contexts: in distributed database systems, where an infrastructure supports the deliberate distribution of a database with some measure of central control (Ozsu and Valduriez, 1999); in federated database systems, which allow multiple autonomous databases to be integrated for use within an application (Hsiao, 1992); and in query-based middlewares, where a query language is used as the programming mechanism for expressing requests over multiple wrapped data sources (e.g., Haas et al., 1997; Garcia-Molina et al., 1997; Davidson et al., 2001). This paper is most closely related to the third category, in that we consider the use of DQP for integrating various Grid resources, including (but not exclusively) database systems.

¹ DEPARTMENT OF COMPUTING SCIENCE,
UNIVERSITY OF NEWCASTLE UPON TYNE, UK

² DEPARTMENT OF COMPUTER SCIENCE,
UNIVERSITY OF MANCHESTER, UK

Another important class of system in which queries run over data that is distributed over a number of physical resources is parallel databases. In a parallel database, the most common pattern is that data from a centrally controlled database are distributed over the nodes of a parallel machine. Parallel databases are now a mature technology, and experience shows that parallel query processing techniques are able to provide cost-effective scalability for data-intensive applications (e.g., DeWitt and Gray, 1992; Graefe, 1993; Hasan et al., 1996; Sampaio et al., 2001). This paper, as well as advocating DQP as a data integration mechanism for the Grid, also shows that techniques from parallel database systems can support data access and analysis for Grid applications.

The claims of this paper with respect to DQP and the Grid are as follows:

1. In providing integrated access to multiple data resources, DQP in and of itself is an important functionality for data intensive Grid applications.
2. The fact that certain database languages can integrate operation calls with data access and combination operations means that DQP can provide a mechanism for integrating data and computational Grid services.
3. Given (1) and (2), DQP can be seen to provide a generic, declarative, high-level language interface for the Grid.
4. By extending technologies from parallel databases, implicit parallelism can be provided within DQP environments on the Grid.

The paper makes concrete how these claims can be supported in practice by describing a prototype DQP system, Polar*, which runs over the Globus toolkit, and illustrates the prototype using an application from bioinformatics.

The remainder of this paper is structured as follows. In Section 2 we present the principal components of a DQP system for the Grid, in particular indicating how this relates to other Grid services. In Sections 3 and 4 we describe, respectively, how queries are planned and evaluated within the architecture. In Section 5 we describe some measurements of an example query in a Grid environment. Finally, in Section 6 we present some conclusions and pointers to future work.

2 Architecture

The two key functions of any DQP system are query compilation and query execution. In the Polar* system described in this paper, both these components are based on those designed for the Polar project (Smith et al., 2000). Polar is a parallel object database server that runs

on a shared-nothing parallel machine. Polar* exploits Polar software components where possible but, as Polar* must provide DQP over data repositories distributed across a Grid, there are a number of key differences between query processing in the two systems. These include:

1. The data dictionary must describe remote data storage and analysis resources available on the Grid – queries act over a diverse collection of application stores and analysis programs.
2. The scheduler must take account of the computational facilities available on the Grid, along with their variable capabilities – queries are evaluated over a diverse collection of computational resources.
3. The data stores and analysis tools over which queries are expressed must be wrapped so that they look consistent to the query evaluator.
4. The query evaluator must use Grid authentication, resource allocation and communication protocols – Polar* runs over Globus, using MPICH-G (Foster and Karonis, 1998).

A representative query over bioinformatics resources is used as a running example throughout the paper. The query accesses two databases: the Gene Ontology Database *GO* (www.geneontology.org; Ashburner et al., 2000) stored in a MySQL (www.mysql.com) RDBMS; and *GIMS* (Cornell et al., 2001), a genome database running on a Polar parallel object database server, based on the *Shore* system (Carey et al., 1994). The query also calls a local installation of the BLAST sequence similarity program (www.ncbi.nlm.nih.gov/BLAST/; Altschul et al., 1990) which, given a protein sequence, returns a set of structs containing protein IDs and similarity scores. For that, the BLAST program accesses a third biological database, which in the example is *SWISS-PROT* (Bairoch and Apweiler, 2000), in a way transparent to the user. Both the BLAST program (and more specifically its *blastall* operation) and its associated database comprise a BLAST server. The query identifies proteins that are similar to human proteins with the GO term 8372:¹

```
select p.proteinId, Blast(p.sequence)
from p in proteins, t in proteinTerms
where t.termID='8372' and
       p.proteinId=t.proteinId
```

In the query, *proteins* is a class extent in *GIMS*, while *proteinTerms* is a table view in *GO*. Before submitting the query, a global database schema has been constructed to describe and combine the relevant views of the participating databases, along with the BLAST program. The ODL (Cattell and Barry, 2000) definition of that schema is as follows:

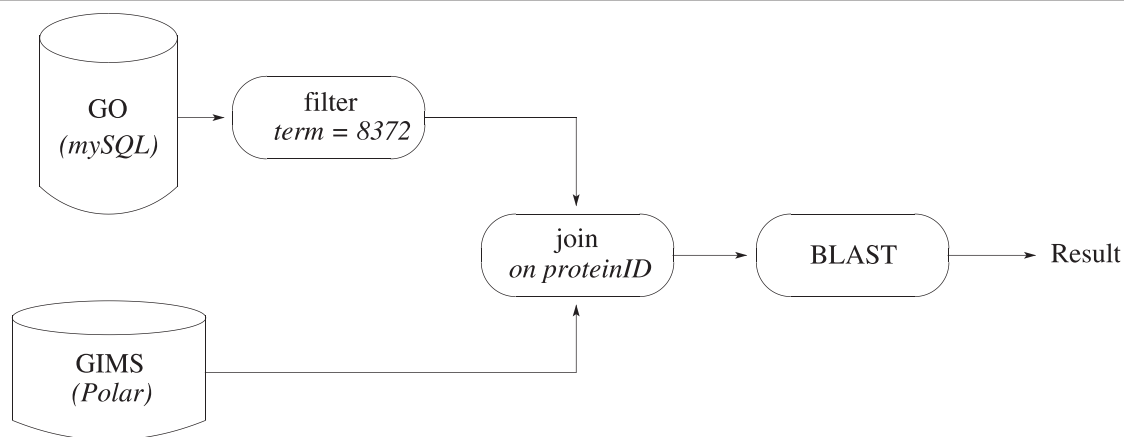


Fig. 1 Evaluating the example query.

```

forward class Protein;
forward class ProteinTerm;

struct Entry {
  string BLASTproteinID;
  long score;
};

static set<Entry> blast(in string protein);

class Protein //from GIMS
(extent proteins) {
  attribute string proteinID;
  attribute string sequence;
};

class ProteinTerm //from GO
(extent proteinTerms) {
  attribute string proteinID;
  attribute string term;
};
  
```

Therefore, as illustrated in Figure 1, the infrastructure initiates two subqueries: one on GIMS, and the other on GO. The results of these subqueries are then joined in a computation running on the Grid. Finally, each protein in the result is used as a parameter to the call to BLAST.

One key opportunity created by the Grid is in the flexibility it offers on resource allocation decisions. In the example in Figure 1, machines need to be found to run both the join operator, and the operation call. If there is a danger that the join will be the bottleneck in the query, then it could be allocated to a system with large amounts of main memory so as to reduce IO costs

associated with the management of intermediate results. Further, a parallel algorithm could be used to implement the join, and so a set of machines acquired on the Grid could each contribute to its execution. Similarly, the BLAST calls could be speeded up by allocating a set of machines, each of which can run BLAST on a subset of the proteins.

