We are IntechOpen, the world's leading publisher of Open Access books Built by scientists, for scientists

6,300

Open access books available

172,000

International authors and editors

190M

Downloads

Our authors are among the

154

Countries delivered to

TOP 1%

most cited scientists

12.2%

Contributors from top 500 universities

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

Toward Optimal Query Execution in Data Grids

Reza Ghaemi¹, Amin Milani Fard², Md. Nasir Bin Sulaiman³
and Hamid Tabatabaee⁴

¹ Academic Member in Department of Computer Engineering, Islamic Azad University,
Quchan branch, Iran and PhD Student in Faculty of Computer Science
and Information Technology Putra University of Malaysia,
E-mail : rezaghaemi@ieee.org

² School of Computing Sciences, Simon Fraser University, BC, Canada,
E-mail: milanifard@cs.sfu.ca

³ Associate Professor and Academic Member in Department of Computer Science, Faculty of
Computer Science and Information Technology Putra University of Malaysia,
E-mail : nasir@fktm.upm.edu.my

⁴ Academic Member in Department of Computer Engineering, Islamic Azad University,
Quchan branch, Iran, and PhD Student in Ferdowsi University of Iran,
E-mail : hamid.tabatabaee@gmail.com.

1. Introduction

Nowadays Grid technology [Foster & Kesselman, 1999] provides us with simultaneous and effective use of distributed computational and informational resources. Three main types of this technological phenomenon are known as resource discovery grids, computational grids, and data grids. In data grids, distributed heterogeneous database systems play important role to provide users access information easily without knowing the resource position. Due to heterogeneity property, communication among subsystems has to be handled well considering different network structure, operating system, and DBMS. In such systems we are also interested in efficient search and retrieval mechanism to speedup traditional relational database queries.

Distributed systems can be thought of as a partnership among independent cooperating centralized systems. Upon this idea number of large scale applications has been investigated during past decades among which distributed information retrieval (DIR) systems gained more popularity due to its high demand. The goal of DIR is to provide a single search interface that provides access to the available databases involving resource descriptions building for each database, choosing which databases to search for particular information, and merging retrieved results into a single result list [Si & Callan, 2003].

A distributed database (DDB) is a collection of multiple, logically interrelated databases distributed over a computer network. This resource distribution improves performance, reliability, availability and modularity that are inherent in distributed systems. As with traditional centralized databases, distributed database systems (DDBS) must provide an

efficient user interface that hides all of the underlying data distribution details of the DDB from the users. The use of a *relational query* allows the user to specify a description of the data that is required without having to know where the data is physically located [Li & Victor, 1981].

Data retrieval from different sites in a DDB is known as *distributed query processing* (DQP). For example, the following query accesses data from the local database as well as the remote sales database. The first table (EMP) found in site1 and the second table (DEPT) found in site2:

```
SELECT  ename, dname
FROM    company.emp e, company.dept@sales.goods d
WHERE   e.deptno = d.deptno
```

So a *distributed query* is one that selects data from databases located at multiple sites in a network and *distributed processing* performs computations on multiple CPUs to achieve a single result. Query processing is much more difficult in distributed environment than in centralized environment because a large number of parameters affect the performance of distributed queries, relations may be fragmented and/or replicated, and considering many sites to access, query response time may become very high [Li & Victor, 1981]. It is quite evident that the performance of a DDBS is critically dependant upon the ability of the query optimization algorithm to derive efficient query processing strategies. DDBMS query optimization algorithms attempts to reduce the quantity of data transferred. Minimizing the quantity of data transferred is a desirable optimization criterion. The distributed query optimization has several problems relate to the cost model, larger set of queries, optimization cost, and optimization interval.

The goal of DQP is to execute such queries as efficiently as possible in order to minimize the response time that users must wait for answers or the time application programs are delayed. And to minimize the total communication costs associated with a query, to improved throughput via parallel processing, sharing of data and equipment, and modular expansion of data management capacity. In addition, when redundant data is maintained, one also achieves increased data reliability and improved response time.

In this work we propose a multi-agent architecture for distributed query processing. The structure of the paper is as follows. Section II describes an overview of query optimization process. An investigation on related works is declared in section III, our proposed approach in section IV, and simulation results in V. We finally concluded the work in the section VI.

