

Hierarchical Decision Tree Induction in Distributed Genomic Databases

Amir Bar-Or, Daniel Keren, Assaf Schuster, and Ran Wolff

Abstract—Classification based on decision trees is one of the important problems in data mining and has applications in many fields. In recent years, database systems have become highly distributed, and distributed system paradigms, such as federated and peer-to-peer databases, are being adopted. In this paper, we consider the problem of inducing decision trees in a large distributed network of genomic databases. Our work is motivated by the existence of distributed databases in healthcare and in bioinformatics, and by the emergence of systems which automatically analyze these databases, and by the expectancy that these databases will soon contain large amounts of highly dimensional genomic data. Current decision tree algorithms require high communication bandwidth when executed on such data, which are large-scale distributed systems. We present an algorithm that sharply reduces the communication overhead by sending just a fraction of the statistical data. A fraction which is nevertheless sufficient to derive the exact same decision tree learned by a sequential learner on all the data in the network. Extensive experiments using standard synthetic SNP data show that the algorithm utilizes the high dependency among attributes, typical to genomic data, to reduce communication overhead by up to 99 percent. Scalability tests show that the algorithm scales well with both the size of the data set, the dimensionality of the data, and the size of the distributed system.

Index Terms—Data mining, distributed algorithms, decision trees, classification.

1 INTRODUCTION

THE analysis of large databases requires automation. Data mining tools have been shown to be useful for this task, in a variety of domains and architectures. It has recently been shown that data mining tools are extremely useful for the analysis of genomic data as well [1]. Since the number of genomic databases and the amount of data in them increases rapidly, there is a dire need for data mining tools designed specifically to target genomic data specifically.

Classification, the separation of data records into distinct classes, is apparently the most common data mining task, and decision tree classifiers are perhaps the most popular classification technique. Some recent works have shown that classification can be used to analyze the effect of genomic, clinical, environmental, and demographic factors on diseases, response to treatment, and the risk of side effects [2]. Providing efficient decision tree induction algorithms suitable for genomic data is therefore an important goal.

One interesting aspect of genomic databases is that they are often distributed over many locations. The main reason for this is that they are produced by a variety of

independent institutions. While these institutions often allow a second party to browse their databases, they will rarely allow this party to *copy* them. There could be a number of reasons for this: the need to retain the privacy of personal data recorded in the database, through questions regarding its ownership, or even because the sheer size of the data makes copying nonpermissively costly in CPU, disk I/O, or network bandwidth.

Our lead example in this paper is the task of mining genomically enriched electronic medical records (EMRs). Within a few years, it is expected that each patient's medical record will contain a genomic fingerprint. This fingerprint will be used mainly to optimize treatment and predict side effects. Existing genomic fingerprinting techniques, such as single nucleotide polymorphisms (SNPs) and Gene Expression Microarrays, yield records with tens of thousands of entries that are usually interpreted as binary (normal/abnormal allele or active/inactive gene, respectively). It is a common perception that an illness or treatment side effect can many times relate to just single SNPs or to the expression of few genes. In this paper, we focus on SNP data mainly because Microarray experiments still lack standardization. This may yield many other problems which call for prior use of additional tools before classification can be applied to the data.

Data mining of genomically enriched EMRs would be needed for the identification of unknown correlations and for the development of new drugs. It would best be performed on a national scale, using EMRs gathered by many different health maintenance organizations (HMOs). This would naturally extend the functionality of systems such as RODS [3] and NRDM [4] which already collect and analyze health data at a regional (RODS) and national (NRDM) scale. RODS, for example, accesses the database of

- A. Bar-Or is with HP Labs—Cambridge, One Cambridge Center, Cambridge, MA 02142-1612. E-mail: amir.baror@hp.com.
- D. Keren is with the Department of Computer Science, Haifa University, Haifa 31905, Israel. E-mail: dkeren@cs.haifa.ac.il.
- A. Schuster is with the Department of Computer Science, Technion—Israel Institute of Technology, Haifa, 32000, Israel. E-mail: assaf@cs.technion.ac.il.
- R. Wolff is with the Computer Science and Electrical Engineering Department, University of Maryland at Baltimore County, 1000 Hilltop Circle, Baltimore, MD 21250. E-mail: ranw@cs.umbc.edu.

Manuscript received 27 July 2004; revised 27 Oct. 2004; accepted 28 Feb. 2005; published online 17 June 2005.

For information on obtaining reprints of this article, please send e-mail to: tkde@computer.org, and reference IEEECS Log Number TKDE-0277-0704.

tens of hospitals using the HL7 [5] protocol to retrieve statistical information and detect disease outbreaks. Nevertheless, it is unlikely that an HMO would allow systems such as RODS to download its entire database. Hence, the need for distributed algorithms.

A distributed decision tree induction algorithm is one that executes on several computers, each with its own database partition. The outcome of the distributed algorithm is a decision tree which is the same as, or at least comparable with, a tree that would be induced were the different partitions collected to a central place and processed using a sequential decision tree induction algorithm. Since decision tree induction poses modest CPU requirements, the performance of the algorithm would usually be dictated by its communication requirements.

Previous work on distributed decision tree induction usually focused on tight clusters of computers, or even on shared memory machines [6], [7], [8], [9], [10], [11]. When a wide area distributed scenario was considered, all these algorithms become impractical because they use too much communication and synchronization. A kind of decision tree induction algorithm which is more efficient in a wide area system employs metalearning [12], [13], [14], [15]. In metalearning, each computer induces a decision tree based on its local data; then, the different models are combined to form the final tree. This final tree is an approximation of the one which would be induced from the entire database. Studies have shown that the quality of the approximation decreases significantly when the number of computers increase, and when the data become sparse. Because genomic databases contain many (thousands) attributes for each data instance and can be expected to be distributed over many distant locations, current distributed decision tree induction algorithms are ill-fit for them.

In this paper, we describe a new distributed decision tree algorithm, Distributed Hierarchical Decision Tree (DHDT). DHDT is executed by a collection of agents which correlate with the natural hierarchy of a national virtual organization. For instance, the leaf level agents may correspond to different HMOs (or clinics within an HMO) while upper levels correspond to regional, state, and national levels of the organization. DHDT focuses on reducing the volume of data sent from each level to the next while preserving perfect accuracy (i.e., the resulting decision tree is not an approximation). When tested on genomic SNP data with 1,000 SNPs in each data record, DHDT usually collects data about only a dozen of the SNPs—a 99 percent decrease in bandwidth requirements. The algorithm is suitable for any high dimension data, provided that the correlations in it are sparse as they are in genomic data. Both the hierarchic organization and the communication efficiency of DHDT give it excellent scalability at no decrease in accuracy.

The rest of the paper is structured as follows: We describe sequential decision tree induction in Section 2 and related work in Section 3. In Section 4, we provide bounds for the Gini index and the information gain functions. The DHDT algorithm is described in Section 5, and our experimental evaluation is given in Section 6. We conclude in Section 7.

2 SEQUENTIAL DECISION TREE INDUCTION

The decision tree model was first introduced by Hunt et al. [16] and the first sequential algorithm was presented by Quinlan [17]. This basic algorithm used by most of the existing decision tree algorithms is given here. Given a training set of examples, each tagged with a class label, the goal of an induction algorithm is to build a decision tree model that can predict with high accuracy the class label of future unlabeled examples. A decision tree is composed of nodes, where each node contains a test on an attribute, each branch from a node corresponds to a possible outcome of the test, and each leaf contains a class prediction. Attributes can be either *numerical* or *categorical*. In this paper, we deal only with categorical attributes. Numerical attributes can be discretized and treated as categorical attributes; however, the discretization process is outside the scope of this paper.

A decision tree is usually built in two phases: a growth phase and a pruning phase. The tree is grown by recursively replacing the leaves by test nodes, starting at the root. The attribute to be tested at a node is chosen by comparing all the available attributes and greedily selecting the attribute that maximizes some heuristic measure, denoted as the *gain function*. The minimal and sufficient information for computing most of the gain functions is usually contained in a two-dimensional matrix called the *crostable* of attribute i . The $[v, c]$ entry of the crostable contains the number of examples for which the value of the attribute is v and the value of the class attribute is c .

The decision tree built in the growth phase can “overfit” the learning data. As the goal of classification is to accurately predict new cases, the pruning phase generalizes the tree by removing subtrees corresponding to statistical noise or variation that may be particular only to the training data. This phase requires much less statistical information than the growth phase; thus, it is by far less expensive. Our algorithm integrates a tree generalization technique suggested in PUBLIC [18], which combines the growing and pruning stages while providing the same accuracy as the postpruning phase. In this paper, we focus on the costly growth phase.

2.1 Gain Functions

The most popular gain functions are information gain [17], which is used by Quinlan’s ID3 algorithm, and the Gini Index [19], which is used by Brieman’s Cart algorithm, among others.