The information needed to make these resource allocation decisions comes from two sources. First, the query optimizer estimates the cost of executing each part of a query and so identifies performance-critical operations. Secondly, the Globus Grid infrastructure provides information on available resources. Once a mapping of operations to resources has been chosen, the single sign-on capabilities of the Grid Security Infrastructure simplify the task of gaining access to these resources. Figure 2 illustrates the high level architecture of the system. The metadata repository, which is needed to compile a query, is populated by both data from the database schema (e.g. types of attributes), and by information about the Grid nodes. Each Grid node has specific computational capacities and data stores. The coordinator node is a usual Grid node with the additional characteristic that it has the Polar* query compiler installed.

3 Query Planning

Polar* adopts the model and query language of the ODMG object database standard (Cattell and Barry, 2000). As such, all resource wrappers must return data using structures that are consistent with the ODMG model. Queries are written using the ODMG standard query

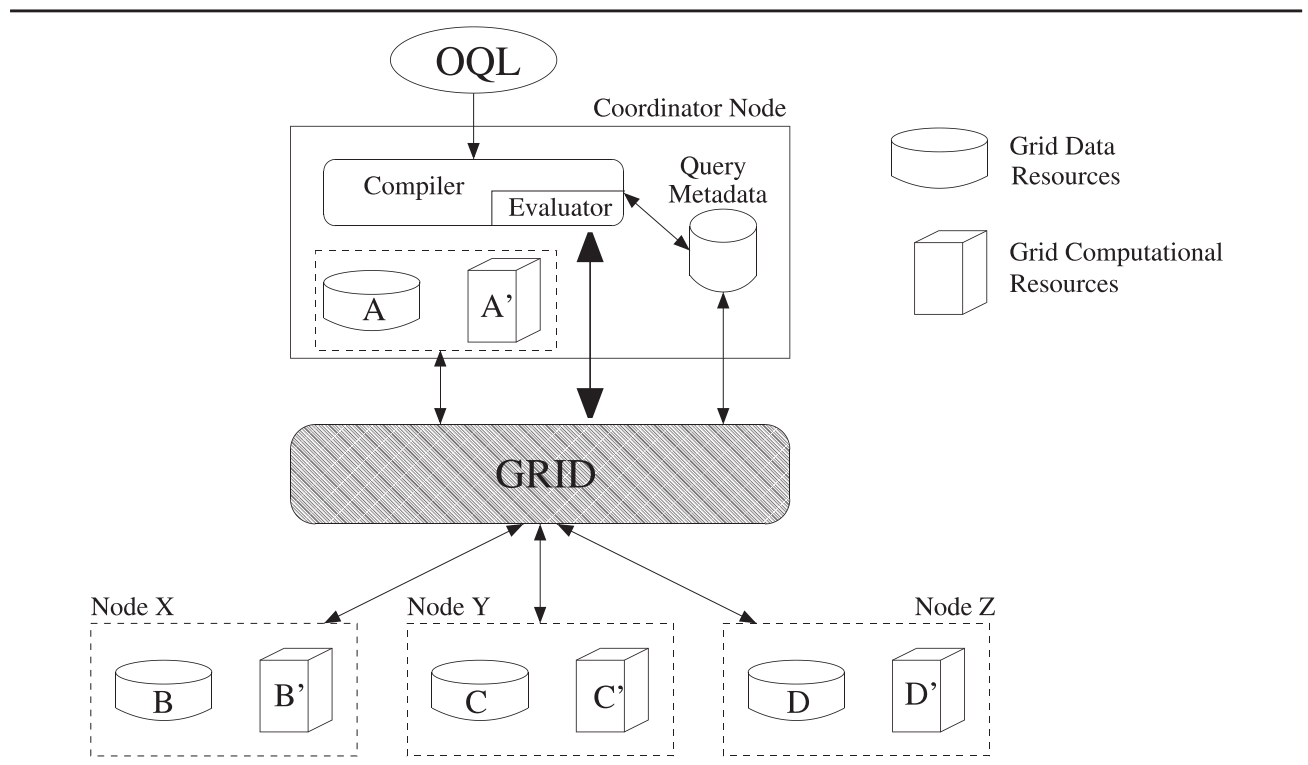


Fig. 2 The Polar* architecture.

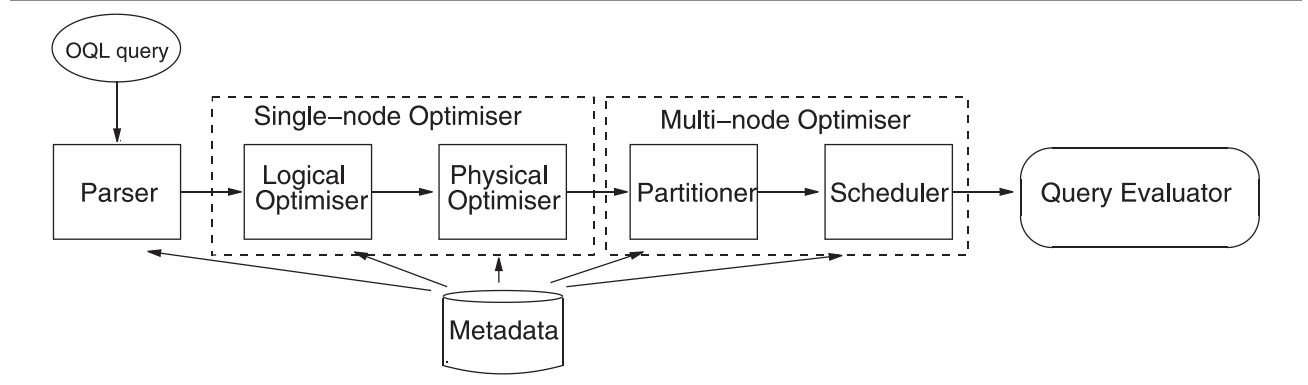


Fig. 3 The components of Polar* query compiler.

language, OQL. The reason for choosing the object data model rather than the relational one is that it provides a richer model for source wrapping and for representing intermediate data.

The main components of the query compiler are shown in Figure 3. The Polar* optimizer has responsibility for

generating an efficient execution plan for the declarative OQL query which may access data and operations stored on many nodes. To do this, it follows the two-step optimization paradigm, which is popular for both parallel and distributed database systems (Kossmann, 2000). In the first phase, the single node optimizer produces a query

plan as if it was to run on one processor. In the second phase, the sequential query plan is divided into several partitions or subplans which are allocated machine resources by the scheduler.

3.1 CONSTRUCTION OF THE SINGLE-NODE PLAN

Figure 4(a) depicts a plan for the example query expressed in the logical algebra of Fegaras and Maier (2000), which is the basis for query optimization and evaluation in Polar*. The logical optimizer performs various transformations on the query, such as fusion of multiple selection operations and pushing *projects* – called *reduce* in Fegaras and Maier (2000) and in the figures – as close to *scans* as possible.

