

Distributed Anonymization: Achieving Privacy for Both Data Subjects and Data Providers

Pawel Jurczyk and Li Xiong

Emory University, Atlanta GA 30322, USA

Abstract. There is an increasing need for sharing data repositories containing personal information across multiple distributed and private databases. However, such data sharing is subject to constraints imposed by privacy of individuals or data subjects as well as data confidentiality of institutions or data providers. Concretely, given a query spanning multiple databases, query results should not contain individually identifiable information. In addition, institutions should not reveal their databases to each other apart from the query results. In this paper, we develop a set of decentralized protocols that enable data sharing for horizontally partitioned databases given these constraints. Our approach includes a new notion, *l-site-diversity*, for data anonymization to ensure anonymity of data providers in addition to that of data subjects, and a distributed anonymization protocol that allows independent data providers to build a virtual anonymized database while maintaining both privacy constraints.

1 Introduction

Current information technology enables many organizations to collect, store, and use various types of information about individuals in large repositories. Government and organizations increasingly recognize the critical value and opportunities in sharing such a wealth of information across multiple distributed databases.

Problem scenario. An example scenario is the Shared Pathology Informatics Network (SPIN)¹ initiative by the National Cancer Institute. The objective is to establish an Internet-based *virtual* database that will allow investigators access to data that describe archived tissue specimens across multiple institutions while still allowing those institutions to maintain local control of the data. There are some important privacy considerations in such a scenario. First, personal health information is protected under the Health Insurance Portability and Accountability Act (HIPAA)^{2,3} and cannot be revealed without de-identification

¹ Shared Pathology Informatics Network. <http://www.cancerdiagnosis.nci.nih.gov/spin/>

² Health Insurance Portability and Accountability Act (HIPAA). <http://www.hhs.gov/ocr/hipaa/>

³ State law or institutional policy may differ from the HIPAA standard and should be considered as well.

or anonymization. Second, institutions cannot reveal their private databases to each other due to confidentiality of the data. In addition, the institutions may not want to reveal the ownership of their records even if the records are anonymized.

These scenarios can be generalized into the problem of privacy-preserving data publishing for multiple distributed databases where multiple data custodians or providers wish to publish an integrated view of the data for querying purposes while preserving privacy for both *data subjects* and *data providers*. We consider two privacy constraints in the problem. The first is the privacy of individuals or *data subjects* (such as the patients) which requires that the published view of the data should not contain individually identifiable information. The second is the privacy of *data providers* (such as the institutions) which requires that data providers should not reveal their private data or the ownership of the data to each other besides the published view.

Existing and potential solutions. Privacy preserving data publishing or data anonymization for a single database has been extensively studied in recent years. A large body of work contributes to algorithms that transform a dataset to meet a privacy principle such as *k*-anonymity using techniques such as generalization, suppression (removal), permutation and swapping of certain data values so that it does not contain individually identifiable information [1].

There are a number of potential approaches one may apply to enable data anonymization for distributed databases. A naive approach is for each data provider to perform data anonymization independently as shown in Fig. 1a. Data recipients or clients can then query the individual anonymized databases or an integrated view of them. One main drawback of this approach is that data is anonymized before the integration and hence will cause the data utility to suffer. In addition, individual databases reveal their ownership of the anonymized data.

An alternative approach assumes an existence of third party that can be trusted by each of the data owners as shown in Fig. 1b. In this scenario, data owners send their data to the trusted third party where data integration and anonymization are performed. Clients then can query the centralized database. However, finding such a trusted third party is not always feasible. Compromise of the server by hackers could lead to a complete privacy loss for all the participating parties and data subjects.




Fig. 1. Architectures for privacy preserving data publishing

In this paper, we propose a distributed data anonymization approach as illustrated in Fig. 1c. In this approach, data providers participate in distributed protocols to produce a *virtual* integrated and anonymized database. Important to note is that the anonymized data still resides at individual databases and the integration and anonymization of the data is performed through the secure distributed protocols. The local anonymized datasets can be unioned using secure union protocols [2,3] and then published or serve as a virtual database that can be queried. In the latter case, each individual database can execute the query on its local anonymized dataset, and then engage in distributed secure union protocols to assemble the results that are guaranteed to be anonymous.

Contributions. We study the problem of data anonymization for horizontally partitioned databases in this paper and present the distributed anonymization approach for the problem. Our approach consists of two main contributions.

First, we propose a *distributed anonymization protocol* that allows multiple data providers with horizontally partitioned databases to build a virtual anonymized database based on the integration (or union) of the data. As the output of the protocol, each database produces a local anonymized dataset and their union forms a virtual database that is guaranteed to be anonymous based on an anonymization principle. The protocol utilizes *secure multi-party computation* protocols for sub-operations such that information disclosure between individual databases is minimal during the virtual database construction.