Consider a set of examples S that is partitioned into M disjoint subsets (classes) C_1, C_2, \dots, C_M such that $S = \bigcup_{i=1}^M C_i$ and $C_i \cap C_j = \emptyset$ for every $i \neq j$. The estimated probability that a randomly chosen instance $s \in S$ belongs to class C_j is

$$p_j = \frac{|C_j|}{|S|},$$

where $|X|$ denotes the cardinality of the set X . With this estimated probability, two measures of impurity are defined: $entropy(S) = -\sum_j p_j \log p_j$ and $Gini(S) = \sum_j p_j^2$.

Given one of the impurity measures defined above, the gain function measures the reduction in the impurity of the

set S when it is partitioned by an attribute \mathbf{A} as follows: $Gain_{\mathbf{A}}(S) = \sum_{v \in Values(\mathbf{A})} \frac{|S_v|}{|S|} Imp(S_v)$, where $Values(\mathbf{A})$ is the set of all possible values for attribute \mathbf{A} , S_v is the subset of S for which attribute \mathbf{A} has the value v , and $Imp(S)$ can be $entropy(S)$ or $Gini(S)$.

3 RELATED WORK

The distributed decision tree algorithm described in [20] perhaps most resembles ours, in both motivation and the assumed distributed database environment (i.e., homogeneous databases with categorical attributes). Caragea et al. decompose the induction algorithm into two components: The first component collects sufficient local statistics and sends them to a centralized site, while the second component aggregates the statistics, computes the gain function, and chooses the best splitting attribute. Obviously, this algorithm has high communication complexity because it sends statistical data for each and every attribute in the database. The size of the crosstable of a single attribute depends on the size of the attribute domain, i.e., the number of distinct values that this attribute can receive, and on the number of distinct classes. For example, assume a genomic data set with $L = 10,000$ SNPs in each entry, each encoded in a single bit (usual or unusual allele). If a central network node in a network with $N = 1,000$ sites learns a decision tree for a binary class attribute with just $D = 100$ decision tree nodes (the maximal number of nodes would be $O(2^L)$), the number of bytes it would receive from the entire network is $O(LND)$, or, in our example, $10,000 * 2 * 2 * 1,000 * 100 * 4 = 16TB$ (assuming a crosstable entry is encoded in 4 bytes). In addition, it is often the case that the process should be repeated many times with different arguments (e.g., different classification goals). Therefore, the above algorithm requires high communication bandwidth between the participating nodes, which clearly does not exist in large-scale distributed systems.

The common approach to reducing the communication overhead would be to sample the distributed data set and collect a small subset of the learning examples for central processing. In addition to this sample, the different sites may also deduce decision trees based on their local data sets and transfer these decision trees to the central site. This has been the theme of an approach called *metalearning* [12], [21]. Beside the fact that by transferring learning examples this approach may violate privacy requirements, it suffers from severe scalability limitations. In order to significantly reduce communication overhead, one would have to collect small samples and possibly not cover all sites. This would cause the quality of the resulting decision tree to deteriorate rapidly whenever the data sets of the different sites vary from one another [22]. In contrast, our approach, which is equivalent to collecting all the data, retains the quality of the result and reduces the communication overhead.

A different *metalearning* induction algorithm was suggested in [14]. The algorithm turns each decision tree classifier into a set of rules and then merges the rules into a single superset of rules while resolving conflicts as suggested in [13]. Kargupta et al. [15] describe a *metalearning* algorithm where the local decision tree classifiers are approximated by a set of Fourier coefficients, which are

then collected to a central site where they are combined into a single model. Although the *metalearning* approach is very scalable in terms of performance, the accuracy and the comprehensibility of the metaclassifier drops sharply as the number of remote sites increases. Thus, these methods are not well-suited for large distributed networks.

Much attention was given to the parallelization of induction algorithms. The parallel algorithms described below were intended for a data warehouse environment, where a control environment and high communication bandwidth are assumed, as opposed to a large distributed network where there is no control over the distribution of the data and a normal Internet bandwidth is assumed.

Three parallel algorithms for decision tree induction were described in [10]. In the first algorithm, called synchronous tree construction, all computing nodes construct a decision tree synchronously in depth-first order by exchanging the class distribution information of the local data. Then, they simultaneously compute the gain function, select the best attribute, and split the decision tree node according to this attribute. In the second algorithm, called partitioned tree construction, one (or more) of the computing nodes is responsible for a portion of the decision tree and data is relocated to the responsible computing node after a split. As a result, the responsible computing node can develop this portion of the decision tree independently. The third algorithm uses a hybrid approach: It starts with synchronous tree construction and switches to partitioned tree construction when the number of active leaves in the tree exceeds a given threshold. Thus, at the top of the decision tree, where there are only a few decision tree leaves and data movement is expensive, synchronous tree construction is used. Then, when the number of developed leaves increases, incurring a high communication cost, partitioned tree construction is used. The hybrid algorithm thus aims to minimize the communication overhead between the computing nodes. However, these straightforward algorithms cannot be used in large-scale distributed systems because data movement is often impractical in distributed networks, for the reasons explained above, and because the communication complexity of synchronous tree construction is similar to that of the algorithm described above [20].

The parallel version of SPRINT, also described in [6], enhances the performance of the algorithm by using a vertical partitioning scheme, where every computing node is responsible for a distinct subset of the data attributes. Thus, by dividing the attribute lists evenly among the computing nodes and finding in parallel the best binary conditions of the attributes, the algorithm boosts the performance of the sequential SPRINT algorithm. However, in order to split the attribute lists, the hash table must be available on all computing nodes. In order to construct the hash table, all-to-all broadcast must be performed, making this algorithm highly unscalable. ScalParC [8] improves upon SPRINT with a distributed hash table that efficiently splits the attribute lists, and is communication-efficient when updating the hash table. The same communication pattern is common to all state-of-the-art parallel decision tree induction algorithms [9], [23], [11].

4 BOUNDS ON THE GAIN FUNCTIONS

The bounds given in this section bound the gain function of a population that is the union of several disjoint subpopulations on which only partial information is available. By using them, we can avoid collecting the crosstables of many of the attributes whose gain, as indicated by the bounds, cannot be large enough to change the result.

4.1 Notations

The bounds given below are defined for a single attribute of a single decision tree leaf node. Therefore, we simplify the notations by removing references to the attribute and the decision tree node. Let P be a population of size n and let $\{P_1, P_2\}$ be a partition of P into two subpopulations of sizes n_1, n_2 , respectively. Let the crosstables of populations P_1, P_2, P be defined as:

$$\begin{aligned}\overrightarrow{P_1}(value, class) &= \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}, \\ \overrightarrow{P_2}(value, class) &= \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix}, \\ \overrightarrow{P}(value, class) &= \begin{pmatrix} a_{11} + b_{11} & a_{12} + b_{12} \\ a_{21} + b_{21} & a_{22} + b_{22} \end{pmatrix},\end{aligned}$$

respectively. Here, $a_{i,j}$ and $b_{i,j}$ denote the number of learning examples with value i and class j in $\overrightarrow{P_1}$ and $\overrightarrow{P_2}$, respectively.

We further use $GiniIndex(P)$, $InfoGain(P)$ as the Gini index and information gain functions computed for population P . For brevity, in the following, we include only the lemma we use in the paper. We provide detailed proofs of these lemmas in the Appendix which can be found on the Computer Society Digital Library at <http://computer.org/tkde/archives.htm>.

4.2 Gini Index Function

Lemma 1. For any $GiniIndex(P_1), GiniIndex(P_2), n_1, n_2$, an upper bound on $GiniIndex(P)$ is given by:

$$Upper\ bound = \frac{n_1 GiniIndex(P_1) + n_2 GiniIndex(P_2)}{n_1 + n_2}. \quad (1)$$

Lemma 2. Let P_1, P_2, n_1, n_2 be given. Furthermore, let the candidate binary split decision divide P_1 into two subsets, P_1^{left} and P_1^{right} , with sizes of n_1^{left} and n_1^{right} ($n_1 = n_1^{left} + n_1^{right}$), respectively. Then, a lower bound on $GiniIndex(P)$ is given by:

$$Lower\ Bound = \frac{GiniIndex(P_1)}{\left[1 + \frac{n_2}{n_1}\right] \left[1 + \max\left\{\frac{n_2}{n_1^{left}}, \frac{n_2}{n_1^{right}}\right\}\right]}. \quad (2)$$

4.3 Information Gain Function

Lemma 3. For any $InfoGain(P_1), InfoGain(P_2), n_1, n_2$, an upper bound on $InfoGain(P)$ is given by

$$\frac{n_1 InfoGain(P_1) + n_2 InfoGain(P_2)}{n_1 + n_2}. \quad (3)$$

Lemma 4. Let P_1, n_1, n_2 be given. Furthermore, let the candidate split decision divide P_1 into two subsets, P_1^{left} and P_1^{right} , with size n_1^{left} and n_1^{right} , respectively. Then, a lower bound on $InfoGain(P)$ is given by:

$$\left[\frac{1}{1 + \frac{n_2}{\min\{n_1^{left}, n_1^{right}\}}} \cdot \frac{1}{1 + \frac{n_2}{n_1}} \right] InfoGain(P_1). \quad (4)$$