The physical optimizer transforms the optimized logical expressions into physical plans by selecting algorithms that implement each of the operations in the logical plan (Figure 4(b)). For example, in the presence of indices, the optimizer prefers *index_scans* to *seq_scans*. Operation calls, such as the call to BLAST, are encapsulated by the *operation_call* physical operator. For each OQL query, many physical plans are produced, and the physical optimizer ranks these according to a cost metric based on the predicted size of the intermediate results produced during the query execution. Changing the order of the operators results in plans with different costs.

3.2 CONSTRUCTION OF THE MULTI-NODE PLAN

A single-node plan is transformed into a multi-node one by inserting parallelization operators into the query plan, i.e., Polar* follows the operator model of parallelization (Graefe, 1990). The *exchange* operator encapsulates flow control, data distribution and inter-process communication. The partitioner first identifies whether an operator requires its input data to be partitioned by a specific attribute when executed on multiple processors (for example, so that the potentially matching tuples from the operands of a join can be compared). These operators are called *attribute sensitive* operators (Hasan and Motwani, 1995), and Sampaio et al. (1999) present the classification of the parallel operators as attribute sensitive or attribute insensitive. Secondly, the partitioner checks whether data repartitioning is required, i.e., whether data need to be exchanged among the processors, for example for joining or for submitting to an *operation_call* on a specific machine. If the children of an attribute sensitive operator in a query plan are partitioned by an attribute other than its partitioning attribute, or there is no data partitioning defined, then data repartitioning needs to take

place. Another case for data repartitioning is when not all the candidate nodes for evaluating parts of the query can evaluate a physical operator, which is the usual case for *operation_calls*.

The physical algebra extended with *exchange* constitutes the parallel algebra used by Polar*. The *exchanges* are placed immediately below the operators that require the data to be repartitioned. For each *exchange* operator, a data distribution policy needs to be defined. Currently, the policies Polar* supports include *round_robin*, *hash_distribution* and *range_partitioning*. The last two policies provide support for non-uniform data distribution among instances of the same physical operator, which is desirable for heterogeneous environments. A multi-node query plan is shown in Figure 4(c), where the *exchanges* partition the initial plan into many subplans (delimited by dashed lines in the figure). In the example query, there is one attribute sensitive operator, the *hash_join*, which needs to receive tuples partitioned by the *proteinID* attribute. Since the children of the *hash_join* are not partitioned in this way, two *exchanges* are inserted at the top of each subplan rooted by the join. *Operation_call* is attribute insensitive, but since it can be called only from specific nodes, an *exchange* is inserted as its child. The optimizer checks if the *exchanges* transmit redundant data. A *reduce* operator ensures that only data that are required by operators higher in the query plan are transferred through the network. If there are no such *reduces* placed by the single-node optimizer, the parallel optimizer inserts them.

The final phase of query optimization is to allocate machine resources to each of the subplans derived from the partitioner, a task carried out by the scheduler in Figure 3 using an algorithm based on that of Rahm and Marek (1995). For running the example query, suppose that six machines are available, and that three of the machines host databases (numbers 2 and 3 for the GIMS database, and 6 for the GO database). The *table_scan* operators are placed on these machines in order to save communication cost. For the *hash_join*, the scheduler tries to ensure that the relation used to construct the hash table can fit into main memory, for example, by allocating more nodes to the join until predicted memory requirements are satisfied or all available memory is exhausted. In the example, nodes 3 and 6 are allocated to run the *hash_join*. As some of the data are already on these nodes, this helps to reduce the total network traffic.

The data dictionary records which nodes support BLAST, and thus the scheduler is able to place the *operation_call* for BLAST on suitable nodes (4 and 5 in Figure 4(c)). The cost of a call to BLAST is much higher than the cost to send a tuple to another node over the network. In any case, the whole set of the results of the *hash_join* needs to be moved from the nodes 3 and 6. For these two reasons, the optimizer has chosen the