Second, we propose a new notion, *l-site-diversity*, to ensure anonymity of data providers in addition to that of data subjects for anonymized data. We present heuristics and adapt existing anonymization algorithms for *l-site-diversity* so that anonymized data achieve better utility.

Organization. The remainder of this paper is organized as follows. Section 2 briefly reviews work related to our research. Section 3 discusses the privacy model we are using and presents our new notion on *l-site-diversity*. Section 4 presents our distributed anonymization protocol. Section 5 presents a set of experimental evaluations and Section 6 concludes the paper.

2 Related Work

Our work is inspired and informed by a number of areas. We briefly review the closely related areas below and discuss how our work leverages and advances the current state-of-the-art techniques.

Privacy preserving data publishing. Privacy preserving data publishing for centralized databases has been studied extensively [1]. One thread of work aims at devising privacy principles, such as *k*-anonymity, *l*-diversity, *t*-closeness, and *m*-invariance, that serve as criteria for judging whether a published dataset provides sufficient privacy protection. Another large body of work contributes to algorithms that transform a dataset to meet one of the above privacy principles (dominantly *k*-anonymity). In this study, our distributed anonymization protocol

is built on top of the k -anonymity and l -diversity principles and the greedy top-down Mondrian multidimensional k -anonymization algorithm [4].

There are some works focused on data anonymization of distributed databases. [5] presented a two-party framework along with an application that generates k -anonymous data from two vertically partitioned sources without disclosing data from one site to the other. [6] proposed provably private solutions for k -anonymization in the distributed scenario by maintaining end-to-end privacy from the original customer data to the final k -anonymous results.

In contrast to the above work, our work is aimed at horizontal data distribution and arbitrary number of sites. More importantly, our anonymization protocol aims to achieve anonymity for both data subjects and data providers.

Secure multi-party computation. Our approach also has its roots in the secure multi-party computation (SMC) problem [7,8,9,10,11]. This problem deals with a setting where a set of parties with private inputs wish to jointly compute some function of their inputs. An SMC protocol is *secure* if no participant learns anything more than the output.

Our problem can be viewed as designing SMC protocols for anonymization that builds virtual anonymized database and query processing that assembles query results. Our distributed anonymization approach utilizes existing secure SMC protocols for subroutines such as computing sum [12], the k th element [13], and set union [2,3]. The protocol is carefully designed so that the intermediate information disclosure is minimal.

3 Privacy Model

In this section we present the privacy goals that we focus on in this paper, followed by models and metrics for characterizing how these goals are achieved, and propose a new notion for protecting anonymity for data providers. As we identified in Section 1, we have two privacy goals. First, the privacy of individuals or data subjects needs be protected, i.e. the published virtual database and query results should not contain individually identifiable information. Second, the privacy of data providers needs to be protected, i.e. individual databases should not reveal their data or their ownership of the data apart from the virtual anonymized database.

Privacy for data subjects based on anonymity. Among the many privacy principles that protect against individual identifiability, the seminal works on k -anonymity [14,15] require that a set of k records (entities) to be indistinguishable from each other based on a quasi-identifier set. Given a relational table T , attributes are characterized into: *unique identifiers* which identify individuals; *quasi-identifier (QID)* which is a minimal set of attributes (X_1, \dots, X_d) that can be joined with external information to re-identify individual records; and *sensitive attributes* that should be protected. The set of all tuples containing identical values for the QID set is referred to as an *equivalence class*. An improved principle,

l -diversity [16], demands every group to contain at least l well-represented sensitive values.

Given our research goals of extending the anonymization techniques and integrating them with secure computation techniques to preserve privacy for both data subjects and data providers, we based our work on k -anonymity and l -diversity to achieve anonymity for data subjects. While we realize they are relatively weak compared to principles such as differential privacy, the reason we chose them for this paper is that they are intuitive and have been justified to be useful in many practical applications such as privacy-preserving location services. Therefore, techniques enforcing them in the distributed environment will still be practically important. In addition, their fundamental concepts serve as a basis for many other principles and there is a rich set of algorithms for achieving k -anonymity and l -diversity. We can study the subtle differences and effects of different algorithms and their interactions with secure multi-party computation protocols. Finally, our protocol structure, and the underlying concepts will be orthogonal to these privacy principles, and our framework will be extensible so as to easily incorporate more advanced privacy principles.

Privacy for data providers based on secure multi-party computation

Our second privacy goal is to protect privacy for data providers. It resembles the goal of secure multi-party computation (SMC). In SMC, a protocol is *secure* if no participant can learn anything more than the result of the function (or what can be derived from the result). It is important to note that for practical purposes, we may relax the security goal for a tradeoff for efficiency. Instead of attempting to guarantee absolute security in which individual databases reveal nothing about their data apart from the virtual anonymized database, we wish to minimize data exposure and achieve a sufficient level of security.