We summarize our results with the following theorems:

Theorem 1. Let P be a population of size n , and $\{P_1, P_2, \dots, P_k\}$ a partition of P into k subpopulations of sizes n_1, n_2, \dots, n_k , respectively. Let $G()$ denote the gain function (information gain or Gini index). Then, an upper bound on $G(P)$ is given by:

$$G(P) \leq \frac{\sum_{i=1}^k n_i G(P_i)}{\sum_{i=1}^k n_i}.$$

Theorem 2. Let P be a population of size n , and $\{P_1, P_2\}$ a partition of P into two subpopulations of sizes n_1, n_2 , respectively. Assume that the candidate split divides P_1 into two subsets, P_1^{left} and P_1^{right} , with sizes n_1^{left} and n_1^{right} , respectively. Let $G()$ denote the gain function (information gain or Gini index). Then, lower bounds on $G(P)$ is given by:

$$G(P) \geq \frac{G(P_1)}{\left[1 + \frac{n_2}{n_1}\right] \left[1 + \frac{n_2}{\min\{n_1^{left}, n_1^{right}\}}\right]}.$$

5 DISTRIBUTED HIERARCHICAL DECISION TREE

The distributed hierarchical decision tree (DHDT) algorithm runs on a group of computers, connected through a wide-area network such as the Internet. Each computer has its own local database, while the goal of the DHDT is to derive the exact same decision tree learned by a sequential decision tree learner on the collection of all data in the network. We assume a homogeneous database schema for all databases, which can be provided transparently, if required, by ordinary federated system services. The algorithm relies on a (possibly overlay) communication tree that spans all computers in the group. The communication tree can be maintained by a spanning tree algorithm such as Scribe [24] or can utilize the natural hierarchy of the network (see Fig. 1). For reasons of locality, communication between nodes in the lower levels of the spanning tree is often cheaper than communication between nodes in the upper levels. Thus, a ‘‘good’’ algorithm will use more communication at the bottom than at the top of the tree. We further assume that during the growth phase of the decision tree, the databases and the communication tree remain static.

For each database, an entity called *Agent* is allocated, which accesses the database through a standard interface such as SQL or HL7 and gathers simple statistics (the joint distribution of each attribute and the class attribute). For performance consideration, the Agent had better reside on either the same host as the database or on a collocated machine. The Agent is in charge of computing the required

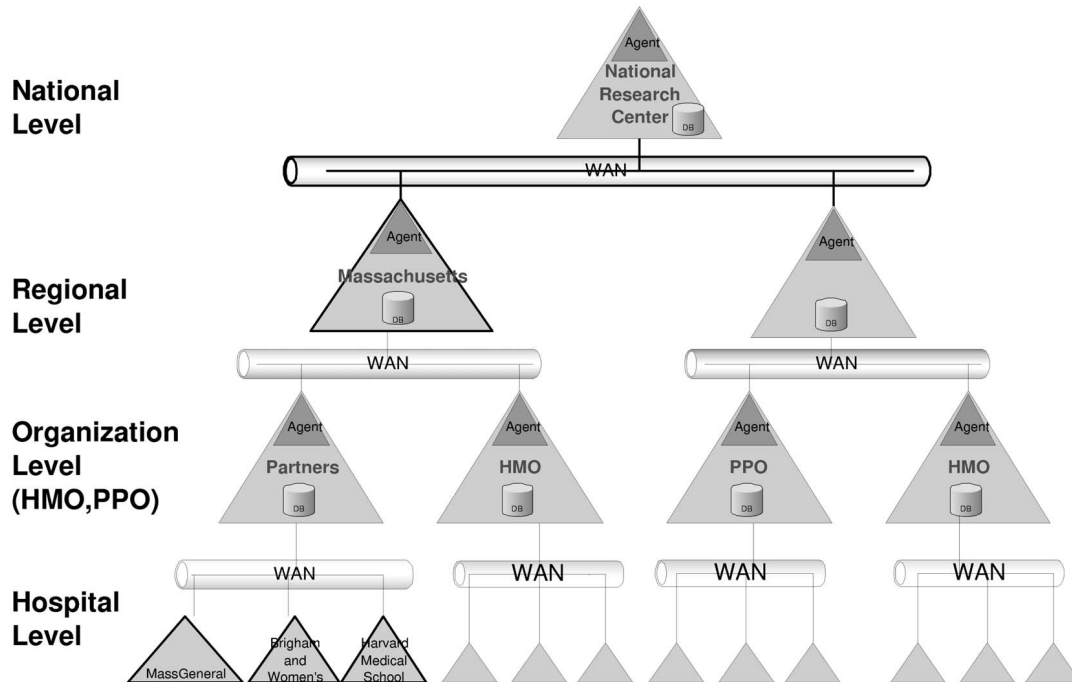


Fig. 1. Hierarchical network of a medical data organization. The network structure in the figure echoes the medical organization's hierarchical structure. Nodes in the lower levels of the hierarchy are hospitals, clinics, medical schools, and research centers. Hospitals and clinics are usually associated with Health Maintenance Organizations (HMOs) such as Partners in the Boston area. Inside these health organizations, medical information is collected independently by each hospital or clinic. However, data is shared on a daily basis between the associated members, to allow patients to get medical care from any one of them. At higher levels of the hierarchy, health organizations collaborate to conduct research on a regional or even national scale.

statistics from the local database and participating in the distributed algorithm. Agents collect statistical data from their children agents and from the local database and send it to their parent agent at its request. All communication is by message exchange.

The root agent is responsible for developing the decision tree and making the split decisions for the new decision tree leaves. First, the root agent decides whether a decision tree leaf has to be split according to one or more stopping conditions (e.g., if the dominance of the majority class has already reached a certain threshold) or according to the PUBLIC method [18], which avoids splitting a leaf once it knows it may be pruned eventually. The class distribution vector, which holds the number of examples that belong to each distinct class in the population, is sufficient for computing these functions and, thus, it is aggregated by the agents over the communication tree to the root agent.

Recall that if a decision tree leaf has to be split, the split must be done by the attribute with the highest gain in the combined database of the entire network. All that is required to decide on the splitting attribute is an agreement as to which attribute has the maximal gain; the actual gain of each attribute does not need to be computed. To reach an agreement, the agents participate in a distributed algorithm called DESAR (Distributed Efficient Splitting Attribute Resolver). For each new leaf that has to be developed, DHDT starts a new instance of DESAR to find the best splitting attribute. The pseudocode for DHDT is given in Algorithm 1:

Algorithm 1: The DHDT Algorithm for the root agent

Initialization

newLeavesList = decision tree root

Algorithm

1. For each $leaf_i$ in newLeavesList do
2. Remove $leaf_i$ from newLeavesList
3. $Attribute_k = \text{run DESAR for } leaf_i$
4. If the gain from splitting $leaf_i$ according to $Attribute_k$ is above the pruning threshold.
5. Split $leaf_i$ by $Attribute_k$
6. Insert new leaves to newLeavesList
7. Endif
8. End

We now proceed to describe the DESAR algorithm.

5.1 Distributed Efficient Splitting Attribute Resolver

To find the best splitting attribute while minimizing communication complexity, DESAR aggregates only a subset of the attribute crosstables over the communication tree to the root agent. The algorithm starts when the agents receive a message from the root that a new leaf has to be developed. Then, each agent waits for messages from its children. When messages are received from all of them, it combines the received crosstables with its own local crosstables, picks the most *promising* splitting attributes on the basis of its aggregated data, and sends to its parent agent only the crosstables for these attributes.

Since different subtrees may choose to send information on different subsets of the attributes to their root, the information eventually collected by the root does not always suffice to decide which attribute maximizes the

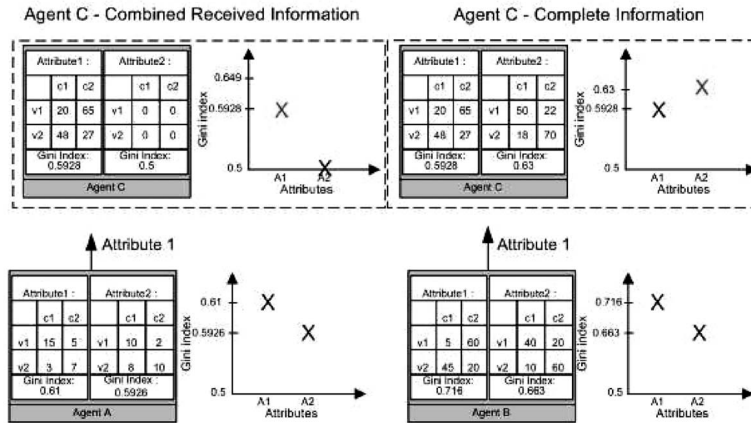


Fig. 2. Example of a wrong decision based on partial statistics. Agent A and Agent B are the children of Agent C. After computing the gain for each of its attributes, Agent A sends only the crosstable of its best splitting attribute, Attribute 1. Similarly, Agent B sends only the crosstable of Attribute 1. Agent C combines these crosstables and chooses Attribute 1 as its best splitting attribute. However, the correct decision is to pick Attribute 2, which has the highest gain in the combined data set.