2. Query Optimization Process

In a relational database all information can be found in a series of tables. A query therefore consists of operations on tables. The most common queries are Select-Project-Join queries. In this paper, we will focus on the join-ordering problem since permutations of the join order have the most important effect on performance of relational queries [Özsu & Valduriez, 1999]. The query optimization process shown in Figure 1, consists of getting a query on n relations and generating the best Query Execution Plan (QEP).

Fig. 1. Query optimization process

For a given query, the *search space* can be defined as the set of equivalent operator trees that can be produced using transformation rules. The example below illustrates 3 equivalent join trees, which are obtained by exploiting the associative property of binary operators. Join tree (c) which starts with a Cartesian product may have a much higher cost than other join trees.

```

SELECT  ENAME, RESP
FROM    EMP, ASG, PROJ
WHERE   EMP.ENO=ASG.ENO
AND     ASG.PNO=PROJ.PNO
  
```

Fig. 2. Query equivalent trees

Regarding different search spaces, there would be different shape of the join tree. In a *linear tree*, at least one operand of each operand node is a base relation. However, a *bushy tree* might have operators whose both operands are intermediate operators. In a distributed environment, bushy trees are useful in exhibiting parallelism [Özsu & Valduriez, 1999].

Fig. 3. Linear vs. bushy join tree

Considering new large scale database applications such as deductive database systems and bioinformatics, it is necessary to be able to deal with larger size queries. The search complexity constantly increases and makes higher demand for better algorithms than our traditional relational database queries.

3. Related Works

Three most common types of algorithms for join-ordering optimization are deterministic, Genetic and randomized algorithms [Zalenay, 2005].

Deterministic algorithm, also known as exhaustive search *dynamic programming* algorithm, produces optimal left-deep processing trees with the big disadvantage of having an exponential running time. This means that for queries with more than 10-15 joins, the running time and space complexity explodes [Zalenay, 2005]. Due to the very large time and space complexity of this algorithm for plan enumeration, iterative dynamic programming approach was proposed which produces reasonable plans with reasonable running time for most network topologies. However, its complexity is not much more than classical DP algorithm.

Genetic and randomized algorithms [Ioannidis & Kang, 2003; Steinbrunn et al., 1997] on the other hand do not generally produce an optimal access plan. But in exchange they are superior to dynamic programming in terms of running time. Experiments have shown that it is possible to reach very similar results with both genetic and randomized algorithms depending on the chosen parameters. Still, the genetic algorithm has in some cases proved to be slightly superior to randomized algorithms.

Layers of distributed query optimization have been depicted in Figure 4.

There are number of Query Execution Plan for DDB such as: row blocking, multi-cast optimization, multi-threaded execution, joins with horizontal partitioning, Semi Joins, and Top n queries [Kossman, 2000; Selinger et al., 1998; Stocker et al., 2001]. In this paper we propose a novel agent based QEP generator for heterogeneous distributed data base systems.

Fig. 4. Distributed query optimization

3.1 Distributed Cost Model

An optimizer cost model includes cost functions to predict the cost of operators, and formulas to evaluate the sizes of results. Cost functions can be expressed with respect to either the total time, or the response time [Ceri & Pelagatti, 1984; Özsu & Valduriez, 1999]. The total time is the sum of all times and the response time is the elapsed time from the initiation to the completion of the query. The total time (TT) is computed as bellow, where T_{CPU} is the time of a CPU instruction, $T_{I/O}$ the time of a disk I/O, T_{MSG} the fixed time of

initiating and receiving a message, and T_{TR} the time it takes to transmit a data unit from one site to another.

$$TT = T_{CPU} * \#_{insts} + T_{I/O} * \#_{I/Os} + T_{MSG} * \#_{msgs} + T_{TR} * \#_{bytes}$$

When the response time of the query is the objective function of the optimizer, parallel local processing and parallel communications must also be considered. This response time (RT) is calculated as bellow:

$$RT = T_{CPU} * seq_ \#_{insts} + T_{I/O} * seq_ \#_{I/Os} + T_{MSG} * seq_ \#_{msgs} + T_{TR} * seq_ \#_{bytes}$$

Most early distributed DBMSs designed for wide area networks have ignored the local processing cost and concentrate on minimizing the communication cost. Consider the following example:

$$TT = 2 * T_{MSG} + T_{TR} * (x + y)$$

$$RT = \max \{T_{MSG} + T_{TR} * x, T_{MSG} + T_{TR} * y\}$$

In parallel transferring, response time is minimized by increasing the degree of parallel execution. This does not imply that the total time is also minimized. On contrary, it can increase the total time, for example by having more parallel local processing (often includes synchronization overhead) and transmissions. Minimizing the total time implies that the utilization of the resources improves, thus increasing the system throughput. In practice, a compromise between the total and response times is desired.