We also adopt the *semi-honest* adversary model commonly used in SMC problems. A semi-honest party follows the rules of the protocol, but it can attempt to learn additional information about other nodes by analyzing the data received during the execution of the protocol. The *semi-honest* model is realistic for our problem scenario where multiple organizations are collaborating with each other to share data and will follow the agreed protocol to get the correct result for their mutual benefit.

Privacy for data providers based on anonymity: a new notion. Now we will show that a method of simply coupling the above anonymization principles and the secure multi-party computation principles is insufficient in our scenario. While the secure multi-party computation can be used for the anonymization to preserve privacy for data providers during the anonymization, the anonymized data itself (considered as results of the secure computation) may compromise the privacy of data providers. The data partitioning at distributed data sources and certain background knowledge can introduce possible attacks that may reveal the ownership of some data by certain data providers. We illustrate such an attack, a homogeneity attack, through a simple example.

Table 1. Illustration of Homogeneity Attack for Data Providers

ID	City	Age	Disease
1	New York	30-40	Heart attack
2	New York	30-40	AIDS

Node 0

ID	City	Age	Disease
3	Northeast	40-43	AIDS
4	Northeast	40-43	Flu

Node 1

Table 1 shows anonymized data that satisfies 2-anonymity and 2-diversity at two distributed data providers (QID: City, Age; sensitive attribute: Disease). Even if SMC protocols are used to answer queries, given some background knowledge on data partitioning, the ownership of records may be revealed. For instance, if it is known that records from *New York* are provided only by *node 0*, then records with ID 1 and 2 can be linked to that node directly. In consequence, privacy of data providers is compromised. Essentially, the compromise is due to the anonymized data and cannot be solved by secure multi-party computation. One way to fix the problem is to generalize the location for records 1 and 2 so that they cannot be directly linked to a particular data provider.

To address such a problem, we propose a new notion, *l-site-diversity*, to enhance privacy protection for data providers. We define a *quasi-identifier set* with respect to data providers as a minimal set of attributes that can be used with external information to identify the ownership of certain records. For example, the location is a QID with respect to data providers in the above scenario as it can be used to identify the ownership of the records based on the knowledge that certain providers are responsible for patients from certain locations. The parameter, l , specifies minimal number of distinct sites that records in each equivalence class belong to. This notion protects the anonymity of data providers in that each record can be linked to at least l providers. Formally, the table T^* satisfies *l-site-diversity* if for every equivalence class g in T^* the following condition holds:

$$\text{count}(\text{distinct nodes}(g)) \geq l \quad (1)$$

where $\text{nodes}(g)$ returns node IDs for every record in group g .

It can be noted that our definition of *l-site-diversity* is closely related to *l-diversity*. The two notions, however, have some subtle differences. *l-diversity* protects a data subject from being linked to a particular sensitive attribute. We can map *l-site-diversity* to *l-diversity* if we treat the data provider that owns a record as a sensitive attribute for a record. However, in addition to protecting the ownership for a data record as in *l-diversity*, *l-site-diversity* also protects the *anonymity* of the ownership for data providers. In other words, it protects a data provider from being linked to a particular data subject. The QID set with respect to data providers for *l-site-diversity* could be completely different from the QID set with respect to data subjects for *k-anonymity* and *l-diversity*. *l-site-diversity* is only relevant when there are multiple data sources and it adds another check when data is being anonymized so that the resulting data will not reveal the ownership of the records. It is worth mentioning that we could also exploit much stronger definitions of *l-diversity* such as *entropy l-diversity* or *recursive (c,l)-diversity* as defined in [16].

4 Distributed Anonymization Protocol

In this section we describe our distributed anonymization approach. We first describe the general protocol structure and then present the distributed anonymization protocol.

We assume that the data are partitioned horizontally among n sites ($n > 2$) and each site owns a private database d_i . The union of all the local databases, denoted d , gives a complete view of all data ($d = \bigcup d_i$). In addition, the *quasi-identifier* of each local database is uniform among all the sites. The sites engage in a distributed anonymization protocol where each site produces a local anonymized dataset a_i and their union forms a virtual database that is guaranteed to be k -anonymous. Note that a_i is not required to be k -anonymous by itself. When users query the virtual database, each individual database executes the query on a_i and then engage in a distributed querying protocol to assemble the results that are guaranteed to be k -anonymous.

4.1 Selection of Anonymization Algorithm

Given our privacy models, we need to carefully adapt or design new anonymization algorithms with additional check for site-diversity and implement the algorithm using multi-party distributed protocols. Given a centralized version of anonymization algorithm, we can decompose it and utilize SMC protocols for sub-routines which are provably secure in order to build a secure distributed anonymization protocol. However, performing one secure computation, and using those results to perform another, may reveal intermediate information that is not part of the final results even if each step is secure. Therefore, an important consideration for designing such protocols is to minimize the disclosure of intermediate information.

