

AGENTS FOR INTEGRATING DISTRIBUTED DATA FOR COMPLEX COMPUTATIONS

Ahmed M. KHEDR

*Mathematics Department, Faculty of Science
Zagazig University, Egypt
e-mail: amkhedr@yahoo.com*

Raj BHATNAGAR

*ECECS Department
Cincinnati University
Cincinnati, USA
e-mail: raj.bhatnagar@uc.edu*

Manuscript received 26 May 2006; revised 5 February 2007
Communicated by Jacek Kitowski

Abstract. Algorithms for many complex computations assume that all the relevant data is available on a single node of a computer network. In the emerging distributed and networked knowledge environments, databases relevant for computations may reside on a number of nodes connected by a communication network. These data resources cannot be moved to other network sites due to privacy, security, and size considerations. The desired global computation must be decomposed into local computations to match the distribution of data across the network. The capability to decompose computations must be general enough to handle different distributions of data and different participating nodes in each instance of the global computation. In this paper, we present a methodology wherein each distributed data source is represented by an agent. Each such agent has the capability to decompose global computations into local parts, for itself and for agents at other sites. The global computation is then performed by the agent either exchanging some minimal summaries with other agents or travelling to all the sites and performing local tasks that can be done at each local site. The objective is to perform global tasks with a minimum of communication or travel by participating agents across the network.

Keywords: Agents, distributed databases, vertical partition, decomposable algorithm, decision tree, association rule

1 INTRODUCTION

The parallel implementations of pattern analysis algorithms take advantage of the high performance of multiprocessor computer systems and work by transferring data from one processor to the other. It is desirable to implement algorithms in parallel and even build specialized hardware chips when large amounts of data are available at a single geographical site. The main focus of most of the parallel and distributed algorithms has been on closely coupled processor systems where data can be easily shared by the processors and the number of processors assigned to an algorithm may depend on the data to be processed.

Distributed Data Sources: Computing situations that are beginning to emerge in the networked environment require data and knowledge from a number of geographically distributed sites to be simultaneously considered. A number of geographically distributed databases together form an implicitly specified global dataset that contains all the data relevant for a computation. For example, some pattern discovery tasks may require simultaneous consideration of data, parts of which reside in census databases, labor statistics databases, and employment related databases. Each of these is a huge database and resides on a different site in a different city. One cannot hope to easily move all these databases to a single computer site, merge or join them, and then execute an algorithm with the tuples in the resulting humongous database. It would be desirable to have algorithms that let the individual databases reside at their own sites and work with an imagined implicit join of the databases by decomposing themselves into localized computations such that each localized computation can be performed locally within a single site using its physical database. A common constraint in these situations is that the data cannot be moved to other network sites due to security, size, privacy and data ownership considerations. Example of such situations is we may need to compute decision trees, association rules, or some complex statistical quantities using data from a census database, a diseases database, a labor statistics database, and a few pollution databases located in ten different cities across the country. It is impossible to bring these databases together and join them for performing some computations. Also, a new instance of some computation may require data from a different set of participating nodes and databases.

Our approach here has been to present a methodology and design of decomposable version of association rules, and induction of decision trees algorithms. In our methodology each data source (a network node) is represented by an agent. This

agent knows all about its underlying database and can access any part of it, as represented in Figure 1, which can be described as follows:

1. Some network node (Init-Node) wants to perform a computation C which requires a body of data D . The entire data D may not be available on Init-Node itself.
2. A search is performed over the network to identify those other nodes that can provide some relevant parts of D for the computation and are willing to cooperate. Init-Node selects a sufficient set of participant nodes that together constitute the body D . Attributes of databases at different nodes may be unique to their sites or may exist at more than one participating sites.
3. The initiator Init-Node determines a decomposition of computations C , compatible with the distribution of attributes of D . It then seeks results of local computations from the participating nodes and composes them to construct the global result. A number of (decomposition, partial computation, composition) iterations may be required for completion of C .

If the computation does not require updates to databases, then the agent also does not need an update privilege for its underlying data. The desired global computation, such as the need to induce a decision tree from underlying databases, is conveyed to the agents of the participating sites. Each agent then determines the local computations that it needs to perform, keeping in mind the constraints of shared data with other sites and also the local results that it needs to share with other agents in order for the global result to evolve at Init-Node. An alternative to communicating with agents at other sites is that a single agent visits each of the participating sites and performs some local computation at each site when it visits. Objectives of the agent's design include minimization of communication across the sites and enough generality of the formulation to permit agents to handle different sets of participating sites and different patterns of knowledge sharing across the participating nodes.

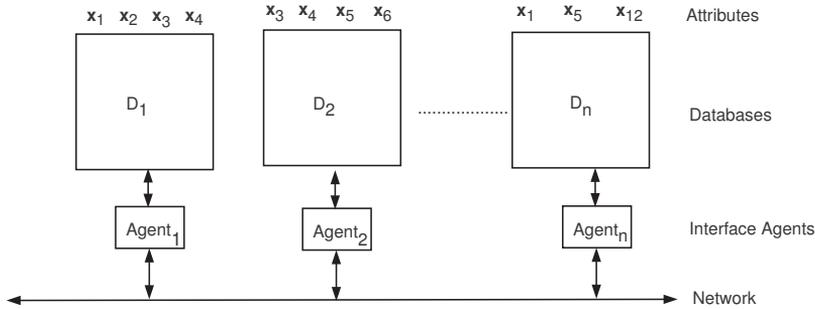


Fig. 1. Databases represented by agents

A desired global computation, such as the need to induce a decision tree from underlying databases, is conveyed to the agents of the participating sites. Each agent also knows about the other participating sites. It then determines the local computations that it needs to perform, keeping in mind the constraints of shared data with other sites and also the local results that it needs to share with other agents, in order for the global result to evolve at either one of, or each of the participating agents. An alternative to communicating with agents at other sites is that a single agent visits each of the participating sites and performs some local computation at each site when it visits. Objectives of the agent's design include minimization of communication across the sites and enough generality of the formulation to permit agents to handle different sets of participating sites and different patterns of knowledge-sharing across the participating nodes.

Our presented algorithms have the capability to decompose their computations to fit the nature of data distribution across the network nodes. The objective of the algorithms is to perform pattern discovery tasks for any arbitrary data distribution across the network by exchanging summaries derived from local databases. Our idea is to decompose the steps of an algorithm into localized computational steps that can be executed at the participating sites and the intermediate results thus obtained are transmitted to the central site. This may have to be repeated a number of times until an algorithm is completed. It thus works by exchanging partial results during the execution of an algorithm. Our algorithms are tailored for situations in which we don't have closely connected processors. There are multiple processors but they are independent and reside at geographically distant sites.