3.2 Database Statistics

The main factor affecting the performance is the size of the intermediate relations that are produced during the execution. When a subsequent operation is located at a different site, the intermediate relation must be transmitted over the network. It is of prime interest to estimate the size of the intermediate results in order to minimize the size of data transfers. The estimation is based on statistical information about the base relations and formulas to predict the cardinalities of the results of the relational operations.

Cartesian product: The cardinality of the Cartesian product of R and S is as (1):

$$card(R \times S) = card(R) \times card(S) \quad (1)$$

Join: There is no general way to estimate the cardinality of a join without additional information. The upper bound of the join cardinality is the cardinality of the Cartesian

product. Some systems, such as Distributed INGRES [Stonebraker, 1986], use this upper bound, which is quite pessimistic. R^* [Selinger & Adiba, 1980] uses this upper bound divided by a constant to reflect the fact that the join result is smaller than the Cartesian product. However, there is a case, which occurs frequently, where the estimation is simple. If relation R is equi-join with S over attribute A from R , and B from S , where A is a key of relation R , and B is a foreign key of relation S , the cardinality of the result can be approximated as (2):

$$card(R \bowtie_{A=B} S) = card(S) \quad (2)$$

In other words, the Cartesian product $R * S$ contains $n_r * n_s$ tuples; each tuple occupies $s_r + s_s$ bytes. If $R \cap S = \emptyset$, then $R \bowtie S$ is the same as $R * S$. If $R \cap S$ is a key for R , then a tuple of s will join with at most one tuple from R , therefore, the number of tuples in $R \bowtie S$ is no greater than the number of tuples in S . If $R \cap S$ in S is a foreign key in S referencing R , then the number of tuples in $R \bowtie S$ is exactly the same as the number of tuples in s . The case for $R \cap S$ being a foreign key referencing S is symmetric.

As discussed earlier, ordering joins is an important aspect of centralized query optimization. This matter in a distributed context is even more important since joins between fragments may increase the communication time. Two main approaches exist to order joins in fragment queries:

- 1) Direct optimization of the ordering of joins (e.g. in the Distributed INGRES algorithm).
- 2) Replacement of joins by combination of *semi-joins* in order to minimize communication costs.

Let R and S are relations stored at different sites and query $R \bowtie S$ be the join operator. The obvious choice is to send the smaller relation to the site of the larger one.

More interesting is the case where there are more than two relations to join. The objective of the join ordering algorithm is to transmit smaller operands. Since the join operations may reduce or increase the size of intermediate results, estimating the size of joint results is mandatory, but difficult. Consider the following query expressed in relate.

algorithm: $PROJ \bowtie_{PNO} EMP \bowtie_{ENO} ASG$ whose join graph is below:

This query can be executed in at least 5 different ways.

1. $EMP \rightarrow ASG ; EMP' = EMP \bowtie ASG \rightarrow PROJ ; EMP' \bowtie PROJ$
2. $ASG \rightarrow EMP ; EMP' = EMP \bowtie ASG \rightarrow PROJ ; EMP' \bowtie PROJ$
3. $ASG \rightarrow PROJ ; ASG' = ASG \bowtie PROJ \rightarrow EMP ; ASG' \bowtie EMP$
4. $PROJ \rightarrow ASG ; PROJ' = PROJ \bowtie ASG \rightarrow EMP ; PROJ' \bowtie EMP$
5. $EMP \rightarrow ASG ; PROJ \rightarrow ASG ; EMP \bowtie PROJ \bowtie ASG$

To select one of these programs, the following sizes must be known or predicted: $size(EMP)$, $size(ASG)$, $size(PROJ)$, $size(EMP \bowtie ASG)$, $size(ASG \bowtie PROJ)$. Furthermore, if it is the response time that is being considered, the optimization must take into account the fact that transfers can be done in parallel with strategy 5. An alternative to enumerating all the solutions is to use heuristics that consider only the size of the operand relations by assuming, for example, that the cardinality of the resulting join is the product of cardinalities. In this case, relations are ordered by increasing sizes and the order of execution is given by this ordering and the join graph. For instance, the order (EMP, ASG, PROJ) could use strategy 1, while the order (PROJ, ASG, EMP) could use strategy 4.

