

Protecting your sensitive values from sequential background knowledge attacks

Author¹ : S.JEEVA
 M.E., Computer Science and Engineering
 Arunai College of Engineering
 Tiruvannamalai-606603
 Email id: jeevavelu13@gmail.com

Author² :Mr.J.SRINIVASAN,
 M.E., Assistant Professor, Computer Science and
 Engineering
 Arunai College of Engineering
 Tiruvannamalai-606603
 Email id: srinivasanj.cs@gmail.com

Abstract—Web queries, credit card transactions, and medical records are examples of transaction data flowing in corporate data stores, and often revealing associations between individuals and sensitive information. The serial release of these data to partner institutions or data analysis centers in a non-aggregated form is a common situation. In this paper, we show that correlations among sensitive values associated to the same individuals in different releases can be easily used to violate users' privacy by adversaries observing multiple data releases, even if state-of-the-art privacy protection techniques are applied. We show how the above sequential background knowledge can be actually obtained by an adversary, and used to identify with high confidence the sensitive values of an individual. Our proposed defense algorithm is based on Jensen-Jensen-Shannon divergence; experiments show its superiority with respect to other applicable solutions. To the best of our knowledge, this is the first work that systematically investigates the role of sequential background knowledge in serial release of transaction data.

institutions; new records are periodically released to data analysis centers in non-aggregated form.

A challenging issue in this scenario is the protection of users' privacy, considering that potential adversaries have access to multiple serial releases and can easily acquire background knowledge related to the specific domain. This knowledge includes the fact that certain sequences of values in subsequent releases are more likely to be observed than other sequences that a sequence of medical exam results within a certain time frame has higher probability to be observed than another sequence.

Privacy protection approaches can be divided in micro data anonymity and differential privacy methods. Micro data anonymity works have focused on techniques dealing either with multiple data releases, or with adversary background knowledge, but limited to a single data release.

In this paper, we formally model privacy attacks based on background knowledge extended to serial micro-data releases. We present a new probabilistic defense technique taking into account adversary's background knowledge and how he can revise it each time new data are released. Similarly to other anonymization techniques, our method is based on the generalization of quasi-identifier (QI) values, but generalization is performed with a new goal: minimizing the difference among sensitive values probability distributions within each QI-group, while considering the knowledge revision process. Jensen-Shannon divergence is used as a measure of similarity. We consider different methods and accuracy levels for the extraction of background knowledge, and we show that

I.INTRODUCTION

Large amounts of data related to individuals are continuously acquired, and stored by corporate and government institutions. Examples include mobile service requests, web queries, credit card transactions, and transit database records. These institutions often need to repeatedly release new or updated portions of their data to other partner institutions for different purposes, including distributed processing, participation in inter-organizational workflows, and data analysis. The medical domain is an interesting example: many countries have recently established centralized data stores that exchange patients' data with medical

our defense is effective under different combinations of the knowledge of the adversary and the defender.

II.MOTIVATING SCENARIO

We consider the case of transaction data representing the results of medical exams taken by patients, and the need to periodically release these transactions for data analysis. Each released view contains one tuple for each patient who performed an exam during the week preceding the publication. We assume that data are published weekly. For the sake of simplicity, we also assume that each user cannot perform more than one exam per week; hence, no more than one tuple per user can appear in the same view. Each generalized tuple includes the age, gender, and zip code of the patient, as well as the performed exam together with its result. We refer to this latter data, represented by the multi-value attribute Exres, as exam result.1 We denote as positive (pos) a result that reveals something anomalous; negative (neg) otherwise. The attribute Exres is considered the sensitive attribute, while the other attributes play the role of quasi identifiers (QI), since they may be used, joined with external information, to restrict the set of candidate respondents.

TABLE 1
Original and Generalized Transaction Data at the First and Second Release (First and Second Week, Respectively)

(a) Original transaction data at time τ_1					(b) Generalized transaction data: 1st release				
Name	Age	Gender	Zip	Exres	QI-group	Age	Gender	Zip	Ex-res
Alice	51	F	12030	MAM-pos	1	[51,52]	F	12030	MAM-pos
Betty	52	F	12030	CX-neg	1	[51,52]	F	12030	CX-neg
Carol	51	F	12031	CX-pos	2	[51,52]	F	12031	CX-pos
Doris	52	F	12031	BS-neg	2	[51,52]	F	12031	BS-neg