The rest of the paper is organized as follows: In the next section, we give an introduction to the basic concepts of our approach. A decomposable version of the association rule algorithm with its complexity computing is given in Section 3. In Section 4, we give a decomposable version of the induction of decision tree algorithm and its complexity computing. In Section 5, we present our simulation results. In Section 6 we give the related works to our handled problems. We conclude our work in Section 7.

2 INTEGRATION OF DISTRIBUTED DATA

In the situation modeled, here we consider n databases located at n different network sites, and all of them together constitute the dataset \mathcal{D} for the global computation. As an abstraction, we model the database D_i at each i^{th} node by a relation containing a number of tuples.

The set of attributes contained in D_i is represented by X_i . For any pair of relations, (D_i and D_j), the corresponding sets X_i and X_j may have a set of shared attributes given by S_{ij} . Since an arbitrary number of independent, already existing, databases may be consulted for a computation, we cannot assume any data normalization to have been performed for their schemas.

The implicit dataset \mathcal{D} with which the computation is to be performed is a subset of the set of tuples generated by a *Join* operation performed on all the participating relations (D_1, D_2, \dots, D_n) . However, the tuples of \mathcal{D} cannot be made explicit at any one network site by any one agent because the D_i 's cannot be moved in their entirety to other network sites. The tuples of \mathcal{D} , therefore, must remain implicitly specified only to one agent. This inability of an agent to make explicit the tuples of \mathcal{D} is the main problem addressed in the generalized decomposition of global algorithms and is discussed in later sections.

To facilitate computations with implicitly specified sets of tuples of \mathcal{D} , we define a set (S) that is the union of all the attribute intersection sets (S_{ij}), that is,

$$S = \bigcup_{i,j,i \neq j} S_{ij}. \quad (1)$$

The set S , thus, contains the names of all those attributes that are visible to more than one agent because they occur in more than one participating D_i . We define a relation *Shared* containing all possible enumerations for the attributes in the set S . This formulation of S facilitates similar treatment for horizontally or vertically partitioned datasets because horizontal partitioning can be seen as the case where all attributes are shared.

2.1 Nature of Data Distribution

Let us say there are n different sites containing databases D_1, D_2, \dots, D_n respectively. Depending on the sets of attributes contained in each D_i , there are two primary ways in which the databases, together, may be seen as forming an implicit global dataset \mathcal{D} .

Horizontally Partitioned Datasets: Figure 2 shows partitioning of \mathcal{D} into components D_1, D_2, \dots, D_n such that each component D_i contains the same attribute set (X_i), but a different set of data tuples. The set of shared attributes (S) is the same as X_i for each database.

Vertically Partitioned Datasets: Figure 3 shows another way in which components of \mathcal{D} may be distributed across a network. In this case, each component (D_i) may share some attributes with other databases (D_j and $j \neq i$). Each D_i may also contain some attributes not shared with any other database.

In effect, each D_i is a projection of an implicit global \mathcal{D} . Vertically partitioned datasets are of more interest because they provide an opportunity to share knowledge across the participating nodes. Our algorithms are designed to work with vertically partitioned databases and can also work with horizontally partitioned databases by considering all attributes are *Shared*.

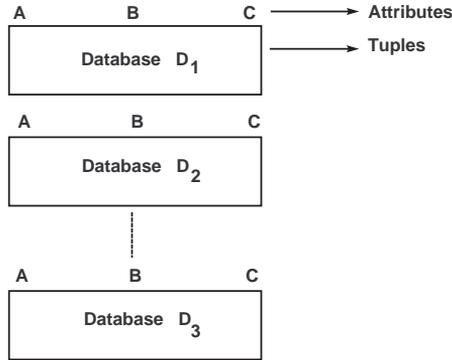


Fig. 2. Horizontally partitioned datasets

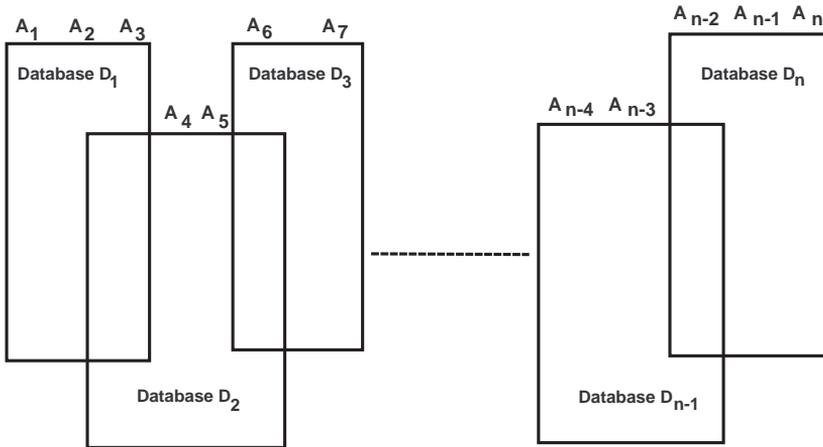


Fig. 3. Vertically partitioned datasets

2.2 Agent's Decomposition Task

The objective of an agent is to perform the global computation by communicating with other similar agents at other sites, and each agent performing some computation with its local database. Each agent should be able to decompose the global computation into local computations – in the context of and as constrained by the sharing of attributes across the participating agents and perform its local part with its own data.

Each *agent_i* in Figure 1 represents a *D_i* and communicates with similar agents at other nodes to exchange the results of its local computations. The decomposition methodologies discussed here can be seen to reside with each individual agent; each

agent is also capable of initiating and completing an instance of a global computation by either exchanging local results with other agents, while stationary at their respective sites, or by launching a mobile agent that visits other network sites. In the case of a mobile agent, the decomposition tools and knowledge reside with the mobile agent.

Let us say a result \mathcal{R} is to be obtained by applying a function \mathcal{F} to the implicit dataset \mathcal{D} . That is:

$$\mathcal{R} = \mathcal{F}(\mathcal{D}). \quad (2)$$

When the global computation is to induce a decision tree from \mathcal{D} , the value of \mathcal{R} is the induced decision tree, and \mathcal{F} corresponds to the implementation of an algorithm for inducing \mathcal{R} from \mathcal{D} .