3.3 Multi Agent Systems

Multi-agent systems (MASs) as an emerging sub-field of artificial intelligence concern with interaction of agents to solve a common problem [Wooldridge, 2002]. This paradigm has become more and more important in many aspects of computer science by introducing the issues of distributed intelligence and interaction. They represent a new way of analyzing, designing, and implementing complex software systems.

In multi-agent systems, communication is the basis for interactions and social organizations which enables the agents to cooperate and coordinate their actions. A number of communication languages have been developed for inter-agent communication, in which the most widely used ones are KIF (Knowledge Interchange Format) [Genesereth & Fikes, 1992], KQML (Knowledge Query and Manipulation Language) [Finin et al., 1994], and ACL (Agent Communication Language) [Labrou et al., 1999]. KQML uses KIF to express the content of a message based on the first-order logic. KIF is a language intended primarily to express the content part of KQML messages. ACL is another communication standard emerging in competition with KQML since 1995. Nowadays, XML (Extensible Markup Language) started to show its performance as a language to encode the messages exchanged between the agents, in particular in agent-based e-commerce to support the next generation of Internet commerce [Korzyk, 2000].

4. Proposed System Architecture

Although the problem of distributed query processing in heterogeneous systems has been investigated before [Huang et al., 1981], a good solution to practical query optimization has not been studied well. To meet so we proposed new multi-agent system architecture based on *Java Agent DEvelopment* (JADE) framework [Bellifemine et al., 2006]. JADE is a software development framework aimed at developing multi-agent systems and applications in

which agents communicate using FIPA¹ Agent Communication Language (ACL) messages and live in containers which may be distributed to several different machines. The Agent Management System (AMS) is the agent who exerts supervisory control over access to and use of the Agent Platform. Only one AMS will exist in a single platform.

Each agent must register with an AMS in order to get a valid AID. The Directory Facilitator (DF) is the agent who provides the default yellow page service in the platform. The Message Transport System, also called Agent Communication Channel (ACC), is the software component controlling all the exchange of messages within the platform, including messages to/from remote platforms.

JADE is capable of linking Web services and agents together to enable semantic web applications. A Web service can be published as a JADE agent service and an agent service can be symmetrically published as a Web service endpoint. Invoking a Web service is just like invoking a normal agent service. Web services' clients can also search for and invoke agent services hosted within JADE containers.

The Web Services Integration Gateway (WSIG) [JADE board, 2005] uses a Gateway agent to control the gateway from within a JADE container. Interaction among agents on different platforms is achieved through the Agent Communication Channel. Whenever a JADE agent sends a message and the receiver lives on a different agent platform, a Message Transport Protocol (MTP) is used to implement lower level message delivery procedures [Cortese et al., 2002]. Currently there are two main MTPs to support this inter-platform agent communication - CORBA IIOP-based and HTTP-based MTP.

Considering large-scale applications over separated networks, agent communications has to be handled behind firewalls and Network Address Translators (NATs), however, the current JADE MTP do not allow agent communication through firewalls and NATs. Fortunately, the firewall/NAT issue can be solved by using the current JXTA implementation for agent communication [Liu et al., 2006].

JXTA is a set of open protocols for P2P networking. These protocols enable developers to build and deploy P2P applications through a unified medium [Gradecki, 2002]. Obviously, JXTA is a suitable architecture for implementing MTP-s for JADE and consequently JADE agent communication within different networks can be facilitated by incorporating JXTA technology into JADE [Liu et al., 2006].

In this work, a query is submitted by user to the Query Distributor Agent and then it will be distributed using the submitted terms in the available ontology. Our proposed architecture uses different types of agents, each having its own characteristics. The proposed system architecture is shown in Figure 5.

¹ Foundation for Intelligent Physical Agents (<http://www.fipa.org>)

Fig. 5. Proposed system architecture

Query Distributor Agent (QDA): After receiving the user query, the QDA sends sub-queries to responsible local optimizer agents. The QDA can also create search agents if needed.

Local Optimizer Agents (LOAs): These agents apply a Genetic algorithm based sub-query optimization which will be discussed later, and return a result table size to the Global Optimizer Agent.