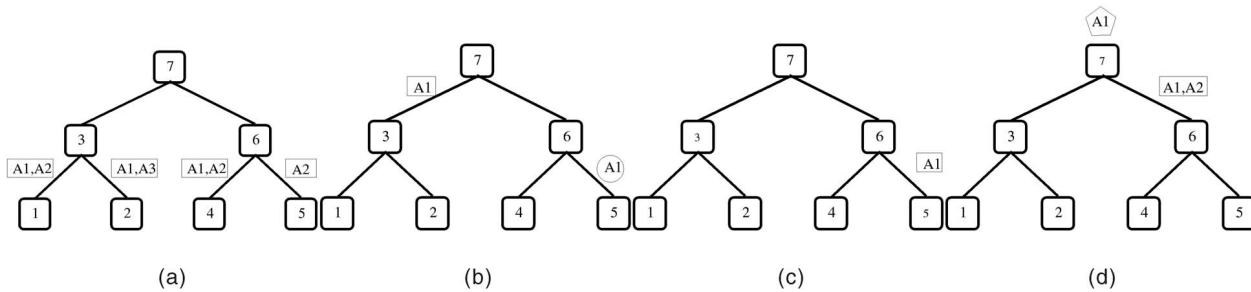


Fig. 3. An overview of DESAR core. In (a), all leaf Agents send their promising attributes to their parents. (b) shows that for Agent 3 the data was sufficient to determine A1 as its promising attribute and, thus, send it to the root. However, Agent 6 requests more data on A1 from Agent 5. In (c), Agent 5 responds. Finally, in (d), Agent 6 sends both A1 and A2 as promising attributes to the root, which finds this data sufficient to split the node according to A1.

gain function. An attribute may have high gain in the part of the tree which sent it to the root, but the gain may drop sharply when data is collected from other parts of the tree. The opposite is also possible, as can be seen in Fig. 2. To overcome this, we use the upper and lower bounds discussed in the previous section and compute for each attribute an *attribute interval*, rather than a single, possibly erroneous value. The interval between the lower bound and the upper bound for the gain of an attribute, based on the known data is denoted the *attribute interval*.

The details of computing the bounds can be found in the next section. The bounds are computed using the information received from all of the agent’s children. Thus, these bounds bound the gain function over the data in the network subtree. In particular, the bounds computed by the root agent bound the gain function over all the data in the network.

Using the notion of attribute intervals, we say that a given threshold defines a *clear separation* of intervals if it separates the attribute intervals into two nonempty disjoint sets of intervals and neither of the intervals crosses the threshold (see Fig. 4).

When the bounds are computed, the agent sets a threshold, denoted *border*, with a minimal number of attributes having their lower bounds larger than the border. If the border obtained defines a clear separation, the attributes whose intervals lay above the border are called *promising*, and their crosstables are sent to the agent’s parent. If a clear separation is not achieved, or if the number of attributes whose interval is above the border is too high,

the agent collects more information from its descendants using request methods which we explain below. Notice that the root agent is special in that it requires, in addition to clear separation, that only a single interval remains above the border. Only if this additional requirement is met can the root agent safely decide to split the leaf according to the attribute whose interval is above the border.

Fig. 3 outlines the possible flow of the DESAR algorithm on a small hierarchy of seven agents. Fig. 5 describes a specific scenario in which data is received by an agent, more data is requested, and a clear separation is achieved.

The simplest way an agent can request more information from its children is to name the attributes for which more information is needed. We call this request method the *naming* method. When a child receives a request for a specific attribute, and if the crosstable of this attribute was




Fig. 4. Example for clear separation. In the above figure, the intervals of five attributes are displayed and a threshold with a value of 0.82 defines a clear separation between the intervals.

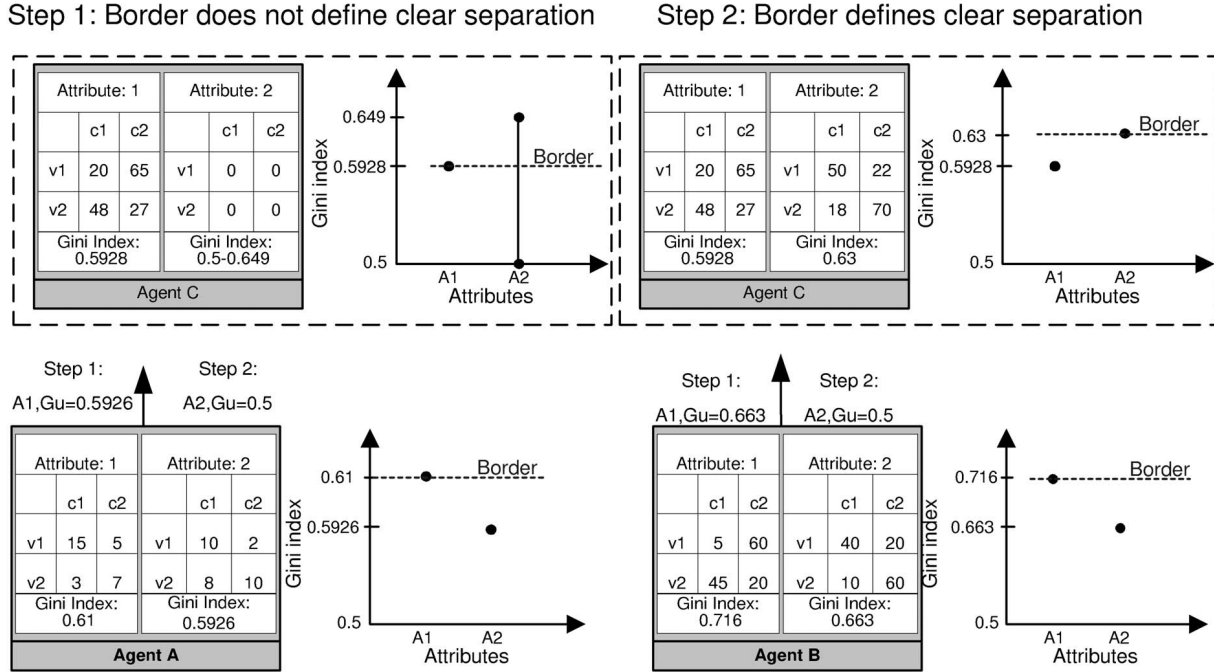


Fig. 5. Two-step example for the DESAR algorithm. Agent A and Agent B are the children of Agent C. Once Agent A and Agent B have computed all attribute intervals, they define the border and their promising attribute sets, which for both agents include only attribute 1. When Agent C receives the messages, it computes its own attribute intervals and discovers that the border does not define a clear separation. Therefore, it requests more information regarding Attribute 2 from its children. When replied, Agent C recomputes the attribute interval of Attribute 2, and this time the border defines a clear separation where the only promising attribute is Attribute 2.

not yet sent to its parent, the child immediately replies by sending its crosstable. If the crosstable was not yet received from the subtree, it first requests that its children send more information regarding this attribute and forwards the information once it arrives. Note that the less information an agent collects from its subtree on an attribute, the larger the interval will be. Therefore, if the naming method is repeated, all information for the attributes whose intervals cross the border will eventually be collected to the requesting parent agent, the lower and upper bounds will be equal to the accurate gain in the agent's combined database, and a clear separation will be defined.

In Section 5.3, we describe an additional request method. We also describe a strategy, aimed at reducing communication complexity, for determining whether this method or the naming method should be used.

Algorithm 2 describes DESAR pseudocode, uniformly executed by all agents.

Algorithm 2: DESAR Algorithm

Definitions

- D1. $border$ = maximal lower bound of all attributes which were not sent to the parent
- D2. $borderAttribute$ = the attribute whose lower bound defines the border
- D3. If agent is root then
- D4. $ExtraCondition$ = There is only a single attribute A_i where $UpperBound(A_i) \geq border$ or $max_i(UpperBound(A_i)) = border$
- D5. Else
- D6. $ExtraCondition = G_u^i < border$ for all children

Algorithm

On initialization of a new leaf is born

01. Receive information from all children
 02. While (not ($border$ defines a clear separation and $ExtraCondition$)) do
 03. If ($G_u^i > border$) then
 04. request $child_i$ to lower its border and send new information
 05. Else if ($border$ does not define a clear separation and crosstable of $borderAttribute$ has only partial information)
 06. request information for $borderAttribute$ from children who did not send complete information
 07. Else
 08. request information for all attributes that cross the $border$
 09. End if
 10. Receive information from all children
 11. End while
 12. Return attributes A_i where $LowerBound(A_i) \geq border$
- ##### On a request for more information from parent
01. If (parent requires more information for attribute $attr_i$) then
 02. If (crosstable of $attr_i$ was not sent to parent) then
 03. Send parent the crosstable of $attr_i$
 04. Else
 05. request information for $attr_i$ from children who sent partial information regarding $attr_i$
 06. Else (the case where parent requests that the border be lowered)
 07. Update $border$ and $borderAttribute$ and start phase 1.
 08. Endif

5.2 Computing Lower and Upper Bounds on the Gain of Attributes

As we did in Section 4, we simplify the notations here by removing indices to the attribute and the decision tree node.