Distributed databases used by the agents cannot make explicit the tuples of \mathcal{D} , which remain implicit in terms of the explicitly known components D_1, D_2, \dots, D_n . The set S of shared attributes determines what explicit \mathcal{D} would be generated by the individual data components. An implementation of \mathcal{F} in Equation (2) for some S can be engineered by a functionally equivalent formulation given as:

$$\mathcal{R}(S) = H[h_1(D_1, S), h_2(D_2, S), \dots, h_n(D_n, S)]. \quad (3)$$

That is, a local computation $h_i(D_i, S)$ is performed by *agent_i* using the database D_i and the knowledge about the attributes shared among all the data sites (S). The results of these local computations are aggregated by an agent using the operation H . However, it may not be possible to decompose a complex computation such as the full algorithm for inducing a decision tree into local computations and an aggregator. We can decompose smaller computational primitive steps of such a complete algorithm and the agent keeps track of the control aspects of sequencing various steps of such an algorithm.

The number and nature of h_i operators and the nature of H would vary with the participating D_i s and the set of attributes (S) shared among them. Hence, a different set of h -operators would need to be generated by the agent for each new instance of D_i 's and S .

Figure 5 shows the process by which the agent would compute \mathcal{R} from the D_i s. That is, a local computation $h_i(D_i)$ is performed at every *Node_i* using the database D_i . The results of these local computations are aggregated using the operation H .

The component operators of a decomposition (H and h_i s), therefore, need to be dynamically determined by the agent for each instance of $\mathcal{F}(\mathcal{D})$, depending on the participating nodes, the attributes contained in their native databases, and the sharing pattern of attributes.

2.3 Stationary and Mobile Agents

We consider two types of agents for computing the decomposed h_i and H functions. Stationary agents that stay at their respective data sites compute local h_i 's and

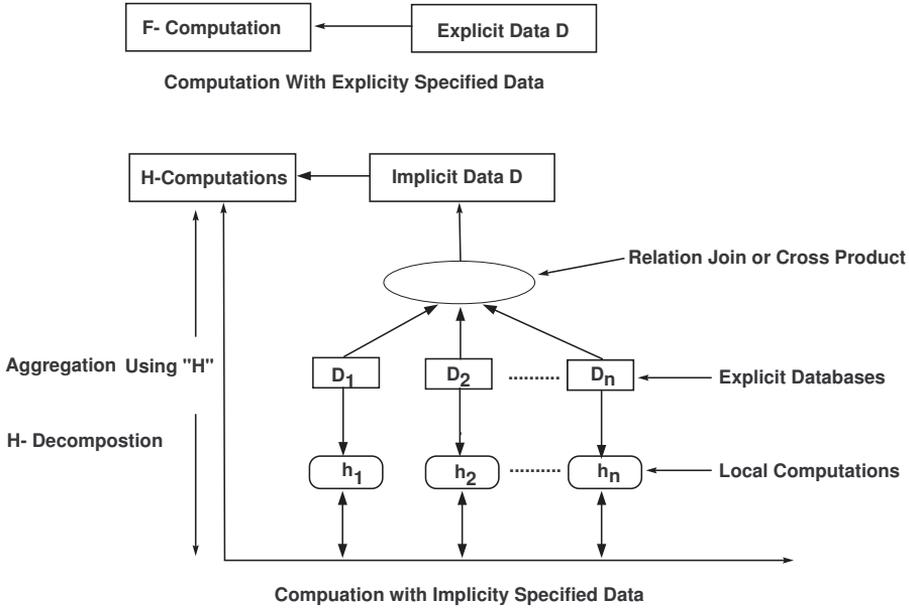


Fig. 4. Computations in explicit vs. implicit data spaces

send them to a coordinating agent who applies H operation to all the local results. Mobile agents move from one site to another, perform local h_i at each site that they visit, and at the end apply the H operation to the gathered results. In the following discussion we present complexity for both of these kinds of agents.

2.4 Cost Models for Algorithmic Complexity

Traditionally, the complexity of algorithms has been measured in terms of CPU time and the required memory. This cost model is well-suited for computations on a single computer and the closely-coupled processors model. When a number of loosely networked nodes are involved in a cooperative computation, the communication cost becomes the overwhelmingly dominant component of the total cost. Complexity for distributed query processing in databases has been discussed in [25], and the cost model used is total data transferred for answering a query. This cost model suits those applications well, where a large amount of data is exchanged during a computation. We have used here, and in other similar work [6, 17], cost models involving the number of messages exchanged and reflecting the efficiency of decomposition carried out by the network algorithm. The following are the three cost models for analyzing the complexity of our algorithms.

1. Communication Cost only: In this cost model we count the number of messages, N_m , that must be exchanged among all the participating sites in order to complete the execution of the algorithm. One *message exchange* includes one message sent by a site requesting some computation from another site and the reply message sent by the responding site. This cost model is relevant in situations where communication cost of messages is orders of magnitude larger than the cost of computations at local sites.
2. Communication Plus Computation Cost: In this model we examine a weighted sum of the number of messages exchanged and the number of local computations performed. If N_m messages are exchanged among the sites and a total of N_c computational units are performed at all the sites combined, then the algorithm's cost is given by: $N_m + N_c$. This cost model is useful when the local computation time within a site is not negligible and must be included within the cost model. When the databases stored at the sites are huge, as in many scientific and data-mining applications, the time to execute a local computation may be comparable to the time taken for exchanging a message across a wide-area network.
3. Elapsed Time Cost: In this model we examine a weighted sum of the number of messages exchanged and the number of local computations performed, while accounting for parallel transmission of messages and simultaneous execution of local computations at the participating sites. If N_m messages are exchanged among the sites and a total of N_c computational units are performed at all the sites combined, then the algorithm's cost is given by: $N_m + N_c/p$ where p is the average number of messages that can be exchanged in parallel. This cost model is useful when our criterion is the total elapsed time for executing the algorithm.

In the case of a Mobile Agent, the decomposition tools and knowledge reside with the mobile agent. This agent has the set *Shared* stored in it. During a visit to a data site, it can compute the local result for that site. Once all the sites have been visited, the aggregator H can be then applied to the local results collected from all the sites. Therefore, the mobile agent needs to visit each site only once in order to compute the global results.

In this paper each step of the algorithm must exchange a number of messages for evaluating the various quantitative values. Each message is generally of a very small length, but the number of messages may grow very fast. We ignore the local computation cost N_c because we perform only queries at the local sites.