Global Optimizer Agent (GOA): This agent has the responsibility to find best join order via network. To do so, GOA receives result table size information from LOAs and again using an evolutionary method finds a semi-optimal join order, however, this time the genetic algorithm fitness function is based on minimizing communication rate among different sites. The final result would be then send to LOAs to perform operation and deliver result to the GOA to show them on the screen.

4.1 Genetic Algorithm based Optimization

The basic concept of the genetic algorithms were originally developed by Holland [Holland, 1975] and later revised by Goldberg [Goldberg, 1989]. Goldberg showed that genetic algorithms are independent of any assumption about the search space and are based on the mechanism of natural genetics. The first step to model this problem as a Genetic Algorithm problem, is determining the *chromosome*, *Genetic A operators*, and *fitness function*.

Chromosome Design

In order to encode different access plans, we considered *ordered list scheme* to represent each processing tree as an ordered list. For instance $((((R1 \bowtie R5) \bowtie R3) \bowtie R2) \bowtie R4)$ is encoded as "15324". This converts the join order to the well-known traveling salesman problem (TSP). Possible query plans are encoded as integer strings. Each string represents the join order from one relation of the query to the next. This mechanism is also used within the PostgreSQL optimizer [PostgreSQL]. There may be other ways to encode processing trees especially bushy trees; however, we have used the above explained ordered list encoding method for implementation simplicity and faster run.

Genetic Algorithm operations

For the crossover, one point in the selected chromosome would be selected along with a corresponding point in another chromosome and then the tails would be exchanged. Mutation processes causes some bits to invert and produces some new information. The only problem of mutation is that it may cause some useful information to be corrupted. Therefore we used elitism which means the best individual will go forward to the next generation without undergoing any change to keep the best information.

Fitness function

Defining fitness function is one of the most important steps in designing a Genetic Algorithm based method, which can guide the search toward the best solution. In our simulation we used a simple random cost estimator as defined below in (3), where $\text{random}(x)$ returns an integer between 0 and x .

$$fitness = \begin{cases} \text{random}(R \times S) + (R + S) & ; \text{if } (R + S) > (R \times S) \\ \text{random}(R + S) + (R \times S) & ; \text{else} \end{cases} \quad (3)$$

Algorithm design

After calculating the fitness function value for each parent chromosome the algorithm will generate N children. The lower a parent chromosome's fitness function value is the higher probability it has to contribute one or more offspring in the next generation. After performing operations, some chromosomes might not satisfy the fitness and as a result the algorithm discards this process and gets M ($M \leq N$) children chromosomes. The algorithm then selects N chromosomes with the lower fitness value from the $M + N$ chromosomes (M children and N parents) to be parent of the next generations. This process would repeat until a certain number of generations are processed, after which the best chromosome is chosen. Figure 6 shows our genetic algorithm based approach.

Fig. 6. Our Genetic Algorithm based query optimization algorithm

5. Results

Experiments were done on a PC Pentium 4 CPU 2 GHz and 1GB RAM. The evolutionary process accomplished in less than a second and seen in Figure 7 a sample query for 20 joins is converged to a near optimal fitness almost after 100 generations. Table 1 shows the parameters value we set for our implementation.

Parameter	Value
Population size	100
Mutation probability	0.02
Crossover probability	0.7
Elitism probability	0.5
Number of generations	100

Table 1. Parameteres settings for Genetic Algorithm based Approach

Fig. 7. Optimization of a 20-joins query in 100 generations

There is no doubt that dynamic programming method always gives us optimal solution, however, since the time and space complexity of the genetic algorithm base optimization is much less, it is not a practical approach for high amount of nested joins. Figure 8 shows T_{GA} / T_{DP} which is rate of average run time for 10 queries, where T_{GA} is genetic algorithm based optimization average run time and T_{DP} is dynamic programming based optimization average run time for the same query. Results as expected shows this rate is always less than 1 which means the genetic algorithm approach has a less overhead in contrast with the DP method.

Fig. 8. Rate of optimization average run time for 10 queries

A computational complexity analysis in table 2 shows the superiority of the evolutionary method in comparison with the dynamic programming approach. The time and space complexity in the evolutionary method is liner with respect to G as the number of generations, and N as the number of chromosomes in the population, while it is exponential in the DP method which is not efficient in handling more than 10 joins.

