

Mine Me but Don't Single Me Out: Differentially Private Event Logs for Process Mining

Gamal Elkoumy
University of Tartu, Tartu, Estonia
gamal.elkoumy@ut.ee

Alisa Pankova
Cybernetica, Tartu, Estonia
alisa.pankova@cyber.ee

Marlon Dumas
University of Tartu, Tartu, Estonia
marlon.dumas@ut.ee

Abstract—The applicability of process mining techniques hinges on the availability of event logs capturing the execution of a business process. In some use cases, particularly those involving customer-facing processes, these event logs may contain private information. Data protection regulations restrict the use of such event logs for analysis purposes. One way of circumventing these restrictions is to anonymize the event log to the extent that no individual can be singled out using the anonymized log. This paper addresses the problem of anonymizing an event log in order to guarantee that, upon disclosure of the anonymized log, the probability that an attacker may single out any individual represented in the original log, does not increase by more than a threshold. The paper proposes a differentially private disclosure mechanism, which oversamples the cases in the log and adds noise to the timestamps to the extent required to achieve the above privacy guarantee. The paper reports on an empirical evaluation of the proposed approach using 14 real-life event logs in terms of data utility loss and computational efficiency.

Index Terms—Process Mining, Event Log, Differential Privacy

I. INTRODUCTION

Process Mining is a family of techniques that helps organizations enhance the performance, conformance, and quality of their business processes. The input of process mining techniques is an event log. An event log captures the execution of a set of instances of a process (herein called *cases*). An event log consists of event records. Each record contains a reference to a case identifier, a reference to an activity, and at least one timestamp. Table I shows an example of an event log of a healthcare process.¹ Each case ID refers to a patient, and each activity label corresponds to a treatment or an event in the healthcare system.

Often, an event log contains private information about some individuals. Data regulations, such as the General Data Protection Regulation (GDPR)², restrict the use of such event logs for analysis purposes. One way to overcome these restrictions is to anonymize the event log such that no individual can be singled out. Singling out an individual happens when they can be distinguished, within a group of people, by evaluating a predicate that discriminates him/her. The legal notion of singling out has been mathematically formalized by Cohen & Nissim [1], who define a type of attack called *Predicate Singling Out (PSO) attack*. The disclosure of the event log in Table I permits singling out individuals. Specifically, the

TABLE I
EVENT LOG

Case	Activity	Timestamp	Source State	Target State
1	A	2020-08-08 10:20:0.000	s_0	s_5
1	B	2020-08-08 10:50:0.000	s_5	s_2
1	C	2020-08-08 16:15:0.000	s_2	s_3
2	D	2020-08-08 12:07:0.000	s_0	s_4
2	A	2020-08-08 13:37:0.000	s_4	s_5
2	E	2020-08-08 14:07:0.000	s_5	s_2
2	C	2020-08-08 19:07:0.000	s_2	s_3
3	A	2020-08-08 13:30:0.000	s_0	s_5
3	B	2020-08-08 13:55:0.000	s_5	s_2
3	C	2020-08-08 20:55:0.000	s_2	s_3
4	D	2020-08-08 15:00:0.000	s_0	s_4
4	A	2020-08-08 17:00:0.000	s_4	s_5
4	B	2020-08-08 17:40:0.000	s_5	s_2
4	C	2020-08-08 23:45:0.000	s_2	s_3
5	A	2020-08-08 16:40:0.000	s_0	s_5
5	E	2020-08-08 17:55:0.000	s_5	s_2
5	C	2020-08-08 23:55:0.000	s_2	s_3

predicate “undergoing treatment E after treatment D and A” and the time difference between these activities can lead to a linkage attack [2]. Given that prior knowledge, the adversary can single out the patient with case ID 2.

Privacy-Enhancing Technologies (PETs), such as differential privacy and k-anonymity [3], protect datasets disclosure, including event logs. Among PETs, differential privacy stands out because it mitigates PSO attacks [1] and due to its composability guarantees [4]. Differential privacy mechanisms inject noise into the data, quantified by a parameter called ϵ . Lee et al. [5] show that “the proper value of ϵ varies depending on individual values”, and that the existence of “outliers also changes the appropriate value of ϵ ”. Dwork et al. [6] state that “we do not know what parameter ϵ is right for any given differentially private analysis, and we do know that the answer can vary tremendously based on attributes of the dataset and the policies and practices that constrain those who query it.” Accordingly, this paper proposes a method to determine the needed ϵ value to disclose an event log in terms of a business-related metric, namely guessing advantage, which captures the increase in the probability that an adversary may guess information about an individual after the disclosure. The probability of leakage is a widely used measure of risk which can be interpreted on its own.