There are a large number of algorithms proposed to achieve k -anonymity. These k -anonymity algorithms can be also easily extended to support l -diversity check [16]. However, given our design goal above, not all anonymization algorithms are equally suitable for a secure multi-party computation. Considering the two main strategies, top-down partitioning and bottom-up generalization, we discovered that top-down partitioning approaches have significant advantages over bottom-up generalization ones in a secure multi-party computation setting because anything revealed during the protocol as intermediate results will in fact have a coarser view than the final result and can be derived from the final result not violating the security requirement.

Based on the rationale above, our distributed anonymization protocol is based on the multi-dimensional top-down Mondrian algorithm [4]. The Mondrian algorithm uses a greedy top-down approach to recursively partition the (multi-dimensional) quasi-identifier domain space. It recursively chooses the split attribute with the largest normalized range of values, and (for continuous or ordinal attributes) partitions the data around the median value of the split attribute. This process is repeated until no allowable split remains, meaning that the data points in a particular region cannot be divided without violating the anonymity constraint, or constraints imposed by value generalization hierarchies.

Algorithm 1. Distributed anonymization algorithm - leading site ($i = 0$)

- 1: **function** split(set d_0 , ranges of QID attributes)
 - 2: **Phase 1:** Determine split attribute and split point
 - 3: Select best split attribute a (see text)
 - 4: If split is possible, send *split attribute* to node 1. Otherwise, send *finish splitting* to node 1 and finish.
 - 5: Compute median of chosen a for splitting (using secure k -th element algorithm).
 - 6: **Phase 2:** Split current dataset
 - 7: Send a and m to node 1
 - 8: Split set d_0 , create two sets, s_0 containing items *smaller than* m and g_0 containing items *greater than* m . Distribute median items among s_i and g_i .
 - 9: Send *finished* to node 1
 - 10: Wait for *finished* from last node (synchronization)
 - 11: **Phase 3:** Recursively split sub datasets
 - 12: Find $size_{left} = |\cup s_i|$ and $size_{right} = |\cup g_i|$ (using *secure sum* protocol)
 - 13: If further split of left (right) subgroup is possible, send split_left=true (split_right=true) to node 1 and call the split function recursively (updating ranges of QID attributes). Otherwise send split_left=false (split_right=false) to node 1.
 - 14: **end function** split
-

4.2 Distributed Anonymization Protocol

The key idea for the distributed anonymization protocol is to use a set of secure multi-party computation protocols to realize the Mondrian method for the distributed setting so that each database produces a local anonymized dataset which may not be k -anonymous itself, but their union forms a virtual database that is guaranteed to be k -anonymous. We present the main protocol first, followed by important heuristics that is used in the protocol.

We assume a leading site is selected for the protocol. The protocols for the leading and other sites are presented in Algorithms 1 and 2. The steps performed at the leading site are similar to the centralized Mondrian method. Before the computation starts, range of values for each quasi-identifier in set $d = \cup d_i$ and the total number of data points need to be calculated. A *secure k th element* protocol can be used to securely compute the minimum ($k=1$) and maximum ($k = n$ where n is the total number of tuples in the current partition) values of each attribute across the databases [13].

In Phase 1, the leading site selects the best split attribute and determines the split point for splitting the current partition. In order to select the best split attribute, the leading site uses a heuristic rule that is described in details below. If required, all the potential split attributes (e.g., the attributes that produce subgroups satisfying *l-site-diversity*) are evaluated and the best one is chosen. In order to determine the split medians, a *secure k th element* protocol is used ($k = \lceil \frac{n}{2} \rceil$) with respect to the data across the databases. To test whether given attribute can be used for splitting, we calculate a number of distinct sites in subgroups that would result from splitting on this attribute using the secure sum algorithm. The split is considered possible if records in both subgroups are

Algorithm 2. Distributed anonymization algorithm - non-leading node ($i > 0$)

```

1: function split(set  $c$ )
2: Read split attribute  $a$  and median  $m$  from node  $(i - 1)$ ; pass them to node  $(i + 1)$ 
3: if finish splitting received then return
4: Split set  $c$  into  $s_i$  containing items smaller than  $m$  and  $g_i$  containing items
   greater than  $m$ . Distribute median items among  $s_i$  and  $g_i$ .
5: Read finished from node  $i - 1$  (synchronization); Send finished to node  $i + 1$ 
6: Read split_left from node  $i - 1$  and pass it to node  $i + 1$ 
7: if split_left then call split( $s_i$ )
8: Read split_right from node  $i - 1$ , Send split_right to node  $i + 1$ 
9: if split_right then call split( $g_i$ )
10: end function split