(c) Original transaction data at time τ_2					(d) Generalized transaction data: 2nd release				
Name	Age	Gender	Zip	Ex-res	QI-group	Age	Gender	Zip	Ex-res
Alice	51	F	12030	BCM-pos	3	51	F	1203*	BCM-pos
Carol	51	F	12031	PNE-pos	3	51	F	1203*	PNE-pos
Elisa	51	F	12044	MAM-neg	4	51	F	1204*	MAM-neg
Fran	51	F	12045	CX-neg	4	51	F	1204*	CX-neg
Grace	51	F	12040	CX-pos	4	51	F	1204*	CX-pos

adversary cannot exploit BKsv (reported in Table 2) to infer whether Alice or Betty is the respondent of the tuple with value MAM-pos. Hence, his posterior knowledge after observing tuples released at τ_1 states that, both for Alice and Betty, the probability of being the respondent of one tuple with private value MAM-pos is the same of being the respondent of one tuple with private value CX-neg, i.e., 0.5. Analogously, Carol and Doris have equal probability of being the respondent of one tuple with private value CX-pos and of one with private value BS-neg.

TABLE 2
Adversary's Background Knowledge

(a) Sensitive values background knowledge at τ_1						(b) Sequential background knowledge		
Name	Age	Gender	Zip	Exres	BK ^{sv}	Exres at τ_1	Exres at τ_2	$p(s_{\tau_2} s_{\tau_1})$
Alice	51	F	12030	MAM-pos	0.002	MAM-pos	BCM-pos	0.6
Betty	52	F	12030	MAM-pos	0.002	CX-neg	BCM-pos	0.02
Alice	51	F	12030	CX-neg	0.05	CX-pos	BCM-pos	0.02
Betty	52	F	12030	CX-neg	0.05	BS-neg	BCM-pos	0.02
Carol	51	F	12031	CX-pos	0.0003	MAM-pos	PNE-pos	0.02
Doris	52	F	12031	CX-pos	0.0003	CX-neg	PNE-pos	0.08
Carol	51	F	12031	BS-neg	0.2	CX-pos	PNE-pos	0.6
Doris	52	F	12031	BS-neg	0.2	BS-neg	PNE-pos	0.02
Alice	51	F	12030	BCM-pos	0.001			

III. MODELING ATTACKS BASED ON BACKGROUND AND REVISED KNOWLEDGE

In this section, we formally model privacy attacks based on background and revised knowledge.

A. Problem Definition

We denote by V_i a view on the original transaction data at time T_i , and by V_i^* the generalization of V_i released by the data publisher. We denote a history of released generalized views by $H_j^* = \langle V_1^*, V_2^*, \dots, V_j^* \rangle$. We assume that the schema remains unchanged throughout the release history, and we partition the view columns into a set $A_{qi} = \{A_1, A_2, \dots, A_m\}$ of quasi-identifier attributes, and into a single private attribute S . For simplicity, we assume that the domain of each quasi-identifier attribute is numeric, but our notions and techniques can be easily extended to categorical attributes. Given a tuple t in a view and an attribute A in its schema, $t[A]$ is the projection of tuple t onto attribute A .

B.SENSITIVE VALUES BACKGROUND KNOWLEDGE (BKSV)

Sensitive values background knowledge represents the apriori probability of associating an individual to a sensitive value. We model the sensitive value referring to a respondent r by means of the discrete random variable S having values in $D[S]$. BKsv is modeled according to the following definition.

Definition1. The sensitive values background knowledge is a function $BKsv : R \rightarrow \Upsilon$, where R is the set of possible respondents' identities, and

$$\Upsilon = \left\{ (p_1, \dots, p_n) \mid \sum_{1 \leq i \leq n} p_i = 1 \ (0 \leq p_i \leq 1) \right\},$$

is the set of possible probability distributions of S , where $D[S] = \{s_1, s_2, \dots, s_n\}$.

For example, if $r \in R$ is a possible respondent of a tuple in a released view, $BK^{sv}(r)$ returns, for each sensitive value $s_j \in D[S]$, the probability p_j of r being actually associated with s_j .

IV. JS-REDUCE DEFENSE

In this section, we illustrate the JS-reduce defense against the identified background knowledge attacks.

A. Defense Strategy

In order to enforce anonymity, it is necessary to limit the adversary's capability of identifying the actual respondent of a tuple in a given QI-group. This means reducing the confidence of the adversary in discriminating a configuration C among the possible ones, based on his revised knowledge RBKsv.