Usually, an adversary has prior knowledge about the individuals in the log before its release. Therefore, the adversary gains an additional advantage to guess personal information following the disclosure successfully. The goal of anonymization is to limit this risk. In this paper, we deal with the problem of limiting the additional risk by a maximum guessing advantage

¹The columns “Source State” and “Target State” are explained later

²<http://data.europa.eu/eli/reg/2016/679/oj>

level δ . Specifically, we address the following problem:

Given an event log L , and given a maximum level of acceptable guessing advantage δ , generate an anonymized event log L' such that the success probability of singling out an individual after publishing L' does not increase by more than δ .

Naturally, we should ensure that the anonymized log is useful for process mining. In other words, that process mining algorithms return similar results with the anonymized log as with the original log. In this respect, a desirable property is that the anonymized log should have the same set of *case variants* as the original one. A case variant is a distinct sequence of activities. For example, the case variants of the log in Table I are $\{\langle A, B, C \rangle, \langle D, A, E, C \rangle, \langle D, A, B, C \rangle, \langle A, E, C \rangle\}$. This property ensures, for example, that the set of directly-follows relations between activities is not altered during anonymization. This set of relations, known as the Directly-Follows Graph, is used by automated process discovery techniques [7]. Moreover, the set of case variants is the main input used by conformance checking techniques. In this setting, having the same set of case variants is critical, as every case variant that is added to the log directly is a potential false positive (a deviation that does not exist in reality but it is reported as such) while every removed case variant is a potential false negative (a deviation that occurs in reality but that is not detected when using the anonymized log). A second desirable property is that the differences between the timestamps of consecutive events in the anonymized log are as close as possible to those in the original log, as these time differences are used by performance mining techniques [7]. Accordingly, we tackle the above problem subject to the following requirements:

- R1** The anonymized event log must have the same set of case variants as the original log.
- R2** The difference between the real and the anonymized time values is minimal given the risk metric δ .

The first requirement should be tackled in a way that keeps the variant frequencies and the log size close to the original log. The second requirement can be tackled w.r.t. different attack models. In this paper, we consider an attack model wherein the attacker seeks to learn a prefix or a suffix of an individual's trace or the timing of a given activity for a given individual.

In this paper, we assume that the activity labels are public. Under this starting point, we tackle the above problem by defining a notion of the differentially private event log. Given a maximum allowed guessing advantage, δ , a differentially private event log is obtained by oversampling the traces in the log and injecting noise to the event timestamps. This ensures that the probability that an attacker may single out any individual, based on the prefixes/suffixes of the individual's trace or based on the event timestamps, is not more than δ .

The proposal relies on a data structure that compactly captures all prefixes and suffixes of a set of traces, namely a Deterministic Acyclic Finite State Automata (DAFSA) [8]. A DAFSA is a lossless representation of an event log, wherein every prefix or suffix shared by multiple traces is represented once. By analyzing the frequency and time differences of each DAFSA transition, we determine the amount of oversampling

and timestamp noise. The estimated noise limits the guessing advantage of an attacker by inspecting the traces traverse this transition in the anonymized log.

The paper is structured as follows. Sect. II introduces background notions and related work. Sect. III presents the notion of the differentially private event log, the attack model, and the risk quantification approach. Sect. IV translates the risk quantification into an algorithm that anonymizes a log. Finally, Sect. V presents an empirical evaluation while Sect. VI draws conclusions and discusses future work.

II. BACKGROUND AND RELATED WORK

A. Differential Privacy

A database D is a set of attributes with values drawn from a universe U . A tuple is an instance in the database, which is a collection of values of a set of attributes $A = A_1, A_2, \dots, A_m$, where m is the number of attributes in D . A tuple corresponds to an individual who requires their privacy to be maintained. An individual may contribute to more than one tuple.

A mechanism $M : D \rightarrow \text{Range}(M)$ maps a database D to a particular distribution of values $\text{Range}(M)$ (e.g., to a vector of real numbers). A privacy mechanism M can be either unbounded or bounded ϵ -differentially private (ϵ -DP). An unbounded ϵ -DP mechanism makes it hard to distinguish two databases that differ in the *presence* of one tuple [4]. On the other hand, a bounded ϵ -DP mechanism makes it hard to distinguish two databases that differ in the *value* of one tuple.