```

original data			anonymized data			original data			anonymized data		
ID	ZIP	Age	ID	ZIP	Age	ID	ZIP	Age	ID	ZIP	Age
1	30030	31	1	30030-36	31-32	3	30045	45	3	30037-56	32-45
2	30033	32	2	30030-36	31-32	4	30056	32	4	30037-56	32-45
node 0						node 1					
ID	ZIP	Age	ID	ZIP	Age	ID	ZIP	Age	ID	ZIP	Age
5	30030	22	5	30030-36	22-30	7	30038	31	7	30037-56	22-31
6	30053	22	6	30037-56	22-31	8	30033	30	8	30030-36	22-30
node 2						node 3					

Fig. 2. Distributed anonymization illustration

provided by at least l sites. In Phase 2, the algorithm performs split and waits for all the nodes to finish splitting. Finally in Phase 3, the node recursively checks whether further split of the new subsets is possible. In order to determine whether a partition can be further split, a *secure sum* protocol [12] is used to compute the number of tuples of the partition across the databases.

We illustrate the overall protocol with an example scenario shown in Figure 2 where we have 4 nodes and we use $k = 2$ for k -anonymization and $l = 1$ for l -site-diversity. Note that the anonymized databases at node 2 and node 3 are not 2-anonymous by themselves. However the union of all the anonymized databases is guaranteed to be 2-anonymous.

Selection of split attribute. One key issue in the above protocol is the selection of split attribute. The goal is to split the data as much as possible while satisfying the privacy constraints so as to maximize discernibility or utility of anonymized data. The basic Mondrian method uses the range of an attribute as a goodness indicator. Intuitively, the larger the spread, the easier the good split point can be found and more likely the data can be further split. In our setting, we also need to take into account the site diversity requirement and adapt the selection heuristic. The importance of doing so is demonstrated in Figure 3. Let's assume that we want to achieve 2-anonymity and 2-site-diversity. In the first scenario, the attribute for splitting is chosen only based on range of the QID attributes. The protocol finishes with 2 groups of 5 and 4 records (further split is impossible due to 2-site-diversity requirement). The second scenario




Fig. 3. Impact of split attribute selection when l -site-diversity ($n = 2$) is considered. Different shades represent different owners of records.

exploits information on records distribution when the decision on split attribute is made (the more evenly the records are distributed across sites in resulting subgroups, the better). This rule yields better results, namely three groups of 3 records each.

Based on the illustration, intuition suggests that we need to select a split attribute that results in partitions with even distribution of records from different data providers. This makes further splits more likely while meeting the l -site-diversity constraint. Similar to decision tree classifier construction [17], information gain can be used as a scoring metric for selecting attribute that results in partitions with most diverse distribution of data providers. Note that this is used in a complete opposite sense from decision tree where the goal is to partition the data into partitions with homogeneous classes. The information gain of a potential splitting attribute a_k is computed through the information entropy of resulting partitions:

$$e(a_k) = - \sum_{i=0}^{n-1} p(i, l_k) \log(p(i, l_k)) - \sum_{i=0}^{n-1} p(i, r_k) \log(p(i, r_k)) \quad (2)$$

where l_k and r_k are partitions created after splitting the input set using attribute a_k (and its median value) and $p(i, g)$ is the portion of records that belong to node i in group g . It is important to note that the calculations need to take into account data on distributed sites and thus secure sum protocol needs to be used.

Our final scoring metric combines the original range value based metric and the new diversity-aware metrics using a linear combination as follows:

$$\forall_{a_i \in Q} s_i = \alpha \frac{range(a_i)}{\max_{a_j \in Q}(range(a_j))} + (1 - \alpha) \frac{e(a_i)}{\max_{a_j \in Q}(e(a_j))} \quad (3)$$

where $range$ function returns range of attribute, $e(a_i)$ returns values of information entropy as defined above when attribute a_i is used for splitting and α is a weighting parameter.

Important to note is that if l -site-diversity is not required (e.g., $l=1$), then the evaluation of the heuristic rule above is limited to checking only the range of attributes, and choosing the attribute with the widest range.

4.3 Analysis

Having presented the distributed anonymization protocol, we analyze the protocol in terms of its security and overhead.

Security. We will now analyze the security of our distributed k -anonymity protocol. Our proofs will show that, given the result, the leaked information (if any), and the site's own input, any site can simulate the protocol and everything that was seen during the execution. Since the simulation generates everything seen during execution of the protocol, clearly no one learns anything new from the protocol when it is executed. The proofs will also use a general composition theorem [7] that covers algorithms implemented by running many invocation of secure computations of simpler functionalities. Let's assume a *hybrid model* where the protocol uses a trusted third-party to compute the result of such smaller functionalities $f_1 \dots f_n$. The composition theorem states that if a protocol in hybrid model is secure in terms of comparing the real computation to the ideal model, then if a protocol is changed in such a way that calls to trusted third-party are replaced with secure protocols, the resulting protocol is still secure.