The goal of JS-reduce is to create QI-groups whose tuple respondents have similar RBKsv (resp. BKsv) distributions. Indeed, if the respondents of tuples in a QI-group are indistinguishable with respect to RBKsv (resp. BKsv), the adversary cannot exploit background knowledge to perform the attack. Of course, defending against background knowledge attacks is not sufficient to guarantee privacy protection against other kinds of attacks. For this reason, JS-reduce also enforces k-anonymity and t-closeness, in order to protect against well-known identity and attribute-disclosure attacks,




Fig.1. Defense mechanisms.

B. The JS-Reduce Algorithm

The pseudocode of the JS-reduce algorithm is shown in Algorithm 1. The algorithm takes as input:

Algorithm 1. JS-reduce algorithm

```

Input: Sequence  $\mathcal{H}_n = \langle V_1, \dots, V_n \rangle$ , the set  $R$  of possible respondents as well as their QI values,  $BK^{sv}, BK^{seq}$ , the minimum level  $k$  of  $k$ -anonymity, threshold  $t$  of  $t$ -closeness, threshold  $j$  of JS divergence.

Output:  $V_n^*$ 

1 JS-reduce( $\mathcal{H}_n, R, BK^{sv}, BK^{seq}, k, t, j$ )
2 begin
3   forall  $r \in R$  do
4      $RBK_1^{sv}(r) \leftarrow BK^{sv}(r)$ 
5   end
6   for  $h = 1$  to  $n$  do
7      $V_h^* \leftarrow \text{Generalize}(V_h, RBK_h^{sv}, t, j, k)$ 
8     forall  $r \in R_h$  do
9        $PK_h^{sv}(r) \leftarrow \text{PKComputation}(V_h^*, RBK_h^{sv}, r)$ 
10       $RBK_{h+1}^{sv}(r) \leftarrow \text{BKRevision}(PK_h^{sv}(r), BK^{seq}, r)$ 
11    end
12  end
13  return  $V_n^*$ 
14 end

13  $p(r, s) \leftarrow \frac{\sum_{e_j \in C | e_j(t)=r \wedge e_j(S)=s} d_j}{\sum_{e_j \in C} d_j}$ 

14 end
15  $PK_h^{sv}(r) \leftarrow \{p(r, \tilde{s}), \forall \tilde{s} \in D[S]\}$ 
16 return  $PK_h^{sv}(r)$ 
17 end
    
```

C. Data Quality-Oriented Generalization

Any anonymization technique based on QI generalization needs to carefully consider the resulting data quality: the more the QI values are generalized, the lower is the quality of released data. Hence, instead of adopting a general purpose anonymization framework such as Mondrian we devised an ad hoc technique. Note that finding the optimal generalization of data that satisfies the privacy requirements of JS-reduce (i.e., the one that minimizes QI generalization) is an NP-hard problem; indeed, it is well known that even optimal k-anonymous generalization is NP-hard.

V. EXPERIMENTAL EVALUATION

In this section, we present an experimental evaluation of the privacy threats due to sequential background knowledge attacks, we compare our defense with other applicable solutions, and we evaluate data quality.

TABLE 5
Values of Privacy Parameters Used in the Experiments

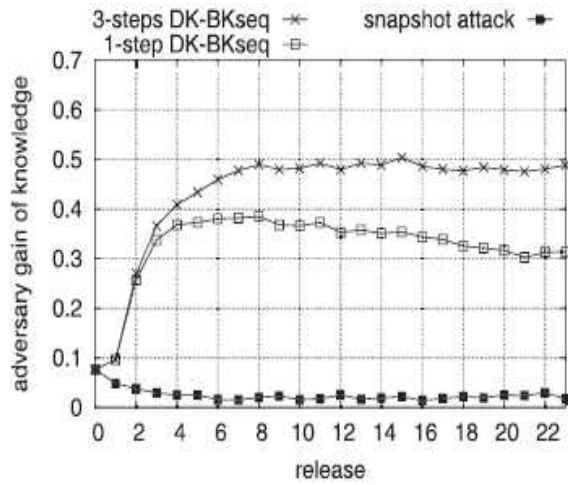
	l	t	B	j
l-div.	[2, 8] 2	-	-	-
t-clos.	-	[0.5, 1] 0.8	-	-
(B,t)-priv.	-	[0.5, 0.8] 0.8	[0.3, 0.7] 0.5	-
JS-red.	-	[0.5, 0.8] 0.5	-	[0.2, 0.8] 0.6

performed during that week. A tuple is composed of three QI attributes age, gender, and weight, and a sensitive attribute Ex-res. Age has values in the interval [45, 74] gender in [1, 2] and weight in [60, 89]. The domain of Ex-res includes 17 different values associated to stages of different diseases (five stages of liver disease, four of the HIV syndrome, three of Alzheimer, and five of sepsis), as well as two sensitive values to describe the deceased and discharged events.