3 ASSOCIATION RULES

Given a set of transactions where each transaction is a set of items, an association rule is an expression $X \Rightarrow Y$, where X and Y are sets of items. Intuitively, it would mean that transactions in the database which contains the items in X also contain the items in Y . Traditionally, most of the algorithms to determine association rules

are confined to a single large database. The main issues have been the size of the data, how many passes that are to be made over the database and the resultant time taken. However, this research is focused on mining the association rules in a distributed database scenario where explicit join is not possible because of the security, privacy, and other issues. Algorithms for discovering association rules in databases [1, 2] have been studied extensively. The main phases of various versions of the algorithm involve iterating the following 2 steps:

1. Enumerate item-sets at level L_k from the frequent item-sets (the sets of items that have minimum support) determined at level L_{k-1} .
2. Determine the *support* and the *confidence* levels for the item-sets and rules at level L_k .

The decomposability of an association rule algorithm can be similarly divided into 2 major tasks: (1) Maintenance of the active item-sets and enumeration of the candidates at the next level from the frequent item sets at the preceding level; (2) Computation of the *support* and *confidence* levels. A general decomposition of the algorithm can be implemented as follows: an agent at any network node initiates the algorithm; this agent performs within itself all the control aspects of the algorithm, such as the tasks of maintaining and generating the candidate item-sets; computation of *support* and *confidence* levels requires consultation with agents at other sites and it is this step that needs to be redesigned by an agent using decomposition; the decomposed version for *support* computations can then be repeated at each iterative step where the algorithm requires it, the control being with the agent.

3.1 Computational Primitives

The most common computational primitive needed in the above algorithm is the count of all tuples in \mathcal{D} to be determined only by obtaining local results from each participating agent. Counts of tuples that satisfy certain attribute-value conditions are a little bit more complex and we describe them below.

3.1.1 Count of Tuples in Implicit Space

When the tuples of \mathcal{D} are explicitly available in a relation then the count of all its tuples can be obtained easily. For our case of an implicitly defined set of tuples, we can decompose the counting process in such a way that various local count requests can be sent to the agents of individual D_i s and their responses can then be composed to construct the total count for the tuples in implicit \mathcal{D} . The decomposition for obtaining the count $N_{total}(\mathcal{D})$ is as follows:

$$N_{total}(\mathcal{D}) = \sum_j \prod_t N(D_t)_{cond_j} \quad (4)$$

where the subscript $cond_j$ specifies a condition composed from the attribute-value pairs of the j^{th} tuple of the relation *Shared*, n is the number of participating

agents (D_i s), and $N(D_t)_{cond_j}$ is the count in relation D_t of those tuples that satisfy the condition $cond_j$.

As per the decomposition expressed in Equation (3), we can see that

$$h_i(D_i, S) = N(D_t)_{cond_j} \quad (5)$$

where j corresponds to the j^{th} tuple of *Shared*. One such summary is needed from each agent for each tuple in the relation *Shared*. The relation *Shared* can be controlled by one agent or maintained by each agent separately, and thus it can reduce the communication among the agents. The function H in this case would perform a sum-of-products from the summaries as per Equation (4). Each term in the product is the count of tuples satisfying condition $cond_j$ in a D_i . The resulting product produces the number of distinct tuples that would be contributed to the implicit *Join* of all the D_i s for the condition specified by $cond_j$. The summations in the above expression amount to selecting each tuple of *Shared* as $cond_j$ and then summing the product terms obtained for each tuple. This expression, therefore, simulates the effect of a *Join* operation performed on all the databases without explicitly enumerating the tuples.

A very desirable aspect of the above decomposition of $N_{total}(\mathcal{D})$ is that each product term $N(D_t)_{cond_j}$ can be translated into an SQL query; select count (*) where $cond_j$ can be performed by the local agent at D_t .

Communication Complexity: When there are k attributes in the relation *Shared* and each attribute has I values on the average, the relation *Shared* would have $k * I$ tuples in all.

First we consider the case of stationary agents at the local sites. If there are n participating agents, then one agent would be sending one request to each of the $(n - 1)$ agents for each tuple in *Shared*, amounting to a total of $(n - 1) * k * I$ messages being exchanged among the agents.

However, it is possible to send a request to an agent for all $h_i(D_i, S)$ values, that is, values corresponding to all tuples $cond_j$ of S in one request, and receive all the summaries in one message. This reduces the number of messages exchanged to n , the same as the number of participating agents. The trade-off between the two approaches is that the first one may be considered more secure for transmission over a network because each message contains only very little information about the participating databases. The second alternative requires very few messages, but each message contains more information about each database. At no time, though, actual data tuples are transferred across the sites; only counts of tuples satisfying certain conditions are transferred.

Now we consider the case of mobile agents. This agent has the set *Shared* stored in it. During a visit to a data site, it can compute the local h_i for that site. Once all the sites have been visited, the sum-of-products (H aggregator) can be applied to the local results collected from all the sites. Therefore, the mobile agent

needs to visit each site only once in order to compute the count of all the tuples in implicit \mathcal{D} .

3.1.2 Support for Candidate Sets

We can easily extend the above decomposition for count to the counts of only those tuples that satisfy a certain new condition by simply changing $cond_j$ in Equation (4):

$$N_{new-condition} = \sum_j \left(\prod_{t=1}^n (N(D_t)_{cond_j.and.new-condition}) \right). \quad (6)$$

This would be required to determine the support level for a candidate frequent item set. The values of attributes in the frequent item set would form the *new-condition*

The way the support measure for a candidate frequent item-set would be computed by an agent can be viewed as follows: In relation *Shared*, it retains only those tuples that match the attribute value pairs for the conditions specified in the candidate set; determine a count of the tuples that is obtained using this reduced *Shared* relation; this resulting candidate count divided by the count N_{total} provides the support level for a candidate set of attribute-value pairs. Each count would require n messages to be exchanged and thus a support level can be computed by exchanging $2 * n$ messages. However, messages for each candidate set currently under considerations can be packed into a single message. Therefore, the support levels for all the candidate sets at a level can be computed by exchanging only $2 * n$ messages among the nodes.

For a mobile agent, the local results for computing all the candidate relevant tuple counts and also the total tuple counts can be gathered during a single visit to a site. Thus, the mobile agent can compute the support levels for all its candidate item sets by visiting each site only once and then aggregating the local results.