We have analyzed the distributed k -anonymity protocol in terms of security and present the following two theorems with sketches of proofs. For complete proofs, we refer readers to [18].

Theorem 1. The distributed k -anonymity protocol privately computes a k -anonymous view of horizontally partitioned data in semi-honest model when $l = 1$.

Proof sketch. The proof needs to show that any given node can simulate the algorithm and all what was seen during its execution given only the final result and its local data. The simulation first analyzes the complete anonymized data and finds a range of each attribute from the quasi identifier. This can be done by looking for the largest and the smallest possible values of all the QID attributes. Next, as no l -site-diversity is required ($l=1$), the site knows that the attribute used for splitting is actually the attribute with the largest range. As the site knows which attribute was used for splitting, it can attempt to identify the split point using the following approach. First, it identifies all possible splitting attributes/points that could be used when the algorithm was executed. Formally, the points of potential split are the distinct values that appear on the bounds of the QID ranges in the anonymized view. To identify the median value (or the value that was used for split) the node checks which of the potential splitting points is actually a median. This can be done by choosing a value that divides the set of records into two sets with the sizes closest to half of the number of records in the database (note that the subsets resulting from spitting might not have equal sizes - for instance if number of records is odd). Now the site is ready to simulate the split. If size of any of the two groups that result from splitting is greater than or equal to $2 * k$, this group can be further split. In such case the node would continue the described simulation with input data being one of the subgroups.

Theorem 2. The distributed k -anonymity protocol privately computes a k -anonymous view of horizontally partitioned data in semi-honest model, revealing at most the following statistics of the data, when $l > 1$.

1. Median values of each attribute from QID for groups of records of size $\geq 2*k$,
2. Entropy of distribution of records for groups resulting from potential splits,
3. Number of distinct sites that provide data to groups resulting from potential splits (the identity of those sites are confidential).

Proof sketch. The proof uses a similar approach as above. The main difference is in the decision step when a node uses the entropy based heuristic designed for l -site-diversity to decide on the split attribute. In this case, not only the range of attribute, but also the distribution of records in groups resulting from potential splitting, need to be considered. Using the final result, information from points 1, 2 and 3, and the range of QID attributes, however, any node can decide on the split attribute in the protocol simulation.

Overhead. Our protocol introduces additional overhead due to the fact that the nodes have to use additional protocols in each step of computation. The time complexity of the original Mondrian algorithm is $O(n \log n)$ where n is the number of items in the anonymized dataset [4]. As we presented in Algorithm 1, each iteration of the distributed anonymization algorithm requires calculation of the heuristic decision rule, median value of an attribute, and the count of tuples of a partition. The secure sum protocol does not depend on the number of tuples in the database. The secure k -th element algorithm is logarithmic in number of input items (assuming the worst - case scenario that all the input items are distinct). As a consequence, the time complexity of our protocol can be estimated as $O(n \log^2 n)$ in terms of number of records in a database.

The communication overhead of the protocol is determined by two factors. The first is the cost for a single round. This depends on the number of nodes involved in the system and the topology which is used and in our case it is proportional to the number of nodes on the ring. As the future work, we are considering alternative topologies (such as trees) in order to optimize the communication cost for each round. The second factor is the number of rounds and is determined by the number of iterations and the sub-protocols used by each iteration of the anonymization protocol. The secure sum protocol involves one round of communication. In the secure k th element protocol, the number of rounds is $\log M$ (M being the range of attribute values) and each round requires secure computations twice. It is important to note that the distributed anonymization protocol is expected to be run offline on an infrequent basis. As a result, the overhead of the protocol will not be a major issue.

5 Experimental Evaluation

We have implemented the distributed anonymization protocol in Java within the DObjects framework [19] which provides a platform for querying data across

distributed and heterogeneous data sources. To be able to test a large variety of configurations, we also implemented the distributed anonymization protocol using a simulation environment. In this section we present a set of experimental evaluations of the proposed protocols.

The questions we attempt to answer are: 1) What is the advantage of using distributed anonymization algorithm over centralized or independent anonymization? 2) What is the impact of the *l-site-diversity* constraint on anonymization protocol? 3) What are the optimal values for α parameter in our heuristic rules presented in equation 3?

5.1 Distributed Anonymization vs. Centralized and Independent Anonymization

We first present an evaluation of the distributed anonymization protocol compared to the centralized and independent anonymization approaches in terms of the quality of the anonymized data.

Dataset and setup. We used the Adult dataset from UC Irvine Machine Learning Repository. The dataset contained 30161 records and was configured as in [4]. We used 3 distributed nodes (30161 records were split among those nodes using round-robin protocol). We report results for the following scenarios: 1) the data is located in one centralized database and classical Mondrian k -anonymity algorithm was run (centralized approach), 2) data are distributed among the three nodes and Mondrian k -anonymity algorithm was run at each site independently (independent or naive approach) and 3) data are distributed among the three nodes and we use the distributed anonymization approach presented in section 4. We ran each experiment for different k values. All the experiments in this subsection used *1-site-diversity*.