Let $agent_1, \dots, agent_q$ be the descendants of $agent_0$. Additionally, let $child_i$ denote the i th immediate child of $agent_0$. Let P_d be the population of $agent_d$ derived from its local database. The combined population for accurate computation of the gain function for the network tree rooted in $agent_0$ is defined by $P = \bigcup_d P_d$.

Without loss of generality, for any given attribute, let $P_U = \bigcup_{d=k+1..q} P_d$ be the combined population of the descendant agents which did not yet send the attribute's crosstable to $agent_0$, and let $P_K = P/P_U$. Furthermore, let P_K^i be the combined population of the descendant agents who did submit the attribute's crosstable and are also descendants of $child_i$ (including $child_i$ itself), and let P_U^i be the combined population of the descendant agents who did not submit the attribute's crosstable and are also descendants of $child_i$. Finally, let $G()$ denote the gain function used by the algorithm and let $|X|$ denote the size of the set X .

5.2.1 Upper Bound

First, the agent computes an upper bound G_U on $G(P_U)$. This bound is computed recursively, where each child agent computes and sends to its parent an upper bound, denoted G_U^i , on its contribution to population P_U . G_U^i is computed by the following recursive rule: If the attribute's crosstable is not sent to the parent, G_U^i is equal to the attribute's upper bound. Otherwise, G_U^i is equal to G_U of the child itself. Note that for the leaf agents $P_U = \emptyset$ and, thus, G_U is set to 0.

Then, by applying Theorem 1, G_U is:

$$G_U \geq G(P_U), \text{ where } G_U = \frac{\sum_i |P_U^i| G_U^i}{\sum_i |P_U^i|}. \quad (5)$$

Now, by applying Theorem 1 again, the agent computes the upper bound as follows:

$$G(P) \leq \frac{|P_K| G(P_K) + |P_U| G_U}{|P|}. \quad (6)$$

Note that the size of the combined database, $|P|$, can be computed from the aggregated class distribution vector and, thus, $|P_K|$, $|P_U^i|$ can easily be computed.

Finally, in order to further reduce communication complexity and make it independent of the number of candidate attributes, a child agent sends the maximal G_U^i of all attributes as a single upper bound (denoted G_U^i) for all of them.

5.2.2 Lower Bound

The lower bound is trivially computed by Theorem 2, where $P_1 = P_K$ and $P_2 = P_U$.

5.3 Efficient Request Methods

The disadvantage of the *naming* method is in the way the bounds are computed: Recall that the parent receives from its $child_i$ a single upper bound, G_U^i , which bounds, for all attributes, the possible contribution to the gain function of all the data beneath $child_i$, which was not sent up. Following (5) and (6), the upper bound of every attribute

is partially based on G_U^i . If G_U^i is high, the weighted upper bounds of many attributes (which are computed using G_U^i and Theorem 1) will be higher than the border, and since their lower bounds remain low, their attribute intervals will cross the border. Consequently, a request that names many crossing attributes causes high communication overhead.

To overcome this, a different request method, called the *independent* method, is used. The new request method asks the child to lower its border independently of its parent. When a child receives this request, it sets its border as the maximal lower bound of all attributes that were not sent up. Then, it tries again to find a clear separation, if necessary, by requesting more information from its children. Consequently, new information is sent to the requesting parent, and the upper bound G_U^i of the child is reduced.

Finally, to minimize the overall communication complexity, DESAR employs the following strategy when using the request methods: If G_U^i of $child_i$ is above the border, the *independent* method is used. Otherwise, if a clear separation does not exist and the highest attribute (i.e., the attribute with the highest lower bound) has partial information, the child uses the *naming* method to request information, for the highest attribute only, from all children who sent partial information regarding this attribute. This is done in the hope that the new information will raise the border and a clear separation will be achieved. If the highest attribute already has full information, i.e., the lower and upper bound of the attribute are equal, then the agent will use the *naming* method to request more information for the attributes that cross the border.

5.4 Reducing Message Complexity with the Hoeffding Bound

In the DESAR algorithm, each Agent sends to its parent attributes whose intervals are above a border which it calculates based on data from its subtree. The parent may later request that the Agent sends additional data, regarding other attributes. One cause for such requests for more data may be slight differences between the data in different subtrees, which cause the gain of an attribute in the subtree database to be exceptionally higher (or lower) for that agent comparing to the global database. Since requests for more data are wasteful, it is beneficial to sometimes send a small number of additional attributes in the first place so as to avoid them.

For this purpose, we enhance DESAR by further lowering the border by a small constant ε (see Fig. 6). We calculate ε using the well-known *Hoeffding* bound [25] in a way that resembles its usage in the VFDT algorithm [26]. We treat the gain of an attribute in a subtree database as a random variable r calculated over an independent sample of size n . The Hoeffding bound then predicts that with probability $1 - \delta$ the actual value of r (i.e., the gain over the whole population) lays within

$$\varepsilon = \sqrt{\frac{R^2 \ln(\frac{1}{\delta})}{2n}}$$

of the value calculated over the sample (R , the range of r , equals one in the case of the gain functions we use). By tuning δ , we can now trade the number of attributes sent in




Fig. 6. Redefinition of *border* for reducing the average number of messages sent. In the above figure, the interval of five attributes are displayed. the border computed by (7) defines three promising attributes: A2, A3, and A5. Only the crosstables of these attributes are sent to the parent again.

the first message against the probability that additional data will be required.

A further complication arises from the fact that DESAR does not really calculate the value of r , but rather bounds it from below. Since the real value is guaranteed to lay above the border, reducing the border by ε means that we may sometimes overestimate δ and send too many attributes. This would not happen if DESAR collects full information on the most prominent attribute, as it usually does.

We conclude that the new definition of *border* is:

$$border = \max_{i \in \text{attributes}} \{LowerBound_i\} - \sqrt{\frac{\ln(\frac{1}{\delta})}{2n}}. \quad (7)$$

5.5 The DESAR Correctness Proof

We now prove that the DESAR algorithm terminates and that it chooses the same attributes that the sequential decision tree algorithm would choose when working on the combined database of the complete network.

The referenced line numbers in the proof refer to the pseudocode of DESAR provided in Algorithm 2.

Proof.

- **Correctness.** Assume that the gain of another attribute, say A_k , is higher in the combined database than the gain of the attribute returned by DESAR algorithm at the root agent, A_{best} . Clearly, this means

$$LowerBound(A_{best}) < UpperBound(A_k).$$

Since A_{best} , by definition, has the highest lower bound,

$$LowerBound(A_{best}) > LowerBound(A_k).$$

This means that the border, which is set to $LowerBound(A_{best})$, must cross A_k 's attribute interval. Hence, there is no clear separation and DESAR does not terminate.

- **Termination.** The algorithm does not end until the condition tested in Line 2 is satisfied. Assume that the algorithm has made N rounds of requesting more information, where N is large enough so that all crosstables from all the descendants have arrived at the root agent. In this case, all the intervals will have size zero since

the gain for the combined database can be computed accurately. Clearly, the condition tested in Line 2 is satisfied and the algorithm terminates. \square

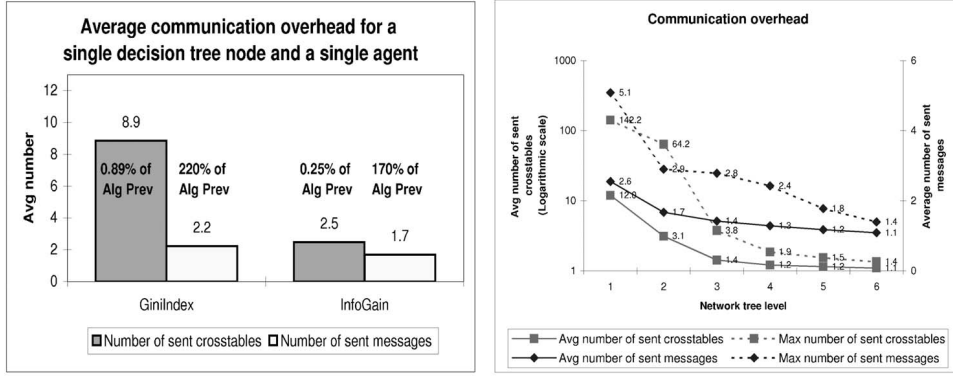
6 EXPERIMENTAL EVALUATION

The DHDT algorithm is designed to run on data sets with a large number of attributes, such as the genomically enriched EMR. However, data which include both genotype and phenotype information is scarce. Therefore, we adopted an approach common in bioinformatics studies, which is to try and predict part of the genotype using the rest of the genotype. This approach can be justified in the case of phenotypical patterns that are correlated strongly with a single SNP variation or an increased activity of a single gene. In such cases, learning a model which predicts an SNP's allele, for instance, is equivalent to learning a model which predicts one of these diseases.