Definition 2.1 (Unbounded ϵ -differentially private mechanism [4]): A mechanism M is said to be ϵ -differentially private if, for all the data sets D_1 and D_2 differing at most on one item, and all $S \subseteq \text{Range}(M)$, we have $\Pr[M(D_1) \in S] \leq \exp(\epsilon) \times \Pr[M(D_2) \in S]$.

Definition 2.2 (Bounded ϵ -differentially private mechanism [4]): A mechanism M is said to be ϵ -differentially private if, for all the data sets D_1 and D_2 differing at most on the value of one item, and all $S \subseteq \text{Range}(M)$, we have $\Pr[M(D_1) \in S] \leq \exp(\epsilon) \times \Pr[M(D_2) \in S]$.

In some cases, it is desired to apply differential privacy to only values of a particular attribute A , e.g., the attribute timestamp in Table I. We apply differential privacy to D_1 and D_2 w.r.t the attribute A , i.e., D_1 and D_2 differ only on A 's value in a single tuple. Moreover, we want to take into account the particular *amount of change* in attribute A .

Definition 2.3 (Bounded ϵ -differentially private mechanism w.r.t. attribute): A mechanism M satisfies ϵ -differential privacy if all the data sets D_1 and D_2 differing at most on the value of one attribute A for one tuple in D_1 and D_2 , and all $S \subseteq \text{Range}(M)$, we have $\Pr[M(D_1) \in S] \leq \exp(\epsilon \cdot |D_1.A - D_2.A|) \times \Pr[M(D_2) \in S]$.

The ϵ -differential privacy restricts the ability to single out an individual (Def. 2.1 and 2.2) or disclose an individual's private attribute (Def. 2.3). We consider an interactive mechanism [4], where a user submits a query function f to a database and receives a noisified result. Formally, there is a mechanism M_f that computes f and injects noise into the result. The amount

of noise depends on the *sensitivity* of f , which quantifies how much change in the input of f affects change in its output.

Definition 2.4 (Global Sensitivity): Let $f : D \rightarrow \mathbb{R}^d$.

- Global sensitivity w.r.t. presence of a tuple is $\Delta f = \max_{D_1, D_2} |f(D_1) - f(D_2)|$;
- Global sensitivity w.r.t. attribute A is $\Delta^A f = \max_{D_1, D_2} \frac{|f(D_1) - f(D_2)|}{|D_1.A - D_2.A|}$,

where max is computed over all databases D_1, D_2 differing in one item at most.

Given the database D and the query function f , a randomized mechanism M_f returns a noisified output $f(D) + Y$, where Y is a value drawn randomly from a particular distribution. E.g., we can draw values from a Laplace distribution $Lap(\lambda, \mu)$, which has a probability density function $\frac{1}{2\lambda} \exp(-\frac{|x-\mu|}{\lambda})$, where λ is a scale factor, and μ is the mean. It is known [4] that, for real-valued f , if we set $\mu = 0$, for $\lambda = \frac{\Delta f}{\epsilon}$ we obtain an ϵ -DP mechanism w.r.t. row presence (Def. 2.1), and for $\lambda = \frac{\Delta^A f}{\epsilon}$, we obtain an ϵ -DP mechanism w.r.t. attribute A (Def. 2.3).

The privacy parameter ϵ ranges from 0 to ∞ , and the desired level of privacy depends on the data distribution. This brings us to the question of how much ϵ is enough [5]? Laud et al. [9] propose a framework to quantify ϵ from a probability value called the *guessing advantage*, which is the increase in the probability that an adversary reveals the value after the data disclosure. In this paper, we adopt the work by Laud et al. [9] to provide a quantification of ϵ based on the guessing advantage threshold, with the assumption that an adversary has background knowledge about all the other instances.

B. Privacy-Preserving Process Mining

The use of PETs for privacy-preserving process mining (PPPM) has been considered in previous studies [3]. PRETSA [10] and TLKC [2] propose k-anonymity mechanisms to anonymize event logs. These studies do not fulfill **R1** because the case variants in the anonymized log are different from the original log due to the suppression of cases or events within cases. In one example in [2], TLKC suppressed 87% of the activities in the output. Furthermore, k-anonymity does not fully mitigate PSO attacks [1].