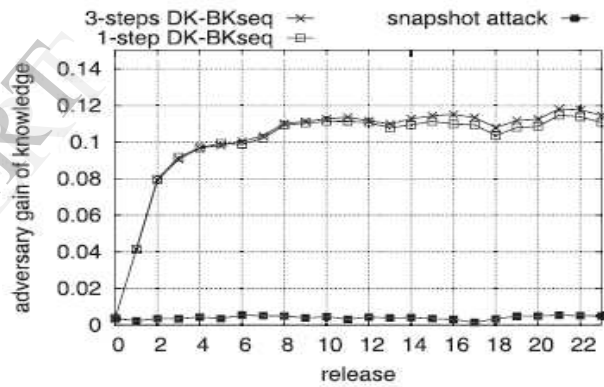
A. The Role of Adversary's Background Knowledge

We performed experiments to evaluate the role of background knowledge on the privacy threats investigated in this paper: Incrementally extracted knowledge IE-BKseq. Since it was the subject of related studies the first kind of background knowledge we consider is the one directly extracted from the data to be released. IE-BKseq can be calculated by applying sequential pattern mining (SPM) techniques on the history of original (i.e., non-anonymized) data at each time t_i , IE-BKseq is calculated based on V_i . Since the size of the corpus is relatively small, we applied a simple SPM algorithm, which is essentially based on a frequency count of sequences appearing in the history.

i.e., it is the "best" knowledge that an adversary may have.

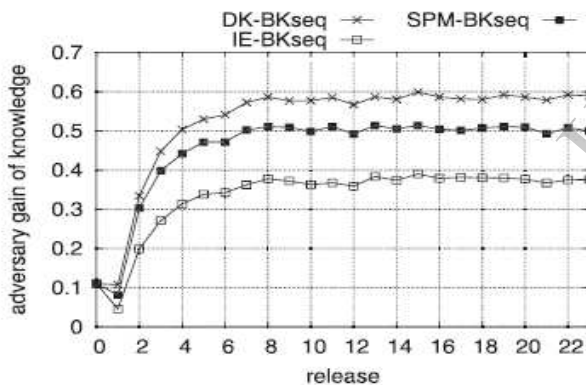


(b) (B,t)-priv.



(c) JS-red.

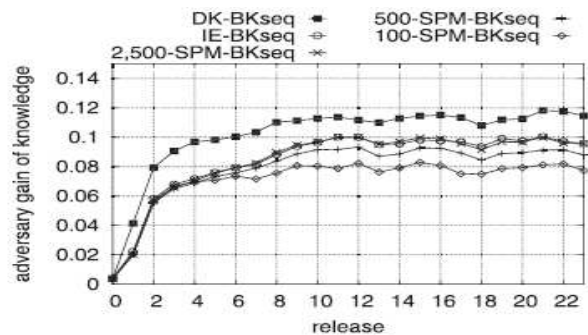
Fig.2. Adversary gain versus accuracy of adversary's domain knowledge DK-BKseq.



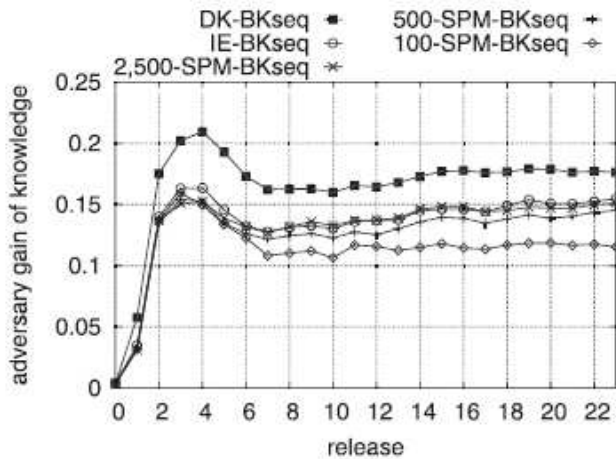
(a) l-div.