Results. Figure 4 shows the average equivalence class size with respect to different values of k . We observe that our distributed anonymization protocol performs the same as the non-distributed version. Also as expected, the naive approach (independent anonymization of each local database) suffers in data utility because the anonymization is performed before the integration of the data.






Fig. 4. Average equivalence class size vs. k **Fig. 5.** Histogram for partitioning using City and Age **Fig. 6.** Average error vs. α (information gain-based)

5.2 Achieving Anonymity for Data Providers

The experiments in this section again use the Adult dataset. The data is distributed across $n = 100$ sites unless otherwise specified. We experimented with distribution pattern that we will describe in detail below.

Metric. The average equivalence group size as shown in previous subsection provides a general data utility metric. The query imprecision metric provides an application-specific metric that is of particular relevance to our problem setting. Given a query, since the attribute values are generalized, it is possible only to return the tuples from the anonymized dataset that are contained in any generalized ranges overlapping with the selection predicate. This will often produce a larger result set than evaluating the predicate over the original table. For this set of experiments, we use summary queries (queries that return count of records) and we use an algorithm similar to the approach introduced in [20] that returns more accurate results. We report a *relative error* of the query results. Specifically, given act as an exact answer to query and est as an answer computed according to algorithm defined above, the relative error is defined as $|act - est|/act$. For each of the tested configurations, we submit 10,000 randomly generated queries, and for each query we calculate a relative error. We report average value of the error. Each query uses predicates on two randomly chosen attributes from *quasi-identifier*. For boolean attributes that can have only two values (e.g. sex), the predicate has a form of $a_i = value$. For other attributes we use predicate in the form $a_i \in R$. R is a random range and has a length of $0.3 * |a_i|$, where $|a_i|$ denotes the domain size of an attribute.

Data partitioning. In a realistic scenario, data is often split according to some attributes. For instance, the patient data can be split according to cities, i.e. the majority records from a hospital located in New York would have a New York address while those from a hospital located in Boston would have a Boston address. Therefore, we distributed records across sites using partitioning based on attribute values. The rules of partitioning were specified using two attributes, City and Age. The dataset contained data from 6 different cities, and every 1/6th of available nodes were assigned to a different city. Next, records within each group of nodes for a given city were distributed using Age attribute: records with age less than 25 were assigned to the first 1/3rd of nodes, records with age between 25 and 55 to the second 1/3rd of nodes, and the remaining records to the remaining nodes. The histogram of the records per node in this setup is presented in Figure 5 (please note the logarithmic scale of the plot).

Results. We now present the results evaluating the impact of α value under this setup. Figure 6 presents the average query error for different α and l values for the heuristic rule we used. We can observe a significant impact of α value on the average error. The smallest error value is observed for $\alpha = 0.3$ and this seems to be an optimal choice for all tested l values. One can observe 30% decrease in error when compared to using only range as in original Mondrian ($\alpha = 1.0$) or using only diversity-aware metrics ($\alpha = 0.0$). It is worth mentioning that we






Fig. 7. Average error vs. k ($l = 30$) **Fig. 8.** Average error vs. l ($k = 200$) **Fig. 9.** Error vs. n ($k = 200$ and $l = 30$)

have also experimented with different distributions of records, and the results were consistent with what we presented above. We do not provide these results due to space limitations.

The next experiment was focused on the impact of k parameter on average error. We present results for $l=30$ in Figure 7 for three different split heuristic rules: using range only, information gain only, and combining range with information gain with $\alpha = 0.3$. We observe that the heuristic rule that takes into account both range and information gain gives consistently the best results and a reduction of error around 30%. These results do not depend on the value of k .

Next, we tested the impact of the l parameter for l -site-diversity. Figure 8 shows an average error for varying l and $k = 200$ using the same heuristic rules as in previous experiment. Similarly, the rule that takes into account range and information gain gives the best results. With increasing l , we observe an increasing error rate because the data needs to be more generalized in order to satisfy the diversity constraints.

So far we have tested only scenarios with 100 nodes ($n = 100$). To complete the picture, we plot the average error for varying n ($k = 200$ and $l = 30$) in Figure 9. One can notice that the previous trends are maintained - the results do not appear to be dependent on the number of nodes in the system. Similarly, the rule that takes into account range and information gain is superior to other methods and the query error is on average 30% smaller than that for others.