3.2 Full Algorithm Complexity

As seen above, frequent item-sets at each level of the association rule algorithm can be determined by exchanging only $2 * n$ messages among the participating nodes. If an association rule algorithm needs to run up to k levels, then we need to exchange a total of $2 * n * k$ messages among the stationary agents to run the association rule algorithm. This number of messages is not dependent on the number of tuples contained in each database and the system, therefore, is easily scalable to large databases. Also, this number of messages is much smaller than the data that may need to be transferred if we were to accumulate all databases at one site and then perform the data mining task.

For a mobile agent, an algorithm performing k iterations would mean visiting each site k times and after that the agent would be able to get the global results.

While it is easy to decompose arithmetic primitives, the step-sequencing and control aspects of an algorithm are more difficult to decompose efficiently and in a generalizable manner. For the results described here we have assumed that the algorithm initiator node executes the control steps of the original algorithm and decomposes each arithmetic computation as it sequences through the algorithm's steps.

The above algorithm decomposition works well for both horizontally partitioned and vertically partitioned databases. The algorithm is especially beneficial for vertically partitioned databases and can run horizontally partitioned databases as a special case. More efficient implementations are possible when we encounter only a horizontal database.

4 INDUCTION OF DECISION TREES

We can also easily perform decomposition of a decision tree induction algorithm using minimization of average entropy because entropy computation depends only on various tuple counts being obtained from the participating databases.

Various tree induction algorithms [7, 22], modeled after Quinlan's entropy-based tree-induction algorithms start by considering the complete dataset \mathcal{D} belonging at the root of the tree and then repeating the following steps until all or a large majority of tuples at each leaf node of the tree belong to some unique class.

1. Pick one such dataset at a leaf node, some large fraction whose tuples belong to different classes.
2. Select an attribute a_j , having m distinct values: $a_{j1}, a_{j2}, \dots, a_{jm}$. The attribute that results in minimum average entropy for the resulting partitions is chosen. This entropy value can be computed by the decomposition primitives described above, mainly the counts with various conditions placed on tuples for determining the appropriate probability values.
3. Split \mathcal{D} into m distinct partitions such that the k^{th} partition contains only those tuples for which $a_j = a_{jk}$.
4. The m distinct partitions are added to the tree as child datasets of the partitioned parent dataset. These child nodes reside at the end of m branches emanating from the parent node.

In the preceding discussion we have included the complexity of performing decomposition of each computational step in terms of the number of messages to be exchanged among the nodes. We show below an expression for the number of messages that need to be exchanged among the stationary agents by transferring only one $h_i(D_i, S)$ summary at a time for generating a simple decision tree using entropy minimization at each step and dealing with the implicit set of tuples. Let us say:

- There are n databases, D_1, D_2, \dots, D_n , residing at n different network sites.
- There are k attributes in set S of shared attributes. Each attribute in this set appears at more than one site.
- There are m distinct attributes in $\mathcal{D} (\cup_{i=1}^n A_i)$ combined.
- There are l possible discrete values for each attribute in set S .

The informational entropy that is computed by the algorithm at each leaf node of the growing tree is given by the expression

$$E = \sum_{b=1}^m \left(\frac{N_b}{N_t} \times \left(\sum_c -\frac{N_{bc}}{N_b} \log \frac{N_{bc}}{N_b} \right) \right) \quad (7)$$

where N_b is the total number of tuples in the database at a parent node of the tree, N_b is the number of tuples at a child node, and N_{bc} is the number of tuples in N_b that belong to class c .

Complexity: To evaluate the entropy for an implicit database \mathcal{D} once, we need to compute the following quantities:

- one N_t count;
- $l * N_b$ counts; and
- $l^2 * N_{bc}$ counts.

The computation of each entropy value, therefore, requires an exchange of $n * l^{k+2}$ messages among the participating nodes. For a dataset at depth \mathcal{D} in the decision tree, the number of messages exchanged would be $(md) * n * l^{k+2}$. If we assume that on the average the decision tree is akin to a filled l -ary tree with p levels, then the total number of message-exchanges needed would be:

$$(n * l^{k+2}) \sum_{d=0}^p (m - d) * l^p. \quad (8)$$

The above expression gives an estimate of the number of messages to be exchanged for constructing a decision tree using the decomposed version of the algorithm, but assuming that only one integer value is exchanged per message.

However, using a more efficient implementation in which all the needed $h_i(D_i, S)$ values are requested from a node in a single message, we need only $c * n$ messages to compute a count or a sum of products, where c is a constant and n is the number of participating nodes. In this case, an entropy computation would require an exchange of only $l * n$ messages and computation of $d * l$ entropy values would require $d * l * l * n$ messages to be exchanged. This is much more efficient than transferring a single summary per message.

A mobile agent can perform all potential computations for a decision tree level once it visits a data site. Therefore, a mobile agent would need to visit each site at most as many times as is the depth of the intended decision tree. During each cycle

of visits it will compute entropies for all nodes that can be split at the current level of the tree. At the end of each cycle it will then compute the actual partitioning to be performed and then it will continue with the cycle of visits for the next level of the tree. It is greatly advantageous that a d -level decision tree can be learned by an agent from distributed data by visiting each data site only d times.

5 SIMULATION RESULTS

We have performed a number of tests to demonstrate that counts of tuples, candidate support levels, informational entropy values and mining association rules can be computed in a distributed knowledge environment without moving all the databases to a single site. These tests have been carried out on a network of workstations connected by a LAN and tested against a number of databases of different sizes. The algorithms have been tested on both test data and real life databases; both flat files and relational databases like MySQL were used to test the algorithms. We have implemented the algorithms using Java and RMI (Remote Method Invocation), and used JDBC (Java Database Connectivity) to interface with the databases. This was done to provide a standard interface and platform independence. We present the results in the form of graphs, which provide a comparative analysis of when the algorithms are run using the non-optimized version, i.e., sending one summary per message, and using the optimized version, i.e., sending all the summaries for a particular site in one message.

Figure 5 shows how the time taken to compute the total number of tuples (N_{total}) in an implicit database \mathcal{D} changes with the size of the individual database. As we can see, when we exchange one summary per message, the time taken to compute the count increases as the size of the database increases. However, when we use the optimized method, the time taken to compute the count reduces considerably and depends on the number of participating nodes.

Figure 6 shows how the number of messages exchanged between the coordinator site and the remote sites varies with the number of tuples in the database. It can be seen easily that the number of messages exchanged varies exponentially with the size of the database when we send one summary per message. The result validates the expression for the total number of messages exchanged as given above. However, in the optimized version, when we receive all the summaries in a single message, the number of messages exchanged was a constant depending upon the total number of participating nodes.