	<i>Space Complexity</i>	<i>Time Complexity</i>
<i>Evolutionary Method</i>	$O(N)$	$O(G.N)$
<i>Dynamic Programming</i>	$O(2^n)$	$O(2^n)$

Table 2. Computational Complexity Comparison

6. Conclusion

An evolutionary query optimization mechanism in distributed heterogeneous systems has boon proposed using multi-agent architecture and genetic algorithm approach. Although deterministic dynamic programming algorithm produces optimal left-deep processing trees, it has the big disadvantage of having an exponential running time. Genetic and randomized algorithms on the other hand do not generally produce an optimal access plan. But in exchange they are superior to dynamic programming in terms of running time. Our practical framework uses hybrid JADE-JXTA framework which allows agent communication through firewalls and NATs in heterogeneous networks.

7. References

- Bellifemine, F.; Caire, G.; Trucco, T. And Rimassa, G. (2006). *JADE Programmer's Guide*, August 2006, published by the Free Software Foundation, USA
- Ceri, S.; Pelagatti, G. (1984). *Distributed Databases: Principles and Systems*, Mcgraw-Hill, ISBN-10: 0070108293, ISBN-13: 978-0070108295, New York, NY, USA
- Cortese, E.; Quarta, F.; Vitaglione, G. and Vrba, P. (2002). Scalability and Performance of the JADE Message Transport System, *Analysis of Suitability for Holonic Manufacturing Systems*, exp, 2002
- Finin, T.; Fritzon, R.; McKay, D. and McEntire, R. (1994). KQML as an Agent Communication Language, *Proceedings of the 3rd International Conference on Information and Knowledge Management (CIKM'94)*, editor N. Adam, B. Bhargava, Y. Yesha, pp. 456-463, ACM Press, Gaithersburg, MD, USA
- Foster, I.; Kesselman, C. (1999). *The Grid: Blueprint for a Future Computing Inf.*, Morgan Kaufmann Publishers, pp. 105-129
- Genesereth, M. R.; Fikes, R. E. (1992). Knowledge interchange format, version 3.0. Technical Report 92-1, Stanford University, Computer Science Department, San Francisco, USA
- Goldberg, D. E. (1989). *The genetic algorithms in search, optimization, and machine learning*, Addison-Wesley, New York, USA
- Gradecki, J. D. (2002). *Mastering JXTA: Building Java Peer-to-Peer Applications*, JohnWiley & Sons, ISBN 0471250848, 9780471250845, Canada

- Holland, J. H. (1975). *Adaptation in natural and artificial systems*, Ann Arbor, MI University of Michigan press, USA
- Huang; Kuan-Tsae; Davenport and Wilbur, B. (1981). Query processing in distributed heterogeneous databases, MIT. Laboratory for Information and Decision Systems Series/Report no.: LIDS-P ;1212
- Ioannidis Y. E.; Kang Y. C. (2003). Randomized Algorithms for Optimizing Large Join Queries, *Conference on Very large data bases*, pp. 488-499, September 2003, Berlin, Germany, ACM Publisher , New York, NY, USA
- JADE Board (2005). JADE WSIG Add-On Guide JADE Web Services Integration Gateway (WSIG) Guide, published by the Free Software Foundation, March 2005, Boston, MA, USA
- Korzyk, A. (2000). Towards XML As A Secure Intelligent Agent Communication Language, *the 23rd National Information Systems Security Conference*, Baltimore Convention Center, Baltimore, October 2000, Maryland, SA
- Kossman, D. (2000). The state of the art in distributed query processing, *ACM Computing Surveys*, Vol. 32, Issue 4, December 2000, pp. 422 – 469, ISSN:0360-0300
- Labrou, Y.; Finin, T. and Peng, Y. (1999). The current landscape of Agent Communication Languages, *The current landscape of Agent Communication Languages*, Intelligent Systems, Vol. 14, No. 2, March/ April 1999, IEEE Computer Society
- Li; Victor, O. K. (1981). Query processing in distributed data bases, MIT. Lab. for Information and Decision Systems Series/Report no.: LIDS-P, 1107
- Liu, S.; Küngas, P. and Matskin, M. (2006). Agent-Based Web Service Composition with JADE and JXTA, *Proceedings of the International Conference on Semantic Web and Web Services (SWWS 2006)*, June 2006, Las Vegas, Nevada, USA
- Özsu, T. M.; Valduriez, P. (1999). *Principles of Distributed Database Systems*, Second Edition, Prentice Hall, ISBN 0-13-659707-6
- PostgreSQL 8.3.0 Documentation, Chapter 49. Genetic Query Optimizer <http://www.postgresql.org/docs/8.3/interactive/geqo.html>
- Selinger, P. G.; Adiba, M. E. (1980). Access Path Selection in Distributed Database Management Systems, *ICOD*, pp. 204-215,
- Selinger, P. G.; Astrahan, M. M.; Chamberlin, D. D.; Lorie, R. A. and Price, T. G. (1998). *Access path selection in a relational database management system*, Morgan Kaufmann Series In Data Management Systems, Readings in database systems (3rd ed.), pp. 141 – 152, ISBN:1-55860-523-1
- Si, L.; Callan, J. (2003). A Semi supervised Learning Method to Merge Search Engine Results, *Proceeding ACM Transactions on Information Systems*, Vol. 21, No. 4, October 2003, pp. 457-491
- Steinbrunn, M.; Moerkotte, G. & Kemper, A. (1997). Heuristic and Randomized Optimization for the Join Ordering Problem, *The VLDB Journal - The International Journal on Very Large Data Bases*, Vol. 6 , Issue 3, August 1997, pp. 191-208, ISSN:1066-8888
- Stocker; Kossman; Braumandl and Kemper (2001). Integrating semi join reducers into state of the art query processors, *ICDE*, ISSN:1063-6382, IEEE Computer Society Washington, DC, USA