Mannhardt et al. [11] define a differential privacy mechanism to anonymize two types of queries: the query “frequencies of directly-follows relations” and “frequencies of trace variants”. PRIPEL [12] propose timestamp shifts to anonymize the timestamp attribute of the event log. These two studies do not address the problem stated in Sect. I because they do not limit the guessing advantage to a certain threshold. Instead, the user has to provide an ϵ value as an input. Furthermore, these approaches do not fulfill **R1**, as the set of case variants in the anonymized log is different from the original.

Other related work of PPPM is orthogonal to our research. For example, privacy quantification methods have been proposed in [13] and [14]. Batista & Solanas [15] present an approach that anonymizes event logs w.r.t. resources. Furthermore, other studies addressed the secure processing of distributed event logs [16], [17].

III. DIFFERENTIALLY PRIVATE EVENT LOG

In this section, we introduce the concepts of event log and DAFSA. We then define a differentially private event log notion, and we present our attack model and risk quantification.

A. Event Log Representation

An event log is a set of event records (events for short) capturing the execution of activities of a process. Each event contains a unique identifier of the process instance in which it occurs (case ID), an activity label, and a timestamp. In addition, an event may contain other attributes, e.g., resources. This paper focuses on anonymizing three attributes: case ID, activity label, and timestamp. If we group the events in an event log by case ID, and we sort the events in each group chronologically, then every resulting group is called a trace. A trace captures the sequence of events that occurred in a case.

Definition 3.1 (Event Log, Event, Trace): An event log $L = \{e_1, e_2, \dots, e_n\}$ of a process is a set of events $e = (i, a, ts)$, each capturing an execution of an activity a , with a timestamp ts , as part of a case i of the process. The trace $t = \langle e_1, e_2, \dots, e_m \rangle$ of a case i is the sequence of all events in L with identifier i , ordered by timestamp.

An event log L may be represented as a set of traces $\{t_1, t_2, \dots, t_k\}$. If we set aside the case ID and the timestamp of a trace and focus on the activity labels, we can represent a trace as a word over the alphabet of activity labels. Each particular word extracted from a log in this way is called a *case variant* of the log.

Definition 3.2 (Case Variant): Given an event log L , a case variant of L is a sequence of activity labels $\langle a_1, a_2, \dots, a_k \rangle$ such that there is a trace $\langle e_1, e_2, \dots, e_k \rangle$ of L such that $\forall j \in [1..k] \text{ id}(e_j) = v_j$, where $\text{id}(e)$ is the case ID of the event e .

We assume that each trace in a log corresponds to the execution of the process pertaining to an individual whose privacy we wish to safeguard. Specifically, our goal is to mitigate singling out of an individual based on any prefix or suffix of their trace. To achieve this, we group the prefixes and suffixes in the log, and we inject independent differentially private noise to each such group. For this, we need a representation of the log that partitions the prefixes and suffixes of the log traces into groups. In other words, this representation should assign each prefix (suffix) in the log to one group such that the union of the groups is equal to the entire set of prefixes (suffixes). In addition, we also require that this representation preserves the set of case variants of the log (cf. **R1**).

The DAFSA provides such a partitioning. Given a set of words, each state in the DAFSA represents a group of prefixes that share the same set of suffixes, and suffixes that share the same set of prefixes [8]. An advantage of the DAFSA over similar representations, such as tries, is that a DAFSA contains a minimal number of groups (states) [8]. By minimizing the number of groups, we obtain larger groups. The larger the group, the smaller the noise generated by differential privacy.

Given a DAFSA constructed from a set of words, every word is a path from an initial state to a final state. Conversely, every path from an initial state to a final state corresponds to

a word in the given set of words [8]. Reissner et al. [18] reuse the algorithm in [8] to represent a log as a DAFSA. Every trace is seen as a word (its corresponding case variant). For example, the DAFSA of the log in Table I is shown in Fig. 1.

Definition 3.3 (DAFSA [8]): Let V be a finite non-empty set of (activity) labels. A DAFSA is an acyclic and deterministic directed graph $D = (S, s_0, A, S_f)$, where S is a finite set of states, $s_0 \in S$ is the initial state, with a set of arcs $A \subset S \times V \times S$, and a set of final states S_f .