Coming to the implementation of the Apriori Association Rules algorithm, we see that Figure 7 gives an analysis of the time taken and the number of messages exchanged between the learner and the remote sites for mining association rules in distributed databases. We note that the graph of the distributed databases indicates exponential complexity. This implies that although an algorithm can be decomposed into sub-parts and run on distributed databases, it will typically incur an extra cost in terms of total execution time and messages exchanged.

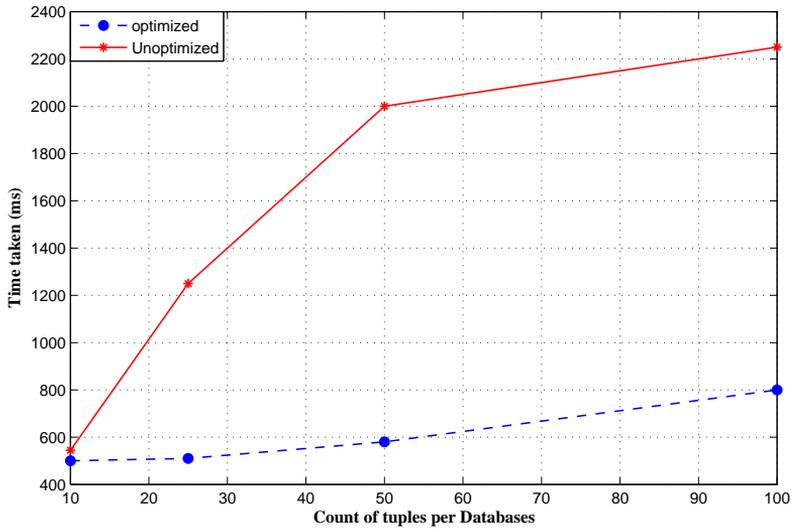


Fig. 5. Time taken to calculate N_t for different database sizes

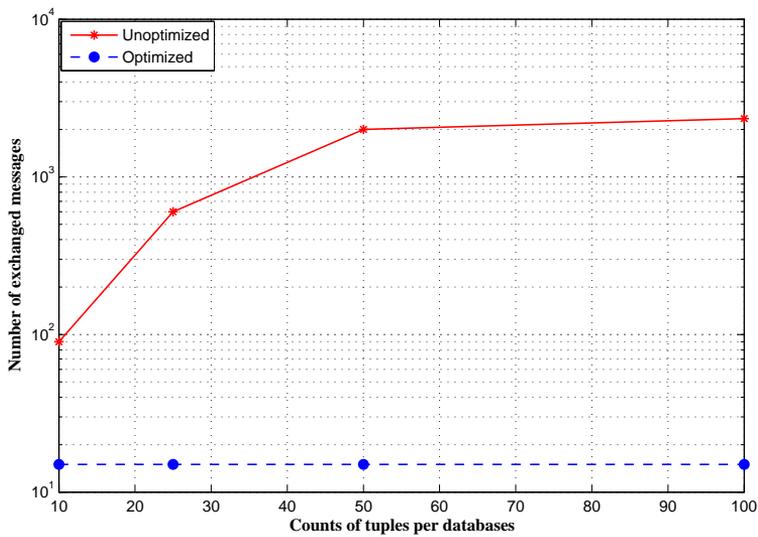


Fig. 6. Number of exchanged messages to calculate N_t for different database sizes

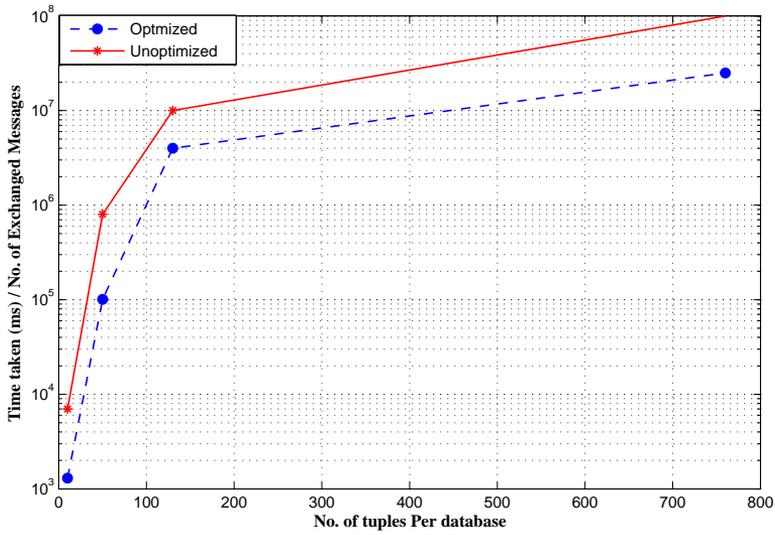


Fig. 7. Number of exchanged messages and time taken for rule mining in distributed database