Another question regarding the data is whether to use real or synthetic data. Real data sets are hard to come by and even harder to merge with one another, mainly because they were usually generated in independent experiments and under different experimental conditions. For instance, data sets of microarray gene expression exhibit two types of variations, variation in microarray technologies and variations in different expression levels of significant genes in various cancer types. In the genomically enriched EMR application we consider, all data would be generated using a standard test and under standard conditions. Thus, we opt for synthetic data created by two accepted data generators ([27], [28]) using two different theoretical models. The performance of the algorithm is almost the same regardless of what generator is used, which indicates that it takes advantage of characteristics of genetic data rather than on characteristics of the data generator.

Each data set contains 250,000 examples describing a single population. A single SNP is described by a binary attribute where "0" denotes the most common allele. An example is composed of 1,000 SNPs. An arbitrary SNP is designated the class attribute. The experiments were performed on a simulation of a communication tree that spans all agents in the system. At the beginning of each experiment, each agent builds its local database by sampling a small fraction of the simulated data set, thus emulating a specific subpopulation.

The DHDT algorithm computes exactly the same result as would be computed by ID3 had it been given the entire data set. For this reason the focus of our experiments is not on evaluating the accuracy of the resulting model, but rather on measuring the reduction in communication overhead. The value of ID3 generated classifiers has been validated over and over in the last two decades on far more data sets than we can possibly cover in this context. To evaluate DHDT performance we compare the overhead it generates to the calculated overhead of previously suggested decision tree algorithms [20], [10], which both collect all the crosstables for all the available attributes to a single agent. Henceforth, we denote these two algorithms as *Algorithm Prev*. *Algorithm Prev* sends $O(LN)$ bytes and $O(L)$ messages per decision tree node, where L is the number of agents and N is the number of attributes in the



(a) (b)

Network spanning tree degree	3
Levels of hierarchy	6
Network size	$\sum_{i=1}^6 3^i = 1093$
Subpopulation size	5000
Total number of attributes	1000
Average number of decision tree nodes (Gini index)	25
Average misclassification rate	3%

(c)

Fig. 7. Average communication overhead for a single split decision. (a) Average communication overhead, comparing to *Alg. prev* which sends 1,000 attributes in a single message. (b) Hierarchical view of the communication overhead (Gini index).

data set. Note that DHDT generates the exact same model *Algorithm Prev* would (which is the same one that which would be computed by a sequential algorithm that accesses all the data). Thus, there is no point in comparing the algorithms' accuracy because it depends only on the data and not of the choice of algorithm.

6.1 Synthesized Genomic Data

We have synthesized two SNP data sets, using two simulations that follow different theoretical models. The first simulation is based on coalescent theory and was conducted with the Hudson simulation engine [27]. The simulation assumes intragenic recombination. Whenever a recombination event occurs, the two separated strands of a DNA sequence become statistically independent, consequently lowering the association between the SNPs of the separated strands. The parameters for this simulation are the number of examples, the number of sites (SNPs), and the probability that recombination will take place. In this simulation, we used a typical recombination probability of 10^{-8} per generation for every 2,500 base pairs.

The second simulation follows a recently developed theory that assumes the existence of DNA blocks. The theory proposes that recombination events occur in narrow hot-spots, resulting in the creation of DNA blocks in the region between two hot-spots. Therefore, SNPs that are contained in a single DNA block are more strongly associated. The simulation was conducted using the simulation tool suggested in [28] with default parameters.

The results of the experiments showed that the communication overhead of the algorithm is the same for both databases. Thus, we present below just one set of results.

6.2 Experiments

Our first experiment measures the average communication overhead of a single split decision (i.e., a single run of the DESAR algorithm) in terms of the number of messages and the number of sent crosstables. These results are compared with previous distributed decision tree algorithms which collect and aggregate the crosstables of all attributes.

Our algorithm demonstrates an average reduction of more than 99 percent in the number of transmitted bytes, with only a small increase in the average number of sent messages (1.2 per Agent per decision tree node). These results are summarized in Fig. 7a. Fig. 7b provides a detailed view of the communication overhead over the levels of the network spanning tree when using the Gini index function. Similar results are achieved for the information gain function. Note that most of the communication takes place in the lower levels of the tree and decreases in the higher levels. Because more data is used to compute the gain function in the higher levels, there is less chance that the best attribute in a child agent will not be the best attribute in its parent agent as well. Recall that, for locality reasons, the communication in higher levels is more expensive and, thus, the DHDT algorithm is able to amortize the cost of communication over the network. In addition, the benefits of the algorithm are emphasized because the maximal communication overhead (out of 20 tests) declines similarly to the average communication overhead, and is sharply reduced in the higher and more costly levels of the network tree.

The above experiment also compares the communication overhead when using the information gain and the Gini index functions. The results show that the information gain function is more efficient communication-wise than the Gini index function. In [29], the authors compared the split

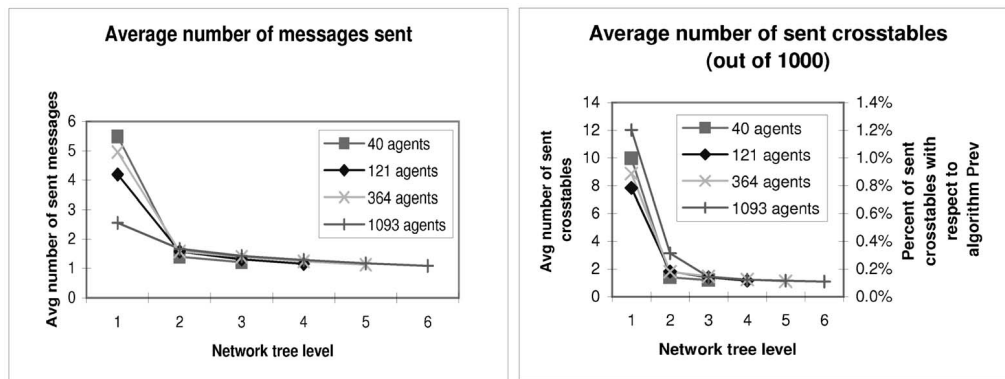


Fig. 8. Scalability in network size. The above figures show the distribution of the average communication overhead over the network tree levels for different network sizes (Gini index). Other experiment parameters remained as described in Fig. 7.

decisions made by the Gini index function with those made by the information gain function. Their results showed a disagreement only on 2 percent of the decisions. However, if different splitting attributes are chosen for a node that resides at the top of the decision tree, then obviously the subtrees below this node will be completely different. Consequently, the effect of the different decision is increased.

Our next experiments examine the scalability of the algorithm with respect to the size of the network, the number of total attributes, and the size of the local databases:

- Scalability in the size of the network.** We examined the communication overhead with an increasing network size (40, 121, 364, and 1,093 agents), where in each step, a new level is added to the network spanning tree and the other parameters remain as in Fig. 7. The results, summarized in Fig. 8, show that the reduction in communication overhead is hardly affected by the network size. Therefore, the algorithm can be deployed on large-scale networks containing even thousands of network nodes and can sharply reduce communication overhead.
- Scalability in data set dimensionality.** The algorithm is intended for highly dimensional data sets. Therefore, we tested the reduction in communication overhead when the number of attributes in the data sets increases. We conducted tests with 125, 250, 500, and 1,000 attributes (SNPs), with a longer segment of the chromosome used each time in order to increase the number of SNPs in our data set. Our results, summarized in Fig. 9, show that when the number of attributes is doubled, only a few additional crosstables are sent up, and the percentage of the sent attributes declines sharply. This behavior is expected due to both the following properties of the SNP data and the DESAR algorithm: SNPs that are near each other on a chromosome are more strongly associated than SNPs that are far away from one another. Thus, as the number of SNPs increases, only a few of the additional SNPs are good predictors of the target SNP (the class attribute). Because the DESAR algorithm sends only attributes that provide good

prediction of the target attribute, only a few additional crosstables are sent up the hierarchy, and the communication overhead remains low.

- Scalability in the size of local databases.** In this experiment, we examined the effect of larger local data sets on communication overhead. As expected, the results in Fig. 10 show that when the size of the local data sets increases, communication overhead decreases because the sampling noise decreases.