6 Conclusion

We have presented a distributed and decentralized anonymization approach for privacy-preserving data publishing for horizontally partitioned databases. Our work addresses two important issues, namely, privacy of data subjects and privacy of data providers. We presented a new notion, l -site-diversity, to achieve anonymity for data providers in anonymized dataset. Our work continues along several directions. First, we are interested in developing a protocol toolkit incorporating more privacy principles and anonymization algorithms. In particular, dynamic or serial releases of data with data updates are extremely relevant in our distributed data integration setting. Such concepts as m -invariance [20] or

l-scarsity [21] are promising ideas and we plan to extend our research in this direction. Second, we are also interested in developing specialized multi-party protocols such as set union that offer a tradeoff between efficiency and privacy as compared to the existing set union protocols based on cryptographic approaches.

Acknowledgement

We thank Kristen LeFevre for providing us the implementation for the Mondrian algorithm and the anonymous reviewers for their valuable feedback. The research is partially supported by a URC and an ITSC grant from Emory and a Career Enhancement Fellowship from the Woodrow Wilson Foundation.

References

1. Fung, B.C.M., Wang, K., Chen, R., Yu, P.S.: Privacy-preserving data publishing: A survey on recent developments. *ACM Computing Surveys* (in press)
2. Kantarcioglu, M., Clifton, C.: Privacy preserving data mining of association rules on horizontally partitioned data. *IEEE Transactions on Knowledge and Data Engineering (TKDE)* 16(9) (2004)
3. Böttcher, S., Obermeier, S.: Secure set union and bag union computation for guaranteeing anonymity of distrustful participants. *JSW* 3(1), 9–17 (2008)
4. LeFevre, K., DeWitt, D., Ramakrishnan, R.: Mondrian multidimensional k-anonymity. In: *Proceedings of the International Conference on Data Engineering (ICDE 2006)* (2006)
5. Jiang, W., Clifton, C.: A secure distributed framework for achieving k-anonymity. *VLDB Journal* 15(4), 316–333 (2006)
6. Zhong, S., Yang, Z., Wright, R.N.: Privacy-enhancing k-anonymization of customer data. In: *Proc. of the Principles of Database Systems (PODS)* (2005)
7. Goldreich, O.: Secure multi-party computation, Working Draft, Version 1.3 (2001)
8. Clifton, C., Kantarcioglu, M., Vaidya, J.: Tools for privacy preserving distributed data mining. *ACM SIGKDD Explorations* 4 (2003)
9. Lindell, Y., Pinkas, B.: Secure multiparty computation for privacy-preserving data mining. *Cryptology ePrint Archive, Report 2008/197* (2008), <http://eprint.iacr.org/>
10. Vaidya, J., Clifton, C.: Privacy-preserving data mining: Why, how, and when. *IEEE Security & Privacy* 2(6), 19–27 (2004)
11. Du, W., Atallah, M.J.: Secure multi-party computation problems and their applications: a review and open problems. In: *NSPW 2001: Proceedings of the 2001 workshop on New security paradigms*, pp. 13–22. ACM, New York (2001)
12. Schneier, B.: *Applied Cryptography*, 2nd edn. John Wiley & Sons, Chichester (1996)
13. Aggarwal, G., Mishra, N., Pinkas, B.: Secure computation of the kth-ranked element. In: Cachin, C., Camenisch, J.L. (eds.) *EUROCRYPT 2004*. LNCS, vol. 3027, pp. 40–55. Springer, Heidelberg (2004)
14. Samarati, P.: Protecting respondents' identities in microdata release. *IEEE Trans. Knowl. Data Eng.* 13(6), 1010–1027 (2001)
15. Sweeney, L.: k-anonymity: a model for protecting privacy. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.* 10(5), 557–570 (2002)

16. Machanavajjhala, A., Gehrke, J., Kifer, D., Venkatasubramanian, M.: l-diversity: Privacy beyond k-anonymity. In: Proceedings of the International Conference on Data Engineering (ICDE 2006), p. 24 (2006)
17. Han, J., Kamber, M.: Data Mining: Concepts and Techniques, 2nd edn. Morgan Kaufmann, San Francisco (2006)
18. Jurczyk, P., Xiong, L.: Distributed anonymization: Achieving privacy for both data subjects and data providers. Technical Report TR-2009-013, Emory University Department of Mathematics and Computer Science (2009)
19. Jurczyk, P., Xiong, L.: Dobjects: Enabling distributed data services for metacomputing platforms. In: Bubak, M., van Albada, G.D., Dongarra, J., Sloot, P.M.A. (eds.) ICCS 2008, Part I. LNCS, vol. 5101, pp. 136–145. Springer, Heidelberg (2008)
20. Xiao, X., Tao, Y.: M-invariance: towards privacy preserving re-publication of dynamic datasets. In: Proc. of the ACM SIGMOD International Conference on Management of Data, pp. 689–700 (2007)
21. Bu, Y., Fu, A.W.C., Wong, R.C.W., Chen, L., Li, J.: Privacy preserving serial data publishing by role composition. Proc. VLDB Endow. 1(1), 845–856 (2008)