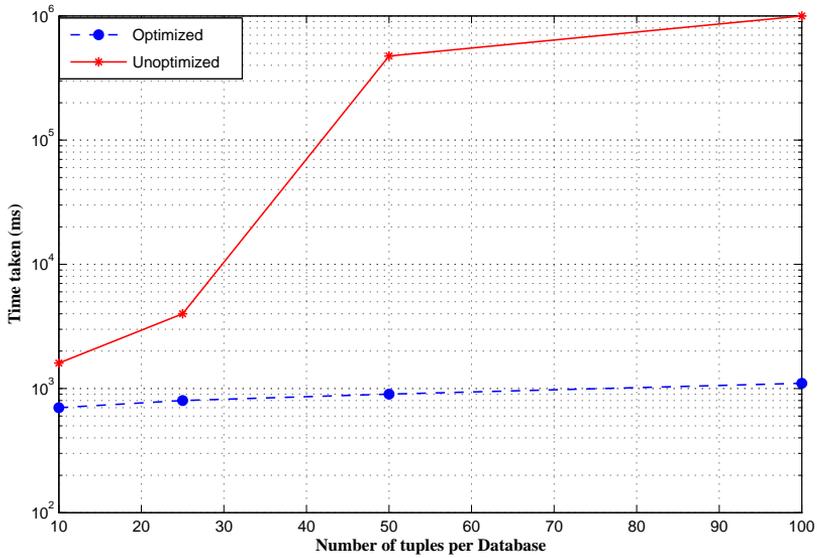


Fig. 8. Time taken to calculate the entropy for different database sizes

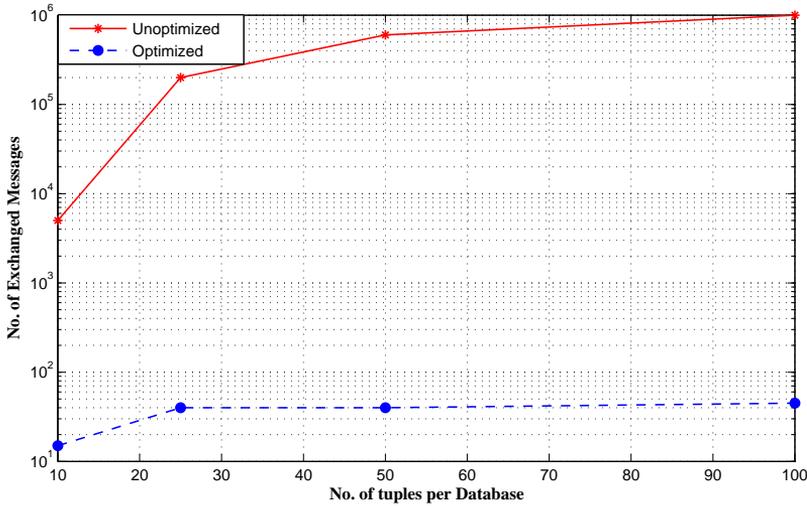


Fig. 9. Number of Exchanged Messages to calculate entropy for different database sizes

Figure 8 shows the variation of the time taken to calculate entropy and the number of tuples in a database. Similar to the analysis in Figure 7, we see that the time taken to calculate entropy varies exponentially when only one summary is sent per message. However, the computing time reduces significantly when using the optimized version. Figure 9 shows the number of exchanged messages to calculate entropy for different database sizes (unoptimized and optimized versions).

6 RELATED WORKS

There has been extensive research in algorithms for sequential and parallel architectures [3, 4, 5, 8, 16]. The main focus of parallel and distributed algorithms has been on systems of closely-coupled processors, where data can be easily shared by the processors. The distributed knowledge environment, where data cannot be shared as easily as a shared memory and must be transmitted over a wide-area network in the form of small message packets, needs a set of network algorithms, which minimize the traffic over the wide-area network.

References [8, 15, 14] offer an attractive approach to implementing modular and extensible distributed computing systems. Each data site has one or more associated agents that process the local data and communicate the results to the other agents or to a Coordinator supervising agent that controls the behavior of the local agents. In [21], a distributed system that takes advantage of data resource migration for transaction processing in ATM networks has been proposed. The proposed system

provides mechanisms to select the transaction processing method, to migrate data resources in a way that reduces the time delay and message traffic in locating and accessing them.

Central to our approach is a clear separation of concerns between hypothesis construction and moving data, and extraction of statistics from data. This separation makes it possible to explore the use of sophisticated techniques for query optimization that yield optimal plans for gathering statistics from distributed data sources under a specified set of constraints describing the query capabilities and operations permitted by the data sources.

In our system, the computing agent at each node does not need to have any uncertainty about the state of data or knowledge of other sites or computing agents. The agent only needs to ask for their local results and it would be truthfully given the needed information sought by other agents. However, this access can be restricted to prohibit any actual data tuples flowing out of a site. All examples given in this paper require only the results of very high level local computations from each site and actual data tuples never leave a site. The goal of the agents in our formulation is also to minimize the exchange of information among themselves for performing the global computations.

Our algorithmic decompositions can be seen as regular data mining algorithms being implemented by a number of coordinated agents either exchanging messages among themselves or visiting participating nodes to gather results of local queries and computations. Multi-agent systems research has addressed many issues relating to the distribution of knowledge and processing capability over a loosely connected communication network. In most of this work [10, 11, 13, 23], agents are modeled as having only a limited view of the global resources and knowledge. The objective of learning by each agent is that the whole society should converge to an optimal operating point after each agent has individually learned about its own optimal performance. In the work on multi-agent learning [23, 26, 27], most of the focus is also on learning about the environment by observing the behavior of other agents in the environment. In contrast, our approach is directed at systems where cooperative agents freely access local results from other agents to evolve concepts from their collective knowledge, while trying to minimize the communication of messages and data among themselves.

Two of our examples in this paper relate to learning and data mining computations, which are widely investigated fields, and we have extensively examined [6] the decomposability of decision-tree induction [7, 18, 19, 22] algorithms. Some other work in distributed data mining [9, 20, 24] seeks to learn local models completely and then resolve their differences at the central coordinating site. This is in contrast to our approach where we seek to decompose every primitive of the global computation and then perform the decomposed steps at local sites. Algorithms for association rules [1, 2] have become very popular, but are designed for cases when the database resides on a single network site. We have adapted these algorithms for distributed knowledge environments. Our focus is on decomposing each primitive computational step of an algorithm and executing it for the same results

that would have been obtained if the databases were to be moved to one single site.

Database researchers have done much work towards the optimization of queries from distributed databases [28]. Databases from which the transfer of large amounts of actual data is not feasible cannot participate in distributed querying, but still may be useful for participating in network algorithms by exchanging local summaries and inferences. Intelligent Query Answering and Data Clustering in large databases have been addressed in [12, 29], and their treatment is also limited to databases residing and available at a single computer site.

7 CONCLUSIONS

We have demonstrated that agents can perform the integration of arbitrarily distributed data and knowledge for performing complex computations. Tasks such as counting tuples in imagined *Joins* of distributed databases, computation of support and confidence levels for candidate item-sets, and informational entropy values of implicit databases can be computed by appropriately coordinating agent actions. These actions are self-determined and self-controlled by the agents in response to the varying sets of participating agents and arbitrary overlaps in the local datasets. Also, for simple arithmetic computations, the number of messages to be exchanged among the n participating agents does not exceed the order of n . This is very significant because it gives us the scalability required for handling large databases. The number of tuples at individual network nodes may keep on increasing but the number of messages that need to be exchanged among the agents for a global computation remains constant. We have demonstrated the adaptability of an association rule learning algorithm, and an informational entropy-driven decision tree induction algorithm. We have shown the complexity of performing these computations in terms of messages that need to be exchanged among the stationary agents for performing these computations. We have also analyzed the number of visits that a mobile agent would need to make to each site for completing the global computation.