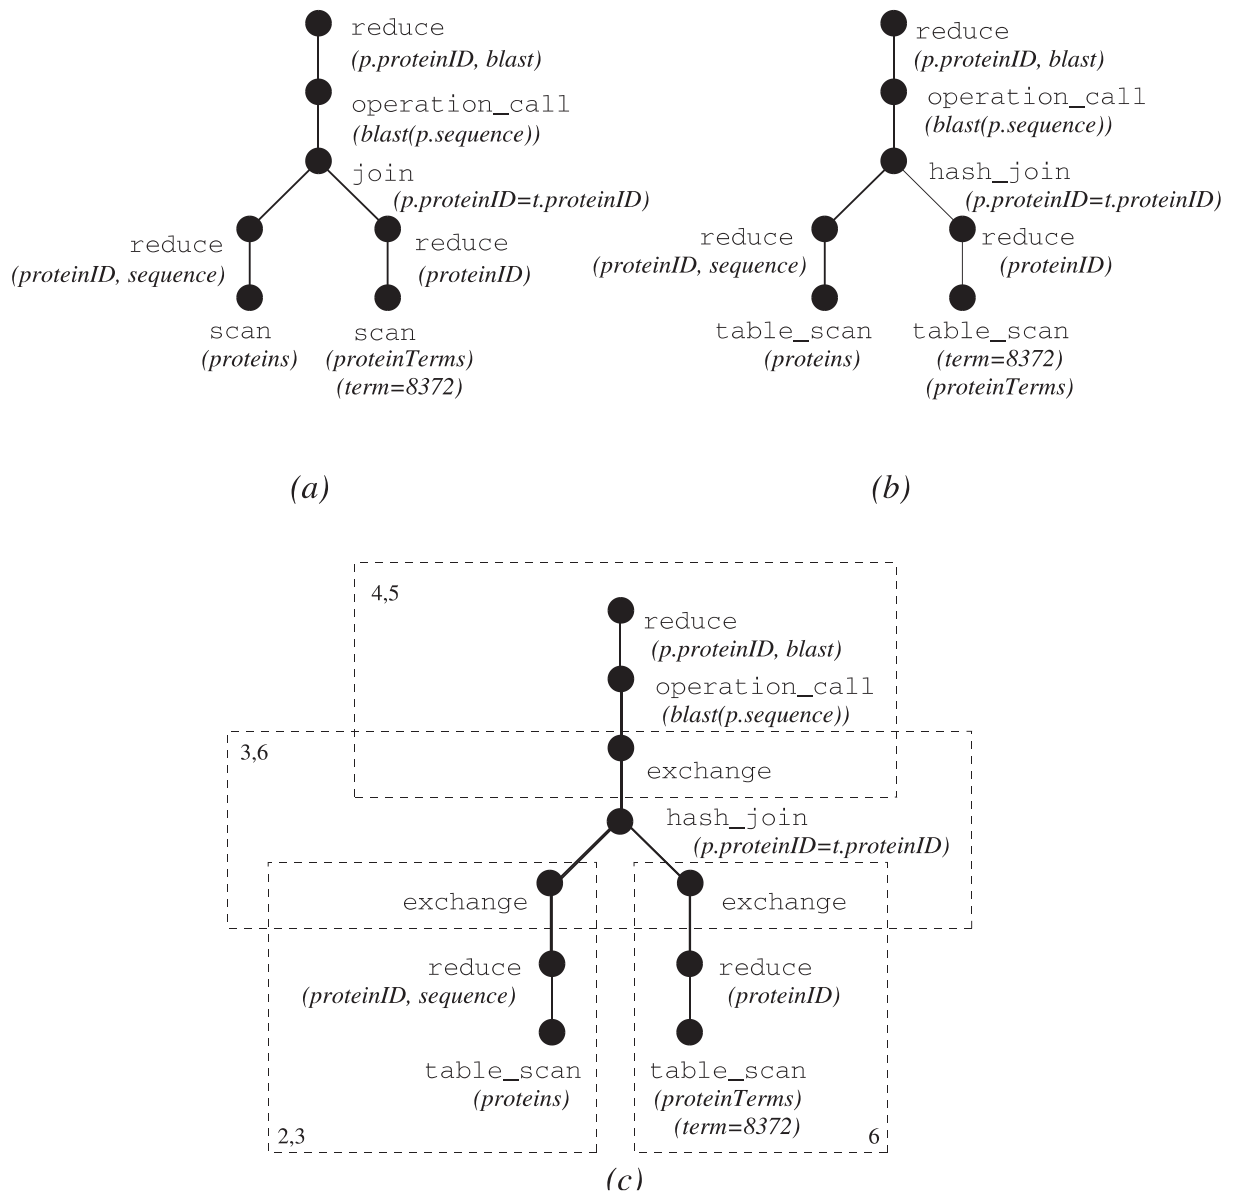


Fig. 4 Example query: (a) single-node logical plan, (b) single-node physical plan, (c) multi-node physical plan.

maximal degree of parallelism for the BLAST operation. The scheduler uses a heuristic that may choose not to use an available evaluator if the reduction in computation time would be less than the increase in the time required to transfer data (e.g., it has decided not to use machine *l* in the example).

4 Query Evaluation

4.1 EVALUATING THE PARALLEL ALGEBRA

The Polar* evaluator uses the *iterator* model of Graefe (1993), which is widely seen as the model of choice for

```

class HashJoin: public Operator {
private:
    Tuple *t; HashTable h; Predicate predicate;
    Operator *left; set<Attribute> hash_atts_left;
    Operator *right; set<Attribute> hash_atts_right;
public:
    virtual void open() {
        left->open(); t = left->next();
        while (! t->is_eof()) {
            h.insert(t, hash_atts_left); t = left->next();
        }
        left->close(); right->open(); t = right->next();
    }
    virtual Tuple *next() {
        while (! t->is_eof()) {
            if (h.probe(t, hash_atts_right) && t->satisfies(predicate))
                return t;
            delete t; t = right->next();
        }
        return t;
    }
    virtual void close() {
        right->close(); h.clear();
    }
};

```

Fig. 5 Implementing *hash-join* as an iterator.

parallel query processing (Kossmann, 2000). In this model, each operator in the physical algebra implements an interface comprising three operations: *open()*, *next()* and *close()*. These operations form the glue between the operators of a query plan. An individual operator calls *open()* on each of its input operators to prompt them to begin generating their result collections. Successive calls to *next()* retrieve every tuple from that result collection. A special *eof* tuple marks the end of a result collection. After receiving an *eof* from an input, an operator calls *close()* on that input to prompt it to shut itself down.

We note that, although the term *tuple* is used to describe the result of a call to *next()*, a tuple in this case is not a flat structure, but rather a recursive structure whose attributes can themselves be structured and/or collection valued.

To illustrate the iterator model, Figure 5 sketches an iterator-based implementation of a *hash_join* operator. *open()* retrieves the whole of the left-hand operand of the join and builds a hash table, by hashing on the attributes for which equality is tested in the join condition. In *next()*, tuples are received from the right input collection and used to probe the hash table until a match is

found and all predicates applying to the join result are satisfied, whereupon the tuple is returned as a result.

The iterator model can support a high degree of parallelism. Sequences of operators can support *pipeline parallelism*, i.e., when two operators in the same query plan are independent, they can execute concurrently. Furthermore, when invocations of an operation on separate tuples in a collection are independent, the operation can be partitioned over multiple machines (i.e., *partitioned* or *intra-operator parallelism*).

Whereas data manipulation operators tend to run in a request-response mode, *exchange* differs in that, once *open()* has been called on it, the producers can run independently of the consumers. The tuples, which may be complex structures, are flattened for communication into buffers whose size can be configured. Underlying an instance of *exchange* is a collection of threads managing pools of such buffers so as to constrain flow to a lagging consumer, but to permit flow to a quicker consumer, within the constraints of the buffer pools. This policy is very conveniently implemented in MPI (Snir et al., 1998) where the tightly defined message completion semantics permit the actual feedback to be hidden within

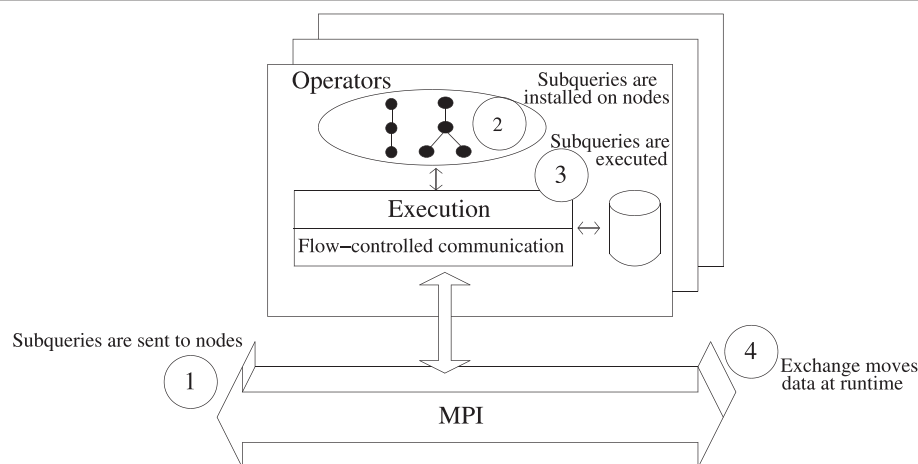


Fig. 6 The different phases involved in planting and executing a query on Grid nodes. Phases 1 and 2 are executed once for each query at the beginning, while phases 3 and 4 are executed concurrently many times.

its layer. The *send* and *recv* calls issued at the MPI level in non-blocking synchronous mode, are managed by multiple user level threads encapsulated in the *exchange* operator. Each Polar* machine contains a global buffer pool, whose size is controlled via the global metadata. During evaluation of a given query, flow control is achieved by restricting each instance of the exchange operator to an allocation of buffers from the global pool; it can only send or receive buffers within the capacity of that allocation. This use of MPI has enabled the Polar *exchange* to port easily to MPICH-G (Foster and Karonis, 1998), for use with Globus.

Since MPICH-G is layered above Globus, a parallel query can be run as a parallel MPI program over a collection of wide area distributed machines, oblivious of the difficulties inherent in such meta-computing, which are handled by the underlying Globus services. A parallel program running over such a Grid environment has to find suitable computational resources, achieve concurrent login, transfer of executables and other required files, and startup of processes on the separate resources. In the Polar* prototype, concurrent login to separate accounts is achieved through GSI (the Grid Security Infrastructure for enabling secure authentication and communication over an open network), and executable staging across wide area connections through GASS (Global Access to Secondary Storage, which is used for remotely accessing data), but all these features are accessed through the MPICH-G interface rather than in a lower level way.