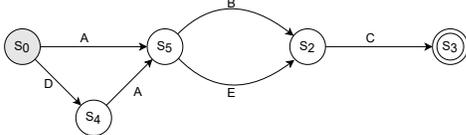


Fig. 1. DAFSA of the event log in Table I

The common prefixes of the DAFSA in Fig. 1 are $\{\langle A, B \rangle \langle D, A \rangle \langle A, A \rangle\}$, and the common suffixes are $\{\langle B, C \rangle \langle E, C \rangle\}$. Cases corresponding to case variants that traverse a given DAFSA state s share the same set of prefixes and suffixes. In this paper, we employ DAFSA states and transitions to group common prefixes and suffixes within cases of the log. We annotate the log with the DAFSA transitions and states to relate such grouping to the event log cases.

Definition 3.4 (State Annotated Event Log): A state annotated event log $L_s = \{r_1, r_2, \dots, r_n\}$ is a set of entries $r = (i, a, ts, s_i, s_e)$, each links an event $e = (i, a, ts) \in L$, where L is the event log, to the DAFSA transition $t = (s_i, a, s_e)$ that represents the occurrence of that event, where t starts from s_i , ends at s_e , and labeled with the same activity a .

The state-annotated event log links every event in the event log, based on its prefix, suffix, and activity label, to the DAFSA transition. Every event is labeled by the source state and target state of the DAFSA transition. Table I (columns source state and target state) shows the state annotated event log.

B. Privacy Mechanism

Given Definitions 3.3, and 3.4, all cases that share a common prefix will traverse a given state s in the DAFSA corresponding to this prefix. The same holds for cases that share a common suffix. We quantify the privacy parameter ϵ using the state annotated event log to mitigate singling out an individual by their prefixes (or suffixes). To study the histogram (count) distribution of the common prefixes and suffixes between traces of an event log, we construct the *DAFSA transition contingency table*. A contingency table is a histogram of group occurrences.

Definition 3.5 (DAFSA Transitions Contingency Table): The DAFSA transition contingency table C is the histogram of counts for each transition $t = (s_i, a, s_e)$ of the DAFSA D , where s_i is the source state, a is the activity label, and s_e is the target state of t .

TABLE II
DAFSA TRANSITIONS CONTINGENCY TABLE

Source State	Activity	Target State	Count
s_0	A	s_5	3
s_5	B	s_2	3
s_2	C	s_3	5
s_0	D	s_4	2
s_4	A	s_5	2
s_5	E	s_2	2

Table II shows the DAFSA transitions contingency table. These counts are usually called marginals. The marginals contain the correlations counts of the common sets of prefixes and suffixes of the DAFSA. We anonymize these marginals to prevent singling out that an individual has been through an activity, using the prefix and the suffix set of activities. The anonymization of the marginals can be directly mapped to the privacy-preserving Online Analytical Processing [19].

Definition 3.6 (differentially private DAFSA Transitions Contingency Table): Let f be a query function that computes a DAFSA transitions contingency table that has a set of transitions $t = (s_i, a, s_e)$ and a count cell c_i for each transition. Let M_f be an unbounded ϵ -differentially private mechanism (by Def. 2.1) that injects noise into the result of f . The differentially private DAFSA transitions contingency table is: $M(C) := \{(t_1, M_f(c_1)), (t_2, M_f(c_2)), \dots, (t_n, M_f(c_n))\}$, where n is the number of DAFSA transitions.

Our goal is not to anonymize a DAFSA but an *event log*. We assume that case IDs have been pseudonymized, the activity labels are public, and the individual cases are independent. We propose a mechanism M that anonymizes two properties of an event log: the activity timestamp and the set of prefixes and suffixes of activity. We can release the time for the first property by applying a bounded ϵ -DP mechanism w.r.t. timestamp attribute (Def. 2.3). We need a mechanism that makes the contingency table ϵ -DP (Def. 3.6) for the second property.

Definition 3.7 (differentially private Event Log): Let L be an event log as defined in Def. 3.1. We say that a log $M(L)$ is ϵ -differentially private if: (1) it is ϵ -differentially private w.r.t. timestamp (Def. 2.3); and (2) its lossless representation is ϵ -differentially private (Def. 3.6).

The mechanism M_f of Def. 3.6 operates on counts and cannot be applied to a log directly. We need to translate the noise injection of the contingency table to the log. Kifer et al. [20] present the notion of a *move* as a way of anonymizing the marginals of contingency tables. A move is a process that adds or deletes a tuple from the contingency table. In order to fulfill the requirement **R1**, we define *oversampling* as increasing a count (positive move) in the contingency table by replicating a random tuple in the log.