Our next experiment examines the optimizations described in Section 5.4. We examine how tuning the δ parameter affects communication overhead. Recall that as δ increases, the border is fixed ever further below the maximal lower bound. Therefore, more crosstables are sent in a child's first message to its parent. Consequently, fewer requests for more information are initiated by the parent, which in turn reduces the number of messages. We performed tests with $\delta = 0, 0.1, \text{ and } 0.2$. Fig. 11 shows that the number of messages decreased at the cost of a higher number of sent attributes. In a detailed view that shows the distribution of the communication overhead (Fig. 12), we see that the optimization mostly effects communication in the lower levels of the network tree, where the number of

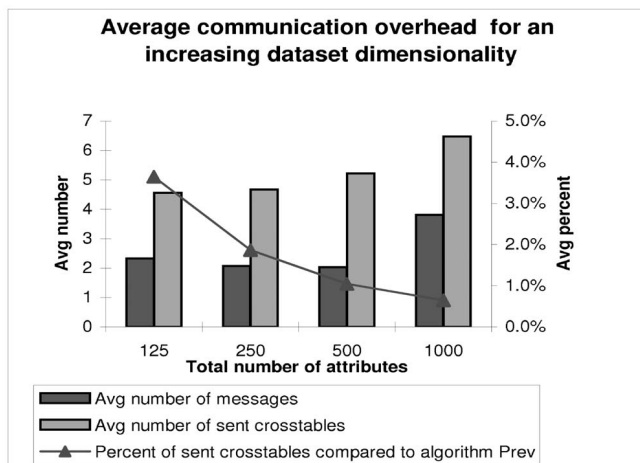


Fig. 9. Scalability in data set dimensionality. The above figure shows the communication overhead for a network of 364 agents, where the number of attributes increases. Other experiment parameters remained as described in Fig. 7.

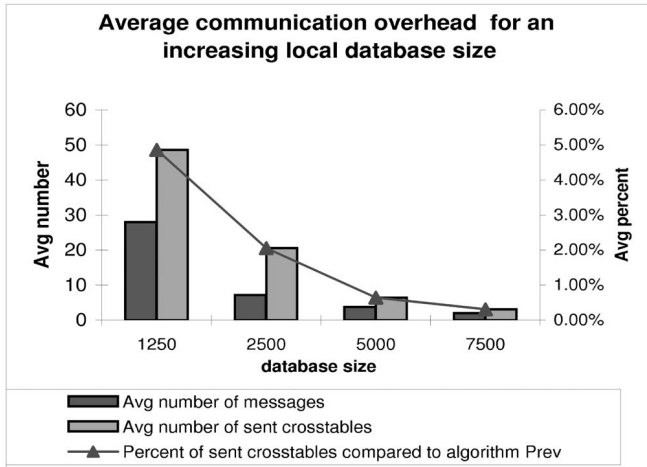


Fig. 10. Effect of smaller databases. When the size of the local databases decreases, the data becomes more noisy. Roughly, the number of message and of crosstables sent is inverse proportional to the number of examples in each database.

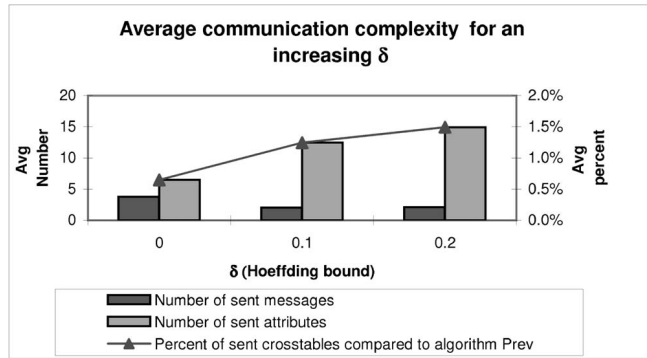


Fig. 11. Optimizing the number of messages. The above figure shows the effect of the optimization described in Section 5.4 on communication overhead for a network of 364 agents. Other experiment parameters remained as described in Fig. 7.

sent messages is effectively reduced, while communication in the upper levels, which already uses an almost minimal number of messages, is hardly affected. This desired behavior is achieved even for small values of δ .

In the experiment described in Fig. 10 we decreased the size of the local database and by that increased the level of “White” noise. In reality, it is often the case that noise stems

from measurement and experimental errors. Such error source is characterized not by “White” noise but rather by outliers—examples that are by far different from the normal population. We generated outliers by agents that contain data generated by splitting chromosomes in the middle and crossing the two halves. Since the data of these agents is so different from the norm, the attributes considered promising by most agents will not seem promising to these agents. Thus, the statistics for the promising attributes will be missing some data. Thus, this experiment measures the effectiveness of our lower bound.

Comparing our lower bound with a trivial lower bound of 0.5, Fig. 13 shows that our lower bound does save part of the communication in the lower levels of the communication tree—where most of the communication is done. The lower bound is less effective at higher levels. This is explained by the larger populations accumulated at higher levels. These larger populations mean that the statistic for the nonaltered data is more accurate and, thus, it is easier for the algorithm to overcome the outliers, even without the help of the lower bound.

7 CONCLUSIONS

Whereas prior decision tree algorithms have had to send statistics for every attribute in the data set in order to make a correct decision, our algorithm sends statistics for only a fraction of the attributes, while eliminating most of the communication overhead. Thus, it does not require the high network bandwidth required by the earlier algorithms—bandwidth that clearly does not exist in wide-area networks. Furthermore, it continues to perform well as the size of the network or the number of attributes increases. Therefore, our algorithm is well-suited for mining large-scale distributed systems with highly dimensional data sets, and especially beneficial for the medical information domain, where clinical and genomic data are distributed across hospital databases and other medical facilities.

In further research, we intend to generalize the basic approach of DHDT, bounding a statistics rather than computing it and computing the bounds hierarchically, to other data mining tasks.

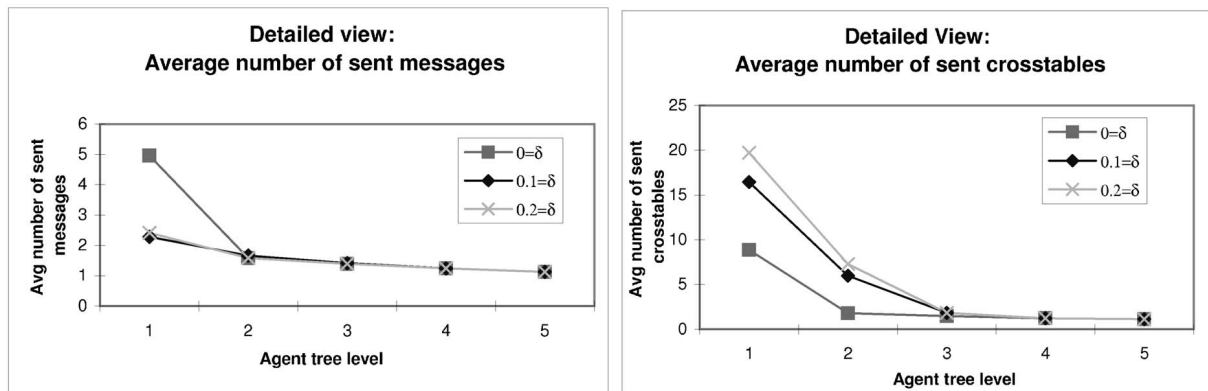


Fig. 12. Hierarchical view of the average communication overhead when the number of messages is optimized.

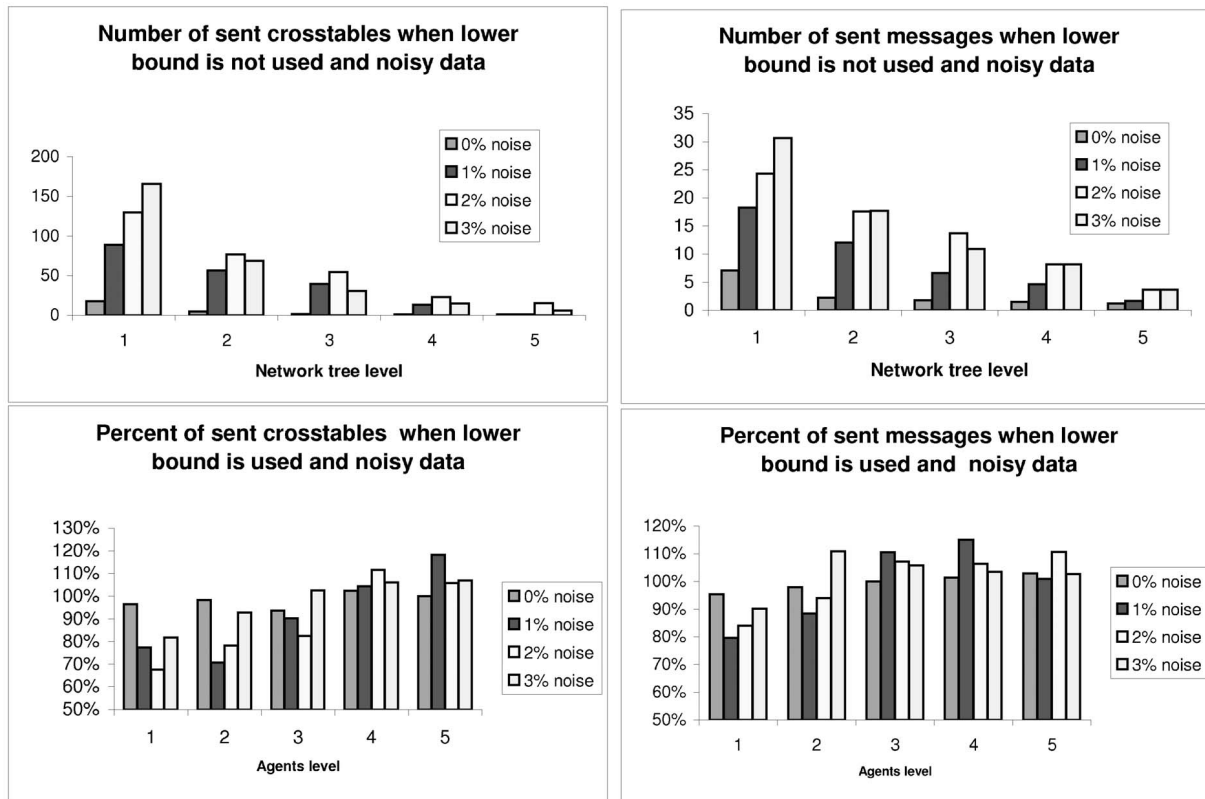


Fig. 13. The effect of the lower bound over the communication overhead with noisy training data.