Because the physical operators in Polar*, such as Polar, are implemented to manipulate generic tuple structures, the implementation of each operator can be present as a C++ object in each installation of the evaluator code. A subquery thus only needs to contain a reference to the specific operator IDs (i.e. *scan*, *hash_join*, *exchange*, etc.) and the parameters for those instances. These parameters include the index of each input operator and the operator specific parameters, such as the arbitrator policy and list of consumers for an *exchange*, the join attributes and predicate for a *hash_join*, etc. While input operators are contained in the same subquery, exchange consumers can be identified by a machine number (i.e. rank number) and index number pair. Thus the description of a query subplan is quite simple, and can be expressed concisely in textual format. By contrast, the dynamically loadable module required by an operation call, containing the user defined code, is installed as part of the procedure of registering that operation call as available on the particular machine.

The collection of machines included in a Polar* distributed system is recorded in the global metadata of that system, located on the coordinator machine. Polar* generates files required by the underlying MPICH-G system from this metadata. Assuming the user has built the coordinator on his own machine, the OQL compiler/optimizer can be called as a local process by the *polar-query* program to generate a collection of subqueries, and a corresponding list of database nodes. The *polar-query* program can thereby construct a single communicator MPI


```

class PhysicalOperationCall: public PhysicalOperator {
private:
    string signature;           // operation to call
    string opslib_name;        // name of library
    list<Expression> expression; // select params from tuple
    Predicate predicate;       // predicate on output tuple
    int key;                   // from cross ref in opslib
    vector<concrete_object*> arg; // sized appropriately
    class operation_library_stub; // functions etc in shared library
    operation_library_stub *stub;
    Tuple *t;
public:
    virtual void open() {
        input->open();          // input is pointer to operator
        stub = load_operation_library(opslib_name);
        key = stub->xref(signature);
    }
    virtual Tuple *next() {
        while (! t->eof()) {
            t = input->next();
            for (list<Expression>::iterator it = expression.begin();
                it != expression.end(); it++)
                arg[i+1] = t->evaluate(*it); // leave arg[0] for result
            stub->call_operation(key, arg);
            t->insert(arg[0]);
            if (t->evaluate(predicate))
                return t;
            delete t;
        }
        return t;
    }
    virtual void close() {
        unload_operation_library(stub);
        input->close();
    }
}

```

Fig. 7 The iterator-based implementation of external *operation calls*.

program with one process, rank 0, running on the local machine and the remainder running on those machines that are required to participate in the query, and supply the subqueries to that MPI program via a text file.

Figure 6 shows the four main phases of query execution. The evaluator receives a set of subqueries from the query compiler as shown in Figure 3. First, the subqueries are sent to the Grid nodes through the network and, secondly, they are installed on them via the MPI interface. The next phase involves the execution of the operators comprising the query plan as iterators. This operator execution may in turn lead to the moving of data between nodes, using a flow-controlled communications infrastructure which itself uses the MPI interface.

4.2 ACCESSING COMPUTATIONAL AND DATA RESOURCES DURING QUERY EVALUATION

From the point of view of the evaluator, both database and analysis operations referred to within a query are external resources, which must be wrapped to enable consistent passing of parameters and returning of results during query evaluation.

To support access to external tools such as BLAST, Polar* implements, in *iterator* style, the *operation_call* operator in Figure 7. While concerns regarding the integrity of the evaluator lead to the isolation of such an operation call within a separate user context in some

databases, the view taken so far in Polar* is that higher performance is achievable when a user can in good faith code an operation to be executed in the server's own context. Thus, such operations are linked into dynamically loadable modules, with stub code generated by the Polar* system to perform the translation of types between tuple format and the format of the programming language used to implement the operation. At runtime, *operation_call* loads the appropriate module in its *open()* call and unloads it in *close()*. Within *next()*, the operator passes attributes from the current tuple to the stub code which assembles required parameters, makes the call and translates the result, which may be collection valued and/or of structured type, back to tuple format before inserting it into the output tuple which is passed back as the result of the operator.

By making such an application available for use in a Polar* schema, the owner does not provide unrestricted, and thereby unmanageable, access on the Internet. By contrast, since Polar* sits above Globus, a resource provider must have granted access rights to a user, in the form of a local login, which they can revoke. Subsequent accesses, while convenient to the user through the single sign-on support of Globus, are authenticated through the GSI.

A specific case that requires access to external operations is the provision of access to external, i.e. non Polar*, repositories.

For example, the runtime interface to a repository includes an *external_scan* operator, which exports an *iterator* style interface in common with other operators (Figure 8). However, below the level of *external_scan* the interface to an arbitrary external repository requires special coding. When an external collection is defined in the schema, the system generates the definition of a class, which has the three operations of the *iterator* interface plus other operations to gather statistics and set attributes of the interface such as a query language and the syntax of results generated by the external repository.

The generated class is the template within which a user can implement an *iterator* style interface to the external repository. The mechanisms used are at the discretion of the user, but the pattern in which the operations of this system-defined class are called is defined by the Polar* system. Fixed attributes, such as the syntax of result tuples returned, are set at the time the completed access class is registered with the Polar* system. The *open()*, *next()* and *close()* operations are simply called by *external_scan* in the usual iterator-based style. However, the results returned by the *next()* operation are translated from the selected syntax into Polar* tuples.

In the running example, the *GO* database is implemented using MySQL, so access to it is through such a

system-specified user-defined class of operations. A delegated subquery such as the access to *proteinTerm* tuples is expressed in SQL, and the results are formatted within the *MySQLAccess* class in Object Interchange Format (OIF).² The *next()* operation of the *external_scan* operator parses each OIF instance to construct its result tuple.

5 Case Study

The experiments described in this section illustrate the behavior of the Polar* system when evaluating the example query introduced in Section 2. The query is run in various different configurations, and in cases where parameters of the query are varied, to indicate how these influence scheduling decisions. For this purpose, it is only query response time that is of concern here.

The case study has been conducted on the heterogeneous collection of machines described below.

– At Newcastle upon Tyne.

| Machine | CPU (MHz) | Memory (Mbyte) | OS | Globus |
|----------------|-----------|----------------|------------------|--------|
| <i>ncl-pc1</i> | 500 | 128 | RedHat Linux 8.0 | GT 2.0 |
| <i>ncl-pc2</i> | 700 | 256 | RedHat Linux 8.0 | GT 2.2 |
| <i>ncl-pc3</i> | 2 × 800 | 256 | RedHat Linux 7.3 | GT 2.0 |

– At Manchester.

| Machine | CPU (MHz) | Memory (Mbyte) | OS | Globus |
|----------------|-----------|----------------|------------------|--------|
| <i>man-pc1</i> | N/A | N/A | N/A | no |
| <i>man-pc2</i> | 1100 | 512 | RedHat Linux 7.1 | GT 2.0 |

Man-pc1 is dedicated as a *MySQL* server for a reference *GO* database, with no direct user access. The wrapper for this database is managed by an *external_scan* operator located on the most local machine available, *man-pc2*, and accesses the *GO* data via a *MySQL* client. The machines in Newcastle are inter-connected via a 100 Mbps network. The bandwidth for the connection between Manchester and Newcastle is constrained by a slow connection between Leeds and Newcastle of 34 Mbps (www.norman.net.uk/old/janet/index.html).