Mined knowledge SPM-BKseq. In practice, an adversary may approximate BKseq by applying SPM techniques on an external corpus of nonanonymized data. We created a data corpus using the same model that we used to generate our data set; the corpus consists in a history of 24 views containing 5,000 tuples each. SPM-BKseq was calculated by applying Algorithm 5 to that corpus.

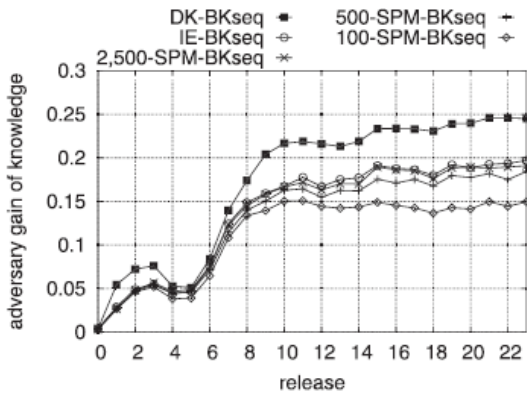
Domain knowledge DK-BKseq. Since our data set was generated based on domain knowledge, in our experiments DK-BKseq corresponds to the exact BKseq;



(a) Defense based on DK-BKseq



(b) Defense based on $IE-BK^{seq}$



(c) Defense based on $SPM-BK^{seq}$

Fig.3. JS-reduce versus different kinds of adversary's BKseq

B. Effectiveness of the JS-Reduce Defense

Results reported in Fig. 4c show that, when views are anonymized with JS-reduce, the adversary gain remains below 0.12, independently from the length of the released history, and on the kind of domain knowledge available to the adversary. This result shows that JS-reduce significantly limits the inference capabilities of the adversary with respect to the other techniques that lead to an adversary gain higher than 0.5. We performed other experiments to evaluate the effectiveness of JS reduce with different combinations of background knowledge available to the defender and to the adversary, respectively. In Fig. 5a, we considered the

case in which the defender has background knowledge DK-BKseq. In this case, the defense is very effective, even when the adversary has the same background knowledge as the defender. When the adversary's background knowledge is extracted from the data, we observe that the adversary gain is lower. With the label n -SPM-BKseq in Fig. 5, we denote that the adversary's SPM-BKseq is extracted based on a history of 24 views containing n tuples each.

The adversary gain is lower with smaller values of n , since the resulting SPM-BKseq is a coarser approximation of the exact BKseq. The adversary gain with incrementally extracted knowledge is comparable to the one obtained with SPM-BKseq. We also considered the unfortunate case in which the adversary has more accurate background knowledge than the defender. Results illustrated in Figs. 5b and 5c show the adversary gain when the defender's background knowledge is IE-BKseq and SPM-BKseq, respectively. As expected, the more accurate the attacker's background knowledge with respect to the defender's one, the more effective the attack.




Fig.4. Adversary confidence.




Fig. 7. Adversary gain versus JS divergence ($t = 0.5$).




Fig.5. Data quality evaluation

VI.CONCLUSIONS

In this paper, we showed that the correlation of sensitive values in subsequent releases can be used as background knowledge to violate users' privacy. We showed that an adversary can actually obtain this knowledge by different methods. We proposed a defense algorithm and we showed through an extensive experimental evaluation that other applicable solutions are not effective, while our defense provides strong privacy protection and good data quality, even when the adversary has more accurate background knowledge than the defender. Our framework is seamlessly extensible with additional forms of probabilistic inference, since the JS-reduce technique relies on a background knowledge revision process that is not tied to a specific inference method.

REFERENCES

[1] R. Agrawal and R. Srikant, "Mining Sequential Patterns," Proc. 11th Int'l Conf. Data Eng. (ICDE '95), pp. 3-14, 1995.

[2] Y. Bu, A. Wai, C. Fu, R.C.W. Wong, L. Chen, and J. Li, "Privacy Preserving Serial Data Publishing by Role Composition," Proc. VLDB Endowment, vol. 1, pp. 845-856, 2008.