One very significant contribution of these results is that many mining and knowledge discovery tasks can be performed by agents on a number of databases residing at different network nodes without having to move the databases to a single site and the communication cost among the performing agents is also very low.

REFERENCES

- [1] AGRAWAL, R.—IMIELINSKI, T.—SWAMI, A.: Mining Association Rules Between Sets of Items in Large Databases. In SIGMOD 93 International Conference on Management of Data, pp. 207–216, 1993.
- [2] AGRAWAL, R.—SRIKANT, R.: Fast Algorithms for Mining Association Rules in Large Databases. In VLDB Conference, 1994.

- [3] AHO, A. V.—HOPCROFT, J. E.—ULLMAN, J. D.: *The Design and Analysis of Computer Algorithms*. Addison Wesley, MA, 1974.
- [4] AKI, S. J.: *Parallel Computation: Models and Methods*. Prentice Hall, New Jersey, 1997.
- [5] BARBOSA, V. C.: *An Introduction to Distributed Algorithms*. MIT Press, 1996.
- [6] BHATNAGAR, R.—SRINIVASAN, S.: Pattern Discovery in Distributed Databases. In *AAAI97*, pp. 503–508, 1997.
- [7] BREIMAN, L.—FRIEDMAN, J. H.—OLSHEN, R. A.—STONE, C. J.: *Classification and Regression Trees*. Belmont, CA: Wadsworth, 1984.
- [8] BRIM, L.—KRETINSKÝ, M.—JACQUET, J.—GILBERT, D.: Modeling Multi-Agent Systems as Synchronous Concurrent Constraint Processes. *Computers and Informatics*, Vol. 21, 2002, No. 6.
- [9] CHAN, P. K.—STOLFO, S. J.: Sharing Learned Models Among Remote Database Partitions by Local Meta-Learning. In *KDD*, pp. 2–7, 1996.
- [10] CHIA, M. H.—NEIMAN, D. E.—LESSER, V. R.: Poaching and Distraction in Asynchronous Agent Activities. In *ICMAS*, pp. 88–95, 1998.
- [11] DURFEE, E. H.—CORKILL, V. R.: Coherent Cooperation Among Communicating Problem Solvers. *IEEE Transactions on Computers*, Vol. 36, 1987, No. 11.
- [12] HAN, J. W.—HUANG, Y.—CERCONE, N.—FU, Y. J.: Intelligent Query Answering by Knowledge Discovery Techniques. *IEEE Transactions on Knowledge and Data Engineering*, pp. 373–390, 1996.
- [13] HUHS, M. N.—SINGH, M. P.—KSIEZYK, T.: *Global Information management via Local Autonomous Agents*. Morgan Kaufmann Publishers, 1997.
- [14] HONAVAR, V.—MILLER, L.—WONG, J. S.: Distributed Knowledge Networks. In *Proceedings of the IEEE Conference on Information Technology*, Syracuse, NY, 1998. IEEE Press.
- [15] JENNINGS, N.—WOOLDRIDGE, M.: Agent-Oriented Software Engineering. In J. Bradshaw, editor, *Handbook of Agent Technology*. MIT Press, 2001.
- [16] JAJA, J.: *An Introduction to Parallel Algorithms*. Addison Wesley Publishers, 1992.
- [17] KHEDR, A. M.—BHATNAGAR, R.: A Decomposable Algorithm for Minimum Spanning Tree. *Distributed Computing – Lecture Notes in Computer Science* Springer-Verlag Heidelberg, Vol. 2918, pp. 33–44, 2004.
- [18] MINGERS, J.: An Empirical Comparison of Pruning Methods for Decision-Tree Induction. *Machine Learning*, Vol. 4, 1989, pp. 227–243.
- [19] MINGERS, J.: An Empirical Comparison of Selection Measures for Decision-Tree Induction. *Machine Learning*, Vol. 3, 1989, pp. 319–342.
- [20] PROVOST, F. J.—HENNESSY, D. N.: Scaling Up: Distributed Machine Learning with Cooperation. *AAAI/IAAI*, Vol. 1, 1996, pp. 74–79.
- [21] SAXENA, P. CH.—CHOUDHURY, D. R.—GABRANT, G.: DSTP-AN A Distributed System for Transaction Processing Based on Data Resource Migration in ATM Networks. *Computers and Informatics*, Vol. 23, 2004, No. 3.
- [22] QUINLAN, J. R.: Induction of Decision Trees. *Machine Learning*, Vol. 1, 1986, pp. 81–106.

- [23] SEN, S.: Adaption Co-Evolution and Learning in Multi-Agent Systems. AAAI Press, 1996.
- [24] STOLFO, S.—FAN, D.—LEE, W.—PRODROMIDIS, A.—CHAN, P.: Jam: Java Agents for Meta-Learning: Issues and Initial Results. In AAAI Workshop on AI Methods in Fraud and Risk Management, 1997.
- [25] WANG, C.—CHEN, M.: On the Complexity of Distributed Query Optimization. IEEE Transactions on Knowledge and Data Engineering, Vol. 8, 1996, No 4, pp. 650–662.
- [26] WEISS, G.: Distributed Artificial Intelligence Meets Artificial Intelligence. Lecture Notes in Artificial Intelligence, Springer-Verlag, Vol. 1237, 1997.
- [27] WEISS, G.—SEN, S.: Adaption and Learning in Multi-Agent Systems. Lecture Notes in Artificial Intelligence, Springer-Verlag, Vol. 1042, 1996.
- [28] YU, C.—OZSOYOGLU, Z. M.—KAM, K.: Optimization of Distributed Tree Queries. Computer System Science, Vol. 29, No. 3, pp.409–445.
- [29] ZHANG, T.—RAMAKRISHNAN, R.—LIVNY, M.: BIRCH: An Efficient Data Clustering Method for Very Large Databases. In SIGMOD Rec. 25, Vol. 2, 1996, pp. 103–114.



Ahmed M. KHEDR received his B. Sc. degree in mathematics in June 1989 and the M. Sc. degree in optimal control in July 1995, both from Zagazig University, Egypt. In March 2003 he received his Ph. D. degree in computer science from University of Cincinnati, Ohio, USA. From March 2003 to January 2004, he was a research assistant professor at ECECS Department University of Cincinnati, USA. From January 2004 till now he is working at the Department of Mathematics, Faculty of Science, Zagazig University, Egypt. He has co-authored 20 works in journals and conferences related to optimal control, distributed databases, sensor network and decomposable algorithms.