The GIMS data (3 Mbytes in total) is located in a Polar parallel object database distributed over two volumes located on *ncl-pc3*. The subquery submitted to the *GO* database retrieves 31006 tuples, from which 2518

```

class PhysicalExternalScan: public PhysicalOperator {
private:
    string opslib_name;           // name of library
    string query_text;           // to be passed to external DB
    enum RESULT_FORMAT;         // e.g. OIF, SQL
    RESULT_FORMAT fmt;
    ObjectType otype;           // given object type and input format
    ObjectWalker *walker;       // walker supports map to/from tuple
    string open_signature;
    string next_signature;
    string close_signature;
    Predicate predicate;        // predicate on output tuple
    int open_key, next_key, close_key;
    vector<concrete_object*> open_arg; // only has string parameter
    vector<concrete_object*> next_arg; // only has result
    vector<concrete_object*> close_arg; // no parameters
    Tuple *t;
public:
    virtual void open() {
        stub = load_operation_library(opslib_name);
        open_key = stub->xref(open_signature);
        next_key = stub->xref(next_signature);
        close_key = stub->xref(close_signature);
        open_arg[1] = query_text;
        stub->call_operation(open_key, open_arg);
        walker = new ObjectWalker(otype, fmt);
    }
    virtual Tuple *next() {
        while (! t->eof()) {
            stub->call_operation(next_key, next_arg);
            t = walker->map(next_arg[0]);
            if (t->evaluate(predicate))
                return t;
            delete t;
        }
        return t;
    }
    virtual void close() {
        stub->call_operation(close_key, close_arg);
    }
}

```

Fig. 8 The iterator-based implementation of *external scans*.

are joined with the GIMS tuples. 35 tuples in the resulting dataset satisfy the join criterion.

The cost of a call to the *BLAST* program is several orders higher than a join between two tuples. Consequently, the query cost is expected to be dominated by the *BLAST* calls, so this is where parallelization is attempted. Under different circumstances, parallelizing other operators as well (such as *hash_joins*) has significant impact, as explained in Section 3. In this case

study, all machines can support a *BLAST* server apart from *man-pc1*, and the SMP *ncl-pc3* can support two. The *SWISS-PROT* dataset referenced in the *BLAST* calls contains 111825 sequences and occupies a total of 70 Mbytes. In itself, then, these data can be accommodated in memory on any of the machines, so that the execution time of a *BLAST* call should be dominated by CPU processing. The CPU cost of a *BLAST* call for a given stored dataset varies significantly depending on the

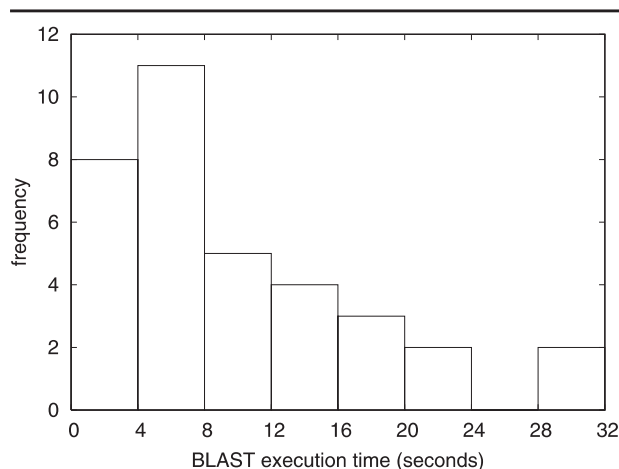


Fig. 9 The variation of the costs of the *BLAST* calls for the *ncl-pc1*.

input sequence. For the 35 sequences resulting from the join, this variation is shown in Figure 9. Since the CPU speeds vary, there is a range of options for increasing the parallelism in the *operation call* operator.

Two policies for parallelizing *operation calls* are considered here. In the first, *fastest-processor-first*, the parallelization of the *operation call* operator is increased by allocating an extra instance to the fastest processor not currently allocated one. In the second case, *slowest-processor-first*, each extra instance is allocated to the slowest processor not currently allocated one. Such a strategy may be adopted if the user is charged for using Grid resources, and the cost of a resource depends on its power.

Orthogonally to the two processor allocation policies, the data distribution may or may not take into account the different machine characteristics. Static load balancing is improved by skewing the *round-robin* distribution of tuples to the multiple *operation_call* operators, according to the speed of the relevant machine.

Figure 10 shows measurements of the performance of the query in Figure 11. The four plots correspond to the four different combinations of processor allocation and data distribution policies. In the *fastest-processor-first*, *man-pc2* is the first node allocated to evaluate the *BLAST*, *ncl-pc3-1* is the second, *ncl-pc3-2* is the third, *ncl-pc2* is the fourth, and *ncl-pc1* is added last. In the *slowest-processor-first* the order is reversed. Each measurement is taken five times, and the average of the last three has been used. Also, the measurements are taken during quiet

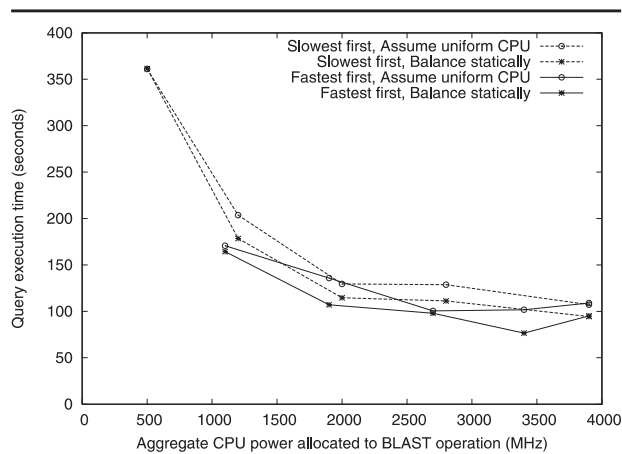


Fig. 10 Measured performance of the example query when the *BLAST* call is parallelized according to *fastest-processor-first* and *slowest-processor-first*.

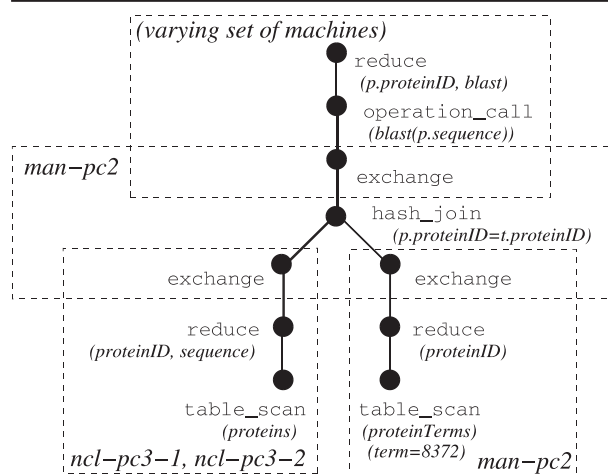


Fig. 11 The query plan used in the measurements. The degree of parallelism of the *operation call* is varying.

periods, at night and weekend, when activity on the network as well as the machines used is low.