- Stonebraker (1986). The Design and Implementation of Distributed INGRES (context), Addison-Wesley Longman Publishing Co., Inc. Boston, MA, USA, ISBN:0-201-07185-1
- Wooldridge, M. (2002). *An Introduction to Multi agent Systems*, Published by John Wiley & Sons., February 2002, ISBN 0 47149691X
- Zelenay, K. (2005). Query Optimization, *Seminar Algorithmen für Datenbanksysteme*, June 2005, ETH Zürich

Advanced Technologies

Edited by Kankesu Jayanthakumaran

ISBN 978-953-307-009-4

Hard cover, 698 pages

Publisher InTech

Published online 01, October, 2009

Published in print edition October, 2009

This book, edited by the Intech committee, combines several hotly debated topics in science, engineering, medicine, information technology, environment, economics and management, and provides a scholarly contribution to its further development. In view of the topical importance of, and the great emphasis placed by the emerging needs of the changing world, it was decided to have this special book publication comprise thirty six chapters which focus on multi-disciplinary and inter-disciplinary topics. The inter-disciplinary works were limited in their capacity so a more coherent and constructive alternative was needed. Our expectation is that this book will help fill this gap because it has crossed the disciplinary divide to incorporate contributions from scientists and other specialists. The Intech committee hopes that its book chapters, journal articles, and other activities will help increase knowledge across disciplines and around the world. To that end the committee invites readers to contribute ideas on how best this objective could be accomplished.

How to reference

In order to correctly reference this scholarly work, feel free to copy and paste the following:

Reza Ghaemi, Amin Milani Fard, Nasir Bin Sulaiman and Hamid Tabatabaee (2009). Toward Optimal Query Execution in Data Grids, *Advanced Technologies*, Kankesu Jayanthakumaran (Ed.), ISBN: 978-953-307-009-4, InTech, Available from: <http://www.intechopen.com/books/advanced-technologies/toward-optimal-query-execution-in-data-grids>

INTECH
open science | open minds

InTech Europe

University Campus STeP Ri
Slavka Krautzeka 83/A
51000 Rijeka, Croatia
Phone: +385 (51) 770 447
Fax: +385 (51) 686 166
www.intechopen.com

InTech China

Unit 405, Office Block, Hotel Equatorial Shanghai
No.65, Yan An Road (West), Shanghai, 200040, China
中国上海市延安西路65号上海国际贵都大饭店办公楼405单元
Phone: +86-21-62489820
Fax: +86-21-62489821

© 2009 The Author(s). Licensee IntechOpen. This chapter is distributed under the terms of the [Creative Commons Attribution-NonCommercial-ShareAlike-3.0 License](https://creativecommons.org/licenses/by-nc-sa/3.0/), which permits use, distribution and reproduction for non-commercial purposes, provided the original is properly cited and derivative works building on this content are distributed under the same license.

IntechOpen

IntechOpen