[3] J.-W. Byun, Y. Sohn, E. Bertino, and N. Li, "Secure Anonymization for Incremental Data Sets," Proc. Third

VLDB Workshop Secure Data Management (SDM '06), pp. 48-63, 2006.

[4] J. Cao, B. Carminati, E. Ferrari, and K.-L. Tan, "CASTLE: Continuously Anonymizing Data Streams," IEEE Trans. Dependable and Secure Computing, vol. 8, no. 3, pp. 337-352, May-June 2011.

[5] T.-H. Hubert Chan, E. Shi, and D. Song, "Private and Continual Release of Statistics," Proc. 37th Int'l Colloquium Conf. Automata, Languages and Programming (ICALP '10), pp. 405-417, 2010.

[6] W. Du, Z. Teng, and Z. Zhu, "Privacy-MaxEnt: Integrating Background Knowledge in Privacy Quantification," Proc. ACM SIGMOD Int'l Conf. Management of Data (SIGMOD '08), pp. 459- 472, 2008.

[7] C. Dwork, "Differential Privacy," Proc. 33rd Int'l Colloquium on Automata, Languages and Programming (ICALP '06), pp. 1-12, 2006.

[8] C. Dwork, M. Naor, T. Pitassi, and G.N. Rothblum, "Differential Privacy under Continual Observation," Proc. 42nd ACM Symp. Theory of Computing (STOC '10), pp. 715-724, 2010.

[9] G. Di Biase et al., "A Stochastic Model for the HIV/AIDS Dynamic Evolution," Math. Problems in Eng., 2007.

[10] J.-L. Fuh, R.-F. Pwu, S.-J. Wang, and Y.-H. Chen, "Measuring Alzheimer' s Disease Progression with Transition Probabilities in the Taiwanese Population," Int'l J. Geriatric Psychiatry, vol. 19, no. 3, pp. 266-270, 2004.

[11] G. Ghinita, P. Karras, P. Kalnis, and N. Mamoulis, "Fast Data Anonymization with Low Information Loss," Proc. 33rd Int'l Conf. Very Large Data Bases (VLDB '07), pp. 758-769, 2007.

[12] D. Kifer and A. Machanavajjhala, "No Free Lunch in Data Privacy," Proc. Int'l Conf. Management of Data (SIGMOD '11), pp. 193-204, 2011.

[13] K. LeFevre, D.J. DeWitt, and R. Raghu, "Mondrian Multidimensional k-Anonymity," Proc. 22nd Int'l Conf. Data Eng. (ICDE '06), 2006.

- [14] J. Li, B.C. Ooi, and W. Wang, "Anonymizing Streaming Data for Privacy Protection," Proc. IEEE 24th Int'l Conf. Data Eng. (ICDE '08), pp. 1367-1369, 2008.
- [15] N. Li, T. Li, and S. Venkatasubramanian, "t Closeness: Privacy beyond k-Anonymity and Diversity," Proc. IEEE 23rd Int'l Conf. Data Eng. (ICDE '07), pp. 106-115, 2007.
- [16] T. Li and N. Li, "Injector: Mining Background Knowledge for Data Anonymization," Proc. IEEE 24th Int'l Conf. Data Eng. (ICDE '08), pp. 446-455, 2008.
- [17] T. Li, N. Li, and J. Zhang, "Modeling and Integrating Background Knowledge in Data Anonymization," Proc. IEEE 25th Int'l Conf. Data Eng. (ICDE '09), pp. 6-17, 2009.
- [18] J. Lin, "Divergence Measures based on the Shannon Entropy," IEEE Trans. Information Theory, vol. 37, no. 1, pp. 145-151, Jan. 1991.
- [19] A. Machanavajjhala, D. Kifer, J. Gehrke, and M. Venkatasubramanian, "l-Diversity: Privacy Beyond k-Anonymity," ACM Trans. Knowledge Discovery from Data, vol. 1, no. 1, article 3, 2007.
- [20] D.J. Martin, D. Kifer, A. Machanavajjhala, J. Gehrke, and J.Y. Halpern, "Worst-Case Background Knowledge for Privacy- Preserving Data Publishing," Proc. 23rd Int'l Conf. Data Eng. (ICDE '07), pp. 126-135, 2007.
- [21] A. Meyerson and R. Williams, "On the Complexity of Optimal k-Anonymity," Proc. 23rd ACM SIGMOD-SIGACT-SIGART Symp. Principles of Database Systems (PODS '04), pp. 223-228, 2004.