The cost of the non-parallelized part of the query, i.e. the cost to retrieve the remote data and perform the join, is 17.6 s. This is relatively high for such volumes of data, and it is mainly due to the cost of sending the GIMS data from *ncl-pc3* to *man-pc2*. Some interesting observations can be made from these measurements. First, the benefits of the resource sharing outweigh the Grid overheads and the communication cost significantly. In

the case in which the fastest processor is chosen first, the first choice is *man-pc2*, which means that the *hash_join* and the *BLAST* call reside on the same processor. Further increases in the parallelism entail movement of data from Manchester to Newcastle. Even in that case, such increases yield better results. Also, for four processors, the aggregate CPU power is 3400 MHz. Assuming that the time of the parallelized part is inversely proportional to the total CPU power available, the performance of an ideal parallel computation would be $17.6 + (164.3 - 17.6)(1100/3400) = 65.1$. The fastest measured execution time is 76.6, suggesting the query processor has done a reasonable job at this point. The most powerful machine, *man-pc2*, would evaluate all the *BLAST* calls with no other processes running in no less than 147 s. It is the responsibility of the scheduler component of the query compiler to choose the best combination of machines and data distribution policy.

A second observation is that even totally static load balancing policies over highly skewed resources can improve the performance of the system. In both cases of processor allocation, assuming uniform CPU capacities leads to worse performance than taking into consideration a very simple resource characteristic such as the maximum CPU power. The difference between the Grid and other distributed systems is that the Grid provides the mechanisms to publish and advertise such characteristics, both static and dynamic.

These measurements barely scratch the surface of investigations into the performance of query processing in a Grid context, but they do hint at the complexity inherent in planning wide area distributed computations over heterogeneous resources, and demonstrate that even simple policies based on static measurements can be followed to arrive at sensible configurations. It is easy to see how the problem of planning this type of computation becomes more difficult in the presence of other usage, or if for instance the *BLAST* dataset were much larger, so that memory availability needs to be considered. In local and distributed database query processing, such considerations have driven the development of sophisticated optimization techniques (Kossmann, 2000). The argument here is that such techniques can be beneficial in a Grid setting and should be extended to take account of the metrics generated by a Grid infrastructure, including both static metrics such as CPU speed and memory size used here and dynamic metrics such as CPU and memory availability, in order to implement a manageable programming model for query style computations over a Grid.

6 Conclusions

One of the main hopes for the Grid is that it will encourage the publication of scientific and other data in a more

open manner than is currently the case. If this occurs then it is likely that some of the greatest advances will be made by combining data from separate, distributed sources to produce new results. The data that applications wish to combine will have been created by different researchers or organizations that will often have made local, independent decisions about both the best database paradigm and design for their data. The role of DQP in such a setting is to provide high-level, declarative facilities for describing requests over multiple data stores and analysis facilities.

The ease with which DQP allows such requests to be phrased has been illustrated through the example query in Section 2. This query is straightforward, but has illustrated how DQP can be used to provide optimization and evaluation of declarative requests over resources on the Grid. We note that the alternative of writing such a request using lower-level programming models, such as MPICH-G or a COG kit (von Laszewski et al., 2001) could be quite time-consuming. We note also that as the complexity of a request increases, it becomes increasingly difficult for a programmer to make decisions as to the most efficient way to express a request. Developing efficient execution plans for such tasks using existing Grid programming environments would take a skilled developer a significant time. We believe that DQP can serve an important role in Grid environments by:

1. increasing the variety of people who can form requests over multiple Grid resources;
2. reducing development times for certain categories of Grid programming task; and
3. enabling typical requests to be evaluated efficiently as a result of system-supported query optimization and support for implicit parallelism.

The Polar* prototype DQP system described in this paper is the result of our ongoing work on query processors for the Grid. The prototype has been implemented over Globus middleware using MPICH-G, and experiments have been conducted over bioinformatics databases and analysis tools at the authors' geographically remote sites. Future work will: (i) extend the range of physical operators in the algebra; (ii) increase the amount of system information used by the scheduler in query planning; (iii) explore the development of more powerful scheduling algorithms; and (iv) conduct performance evaluations over more and larger databases. We also work to evolve the Polar* system to be compliant with the emerging Open Grid Services Architecture (Foster et al., 2002), and to make use of standard service interfaces to databases (Watson, 2001) to reduce the cost of wrapper development.

ACKNOWLEDGEMENTS

This work was funded by the EPSRC Distributed Information Management Initiative and the UK e-Science Programme, whose support we are pleased to acknowledge. We are grateful for the contribution of Sandra Sampaio, whose work on Polar has been built upon extensively in Polar*. Also, we want to thank Dr. Mike Cornell and Dr. Phil Lord for helping with the bioinformatics.

AUTHORS BIOGRAPHIES

Jim Smith is a research associate working at the University of Newcastle upon Tyne on the EPSRC project High Performance Query Processing on the Grid.

Anastasios Gounaris (BSc, MPhil) is a research associate at the Information Management Group of the Computer Science Department of the University of Manchester. He is also pursuing a PhD degree at the same department. His research interests lie in the area of information management, Grid computing, distributed DBMSs and query processing. He is currently working on the Polar* project.

Paul Watson is Professor of Computing Science at the University of Newcastle, and Director of the North-East Regional e-Science Centre. He obtained his BSc (Computer Engineering, 1983) and PhD (Computer Science, 1986) from Manchester University. After time as a lecturer in the Manchester Computer Science Department he moved to ICL High Performance Systems Division where he worked on the design of parallel database systems from 1990 to 1995. In 1995 he joined Newcastle University where he has pursued his research interests in data-intensive parallel systems. This has included work on the Polar parallel object database server, in collaboration with Manchester University. In 1999, he became interested in the Grid and is currently working on a number of projects that are investigating how to integrate data and computation over a Grid. The technology being developed, including that described in this paper, is being exploited in e-science projects in a number of domains including bioinformatics, neuroinformatics and network analysis.

Norman W. Paton is a Professor of Computer Science at the University of Manchester. He obtained a BSc in Computing Science from Aberdeen University in 1986, and a PhD from the same institution in 1989. He was a lecturer in Computer Science at Heriot-Watt University, 1989–1995, before moving to Manchester, where he co-leads the Information Management Group. His research

interests have principally been in databases, in particular active databases, spatial databases, deductive object-oriented databases and user interfaces to databases. He is currently working on parallel object databases, spatio-temporal databases, distributed information systems, and on information management for bioinformatics. In terms of Grid Computing, he chairs the UK e-Science Programme's Database Task Force, and is co-Chair of the Database Access and Integration Services Working Group of the Global Grid Forum.

Alvaro A. A. Fernandes (BSc, MSc, PhD) is a Lecturer in the Department of Computer Science at the University of Manchester. From 1990 to 1995, at Heriot-Watt University in Edinburgh, he was a member of the team that designed and implemented the ROCK & ROLL deductive object-oriented database system. He also participated in the work that extended ROCK & ROLL with active rules and with spatial capabilities. In 1996 he was appointed to a lectureship at Goldsmiths College (University of London), where he worked on personalization and adaptivity mechanisms. In 1998 he joined the Information Management Group in Manchester. Since then he has retained an interest in object databases, spatial and temporal capabilities and on adaption mechanisms. He has worked on MOVIE (an incremental view maintenance system for object databases), on Tripod (an ODMG-compliant spatio-historical database system), and on Polar*. His principal current interests are in distributed query processing in computational Grids, spatio-temporal databases and logic-based approaches to knowledge discovery and their application to knowledge management.