REFERENCES

- [1] W. Stlinger, O. Hogl, H. Stoyan, and M. Muller, "Intelligent Data Mining for Medical Quality Management," *Proc. Fifth Workshop Intelligent Data Analysis in Medicine and Pharmacology (IDAMAP-2000), Workshop Notes of the 14th European Conf. Artificial Intelligence (ECAI-2000)*, pp. 55-67, 2000.
- [2] N.J. Risch, "Searching for Genetic Determinants in the New Millennium," *Nature* 405, pp. 847-856, 2005.
- [3] F.-C. Tsui, J.U. Espino, V.M. Dato, P.H. Gesteland, J. Hutman, and M.M. Wagner, "Technical Description of Rods: A Real-Time Public Health Surveillance System," *J. Am. Medical Informatics Assoc. (JAMIA)*, vol. 10, no. 5, pp. 399-408, Sept./Oct. 2003.
- [4] M.M. Wagner, J.M. Robinson, F.-C. Tsui, J.U. Espino, and W.R. Hogan, "Design of a National Retail Data Monitor for Public Health Surveillance," *J. Am. Medical Informatics Assoc. (JAMIA)*, vol. 10, no. 5, pp. 409-418, Sept./Oct. 2003.
- [5] <http://www.hl7.org>, 2005.
- [6] J.C. Shafer, R. Agrawal, and M. Mehta, "SPRINT: A Scalable Parallel Classifier for Data Mining," *Proc. 22nd Int'l Conf. Very Large Databases for Data Mining*, 1996.
- [7] M. Mehta, R. Agrawal, and J. Rissanen, "SLIQ: A Fast Scalable Classifier for Data Mining," *Proc. Fifth Int'l Conf. Extending Database Technology*, 1996.
- [8] M.V. Joshi, G. Karypis, and V. Kumar, "A New Scalable and Efficient Parallel Classification Algorithm for Mining Large Datasets," *Proc. Int'l Parallel Processing Symp.*, 1998.
- [9] K. Alsabti, S. Ranka, and V. Singh, "CLOUDS: A Decision Tree Classifier for Large Datasets," *Knowledge Discovery and Data Mining*, pp. 2-8, 1998.
- [10] A. Srivastava, E.-H.S. Han, V. Kumar, and V. Singh, "Parallel Formulations of Decision-Tree Classification Algorithms," *Data Mining and Knowledge Discovery: An Int'l J.*, vol. 3, pp. 237-261, 1999.
- [11] R. Jin and G. Agrawal, "Communication and Memory Efficient Parallel Decision Tree Construction," *Proc. Third SIAM Int'l Conf. Data Mining (SDM)*, 2003.
- [12] P.K. Chan and S.J. Stolfo, "Toward Parallel and Distributed Learning by Meta-Learning," *Working Notes AAAI Work. Knowledge Discovery in Databases*, pp. 227-240, 1993.
- [13] F.J. Provost and D.N. Hennessy, "Scaling Up: Distributed Machine Learning with Cooperation," *Proc. 13th Nat'l Conf. Artificial Intelligence*, 1996.
- [14] L.O. Hall, N. Chawla, and K.W. Bowyer, "Combining Decision Trees Learned in Parallel," *Proc. Distributed Data Mining Workshop at the Int'l Conf. Knowledge Discovery and Data Mining*, 1998.
- [15] H. Kargupta, B. Park, D. Hershberger, and E. Johnson, "Collective Data Mining: A New Perspective Toward Distributed Data Mining," *Advances in Distributed and Parallel Knowledge Discovery*, 1999.
- [16] E.B. Hunt, J. Marin, and P.T. Stone, *Experiments in Induction*. Academic Press, 1966.
- [17] J.R. Quinlan, "Induction of Decision Trees," *Machine Learning*, vol. 1, pp. 81-106, 1986.
- [18] R. Rastogi and K. Shim, "PUBLIC: A Decision Tree Classifier That Integrates Building and Pruning," *Data Mining and Knowledge Discovery*, vol. 4, no. 4, pp. 315-344, 2000.
- [19] L. Breiman, J. Friedman, R. Olshen, and C. Stone, *Classification and Regression Trees*. Monterey, Calif.: Wadsworth and Brooks, 1984.
- [20] D. Caragea, A. Silvescu, and V. Honavar, "A Framework for Learning from Distributed Data Using Sufficient Statistics and Its Application to Learning Decision Trees," *Int'l J. Hybrid Intelligent Systems*, invited paper, 2003.
- [21] S.J. Stolfo, A.L. Prodromidis, S. Tselepis, W. Lee, D.W. Fan, and P.K. Chan, "JAM: Java Agents for Meta-Learning over Distributed Databases," *Knowledge Discovery and Data Mining*, pp. 74-81, 1997.
- [22] J. Catlett, "Megainduction: Machine Learning on Very Large Databases," PhD dissertation, Univ. of Sydney, 1991.
- [23] J. Gehrke, V. Ganti, R. Ramakrishnan, and W.-Y. Loh, "BOAT—Optimistic Decision Tree Construction," *Proc. ACM SIGMOD Int'l Conf. Management of Data*, 1999.
- [24] M. Castro, P. Druschel, A. Kermarrec, and A. Rowstron, "Scribe: A Large-Scale and Decentralized Application-Level Multicast Infrastructure," *IEEE J. Selected Areas in Comm.*, vol. 8, p. 20, 2002.
- [25] W. Hoeffding, "Probability Inequalities for Sums of Bounded Random Variable," *J. Am. Statistical Assoc.*, vol. 58, pp. 13-30, 1963.
- [26] P. Domingos and G. Hulten, "Mining High-Speed Data Streams," *Proc. Sixth ACM SIGKDD Int'l Conf. Knowledge and Discovery and Data Mining*, pp. 71-80, 2000.

- [27] R.R. Hudson, "Generating Samples Under a Wright-Fisher Neutral Model of Genetic Variation," *Bioinformatics*, vol. 18, pp. 337-338, 2002.
- [28] G. Greenspan and D. Geiger, "Model-Based Inference of Haplotype Block Variation," *RECOMB*, pp. 131-137, 2003.
- [29] L. Raileanu and K. Stoffel, "Theoretical Comparison between Gini Index and Information Gain Criteria," *Annals of Math. and Artificial Intelligence*, vol. 41, no. 1, pp. 77-93, May 2004.



Amir Bar-Or received the BSc and MS degrees from the Computer Science faculty of the Technion, Haifa, Israel. He is currently a researcher at HP Cambridge Research Lab. Previously, he was a member of HP Haifa Research Lab for 5 years, developing advanced features for the Non-Stop Tandem database. His current research interests include distributed/parallel computing, stream computing, data management, and data mining.

Daniel Keren (<http://www.cs.haifa.ac.il/~dkeren>) received his PhD degree from the Department of Computer Science in the Hebrew University in 1991. He spent three years doing postdoctoral work at Brown University, and since 1995 he has been with the Department of Computer Science at Haifa University, Israel. His main research interests include machine learning and computer vision.



Assaf Schuster (<http://www.cs.technion.ac.il/~assaf>) received the BSc, MSc, and PhD degrees in mathematics and computer science from the Hebrew University of Jerusalem, Israel, in 1984, 1986, and 1991, respectively. Since being awarded his PhD degree, he has been with the Computer Science Department at the Technion—The Israel Institute of Technology. His interests include all aspects of parallel and distributed computing.



Ran Wolff (<http://www.cs.technion.ac.il/~ranw>) received the BA and PhD degrees in computer science from the Technion—Israel Institute of Technology, in 1996 and 2004, respectively. Currently, he is holding a postdoctoral position with the University of Maryland at Baltimore County. His expertise is in large-scale and high-performance data mining. He has authored several papers on data mining in grid systems, peer-to-peer networks, sensor networks, and other distributed environments.

▷ **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.**