Rizos Sakellariou received his PhD degree, in computer science, from the University of Manchester in 1997 for a thesis on symbolic analysis techniques with applications to loop partitioning and scheduling for parallel computers. Since January 2000, he has been a Lecturer in Computer Science at the University of Manchester. Prior to his current appointment, he was a visiting Assistant Professor at the University of Cyprus (fall 1999), and a postdoctoral research associate at Rice University (1998–1999) and the University of Manchester (1996–1998). His research interests are in the areas of parallel and distributed systems, performance evaluation and modeling, optimizing and parallelizing compilers, and software for Grid and Internet computing.

NOTES

1. For what it is worth, the GO term GO:0008372 is used for "the annotation of gene products whose localization is not known or cannot be inferred." (See www.godatabase.org.)

2. OIF is a standard textual representation for ODMG objects.

REFERENCES

- Altschul, S. F., Gish, W., Miller, W., Myers, E.W., and Lipman, D. J., 1990. Basic local alignment search tool. *Journal of Molecular Biology* 215: 403–410.
- Ashburner, M., Ball, C., Blake, J., Botstein, D., Butler, H., Cherry, J., Davis, A., Dolinski, K., Dwight, S., Eppig, J., Harris, M., Hill, D., Issel-Tarver, L., Kasarskis, A., Lewis, S., Matese, J., Richardson, J., Ringwald, M., Rubin, G., and Sherlock, G., 2000. Gene ontology: tool for the unification of biology. *Nature Genetics* 25(1): 25–29.
- Bairoch, A., and Apweiler, R., 2000. The swiss-prot protein sequence database and its supplement trembl in 2000. *Nucleic Acids Research* 28: 45–48.
- Carey, M., DeWitt, D.J., Franklin, M., Hall, N., McAuliffe, M., Naughton, J., Schuh, D., Solomon, M., Tan, C., Tsatalos, O., White, S., and Zwilling, M., 1994. Shoring up persistent applications. In R. Snodgrass and M. Winslett, editors, *Proceedings of the 1994 ACM SIGMOD International Conference on Management of Data*, Minneapolis, MN, 24–27 May 1994, ACM Press, New York, pp. 383–394.
- Cattell, R.G.G., and Barry, D.K., 2000. *The Object Database Standard: ODMG 3.0*, Morgan Kaufmann, San Mateo, CA.
- Cornell, M., Paton, N.W., Wu, S., Goble, C.A., Miller, C.J., Kirby, P., Eilbeck, K., Brass, A., Hayes, A., and Oliver, S.G., 2001. GIMS – a data warehouse for storage and analysis of genome sequence and functional data. In *Proceedings of the 2nd IEEE Symposium on Bioinformatics and Bioengineering (BIBE)*, IEEE Press, Piscataway, NJ, pp. 15–22.
- Davidson, S.B., Crabtree, J., Brunk, B.P., Schug, J., Tannen, V., Overton, G.C., and Stoeckert, C.J., 2001. K2/kleisli and gus: Experiments in integrated access to genomic data sources. *IBM Systems Journal* 40(2): 512–531.
- DeWitt, D., and Gray, J., 1992. Parallel database systems: The future of high performance database systems. *Communication of ACM (CACM)* 35(6): 85–98.
- Dinda, P., and Plale, B., 2001. A unified relational approach to Grid information services. *Technical Report GWD-GIS-012-1*, Global Grid Forum.
- Fegaras, L., and Maier, D., 2000. Optimizing object queries using an effective calculus. *ACM Transactions on Database Systems* 24(4): 457–516.
- Foster, I., and Karonis, N.T., 1998. A Grid-enabled MPI: Message passing in heterogeneous distributed computing systems. In *Proceedings of Supercomputing (SC)*, IEEE Computer Society, Los Alamitos, CA. Online at www.supercomp.org/sc98/proceedings/.
- Foster, I., Kesselman, C., Nick, J., and Tuecke, S., 2002. Grid services for distributed system integration. *IEEE Computer* 35: 37–46.
- Garcia-Molina, H., Papakonstantinou, Y., Quass, D., Rajaraman, A., Sagiv, Y., Ullman, J., Vassalos, V., and Widom, J., 1997. The TSIMMIS approach to mediation: data models and languages. *Journal of Intelligent Information Systems* 8(2): 117–132.
- Graefe, G., 1990. Encapsulation of parallelism in the Volcano query processing system. In *ACM SIGMOD*, pp. 102–111.
- Graefe, G., 1993. Query evaluation techniques for large databases. *ACM Computing Surveys* 25(2):73–170.
- Haas, L., Kossmann, D., Wimmers, E.L., and Yang, J., 1997. Optimizing queries across diverse data sources. In *Proceedings of the VLDB*, Morgan-Kaufmann, San Mateo, CA, pp. 276–285.
- Hasan, W., and Motwani, R., 1995. Coloring away communication in parallel query optimization. In Umeshwar Dayal, Peter M. D. Gray, and Shojiro Nishio, editors, *VLDB '95, Proceedings of 21th International Conference on Very Large Data Bases*, 11–15 September 1995, Zurich, Switzerland, Morgan Kaufmann, San Mateo, CA, pp. 239–250.
- Hasan, W., Florescu, D., and Valduriez, P., 1996. Open issues in parallel query optimization. *SIGMOD Record* 25(3): 28–33.
- Hsiao, D., 1992. Tutorial on federated databases and systems. *The VLDB Journal* 1(1): 127–179.
- Kossmann, D., 2000. The State of the Art in Distributed Query Processing. *ACM Computing Surveys* 32(4): 422–469.
- Ozsu, M.T., and Valduriez, P., editors, 1999. *Principles of Distributed Database Systems*, 2nd edition, Prentice-Hall, Englewood Cliffs, NJ.
- Rahm, E., and Marek, R., 1995. Dynamic multi-resource load balancing in parallel database systems. In *Proceedings of the 21st VLDB Conference*, pp. 395–406.
- Sampaio, S.F.M., Paton, N.W., Watson, P., and Smith, J., 1999. A parallel algebra for object databases. In *Proceedings of the 10th International Workshop on Database and Expert Systems Applications*, Florence, Italy, September, IEEE Computer Society, Los Alamitos, CA, pp. 56–60.
- Sampaio, S.F.M., Smith, J., Paton, N.W., and Watson, P., 2001. An experimental performance evaluation of join algorithms for parallel object databases. In R. Sakellariou et al., editors, *Proceedings of the 7th International Euro-Par Conference*, Springer-Verlag, Berlin, pp. 280–291.
- Smith, J., Sampaio, S.F.M., Watson, P., and Paton, N.W., 2000. Polar: An architecture for a parallel ODMG compliant object database. In *Conference on Information and Knowledge Management (CIKM)*, ACM press, New York, pp. 352–359.
- Snir, M., Otto, S., Huss-Lederman, S., Walker, D., and Dongarra, J., 1998. *MPI – The Complete Reference*, MIT Press, Cambridge, MA.
- Szalay, A., Kunszt, P.Z., Thakar, A., Gray, J., and Slut, D.R., 2000. Designing and mining multi-terabyte astronomy archives: The SLOAN digital sky survey. In *Proceedings of ACM SIGMOD*, ACM Press, New York, pp. 451–462.
- von Laszewski, G., Foster, I., Gawor, J., and Lane, P., 2001. A Java commodity Grid kit. *Concurrency and Computation: Practice and Experience* 13(8–9):643–662.
- Watson, P., 2001. Databases and the Grid. *Technical Report CS-TR-755*, University of Newcastle.