Definition 3.8 (DAFSA Transition Oversample): Given a DAFSA transition contingency table C_i , an oversample O is a transformation that adds a DAFSA transition instance to C_i , producing a contingency table $C_j = O(C_i)$, with an increase of only one count cell by 1.

We define a mechanism M that oversamples tuples of a log L so that, if we computed the contingency table of $M(L)$, it would be ϵ -DP. We must be careful that $M(L)$ does not leak anything that the contingency table would not leak. For example, if we keep timestamps unchanged, it suffices to remove the tuples with repeating timestamps to eliminate duplicates, so we need to make the times ϵ -DP as well. Since duplicated timestamps are correlated, we have to divide the ϵ values of the time queries of the replicated cases by their oversampling ratio. To keep the counts in the DAFSA transitions consistent, we oversample the prefix and the suffix of the oversampled

transition, i.e. we always oversample entire cases.

Definition 3.9 (Case Oversample): Given an event log L , a case oversample O_c is a transformation that duplicates a case c_i of log L , in such a way that c_j , the duplicated case in log $O_c(L)$, and c_i have the same sequence of activities.

C. Risk Quantification and Attack Model

Given Def. 3.7, we seek to quantify ϵ 's value to publish a differentially private log. We consider a scenario where an organization shares its event log with an analyst, who should not be able to infer an individual's information. The analyst has background knowledge about all the individuals in the log except for the individual of their interest. Specifically, we consider the case where an attacker (the analyst) uses the event log to single out an individual. The attacker has a goal $h(L)$ that captures their interest in an event log L . We specifically consider the following attacker's goals:

- h_1 : Has the individual been through a specific subtrace (prefix or suffix)? The output is a bit with a value $\in \{0, 1\}$ that represents yes or no.
- h_2 : What is the execution time of a particular activity that has been executed for the individual? The output is a real value to be guessed with precision.

We use an ϵ -differential privacy mechanism to mitigate the above attacker's goals. To determine the privacy parameter ϵ , we adopt the guessing advantage framework [9], which quantifies the risk of publishing a dataset as the difference between two probabilities: the prior and the posterior guessing probability. Even without publishing the event log, attackers can use their knowledge to guess information about a specific individual. The guess is considered successful if it falls within a range of values H_p , which is the actual value \pm a precision.

Definition 3.10 (Prior Guessing Probability): An attacker's prior guessing probability is defined as $P := Pr[h(L) \in H_p]$

A guessing precision p is a percentage value representing the range of a successful guess H_p . For example, if the true value is 0.5 and $p = 0.2$, the guessed value is successful if it falls in range $H_p = [0.3..0.7]$. Our method pre-processes the log to normalize the values (relative timestamps) to be between 0 and 1. The precision is interpreted within this [0,1] range.

The guessing advantage is defined as the difference between the posterior probability (after publishing $M(L)$) and the prior probability (before publishing $M(L)$) of an attacker making a successful guess in H_p . Let δ be the maximum allowed guessing advantage, stated by the event log publisher.

Definition 3.11 (Guessing Advantage): Attacker's advantage in achieving the goal h with precision p is at most δ if, for any published event log L' ,

$$Pr[h(L') \in H_p \mid M(L) = L'] - Pr[h(L) \in H_p] \leq \delta.$$

Laud et al. [9] quantify the maximum value of ϵ , that achieves the upper bound δ .

Proposition 1: The maximum possible ϵ (i.e., the minimum noise) that achieves the upper bound δ , w.r.t. the above attack model, and fulfills the requirement **R2** is

$$\epsilon_k = -\ln \left(\frac{P_k}{1 - P_k} \left(\frac{1}{\delta + P_k} - 1 \right) \right) \cdot \frac{1}{r}, \quad (1)$$

where r is the maximum value in the range of values ($r = 1$ after the normalization of values), and P_k is the prior guessing probability for an instance k .

The proof of Prop. 1 is available in the supplementary material [21], where it appears with heading "Proposition 2".

Given Prop. 1, if there is enough data, we can estimate P_k based on the data distribution as:

$$P_k = CDF(t_k + p \cdot r) - CDF(t_k - p \cdot r), \quad (2)$$

where t_k is the value of an instance, and p is the precision. Suppose we cannot estimate the probability distribution of input values (e.g., the likelihood of participating in a subtrace). In that case, we can compute the worst-case scenario P_k as [9]:

$$P_k = (1 - \delta)/2 \text{ for all } k. \quad (3)$$

IV. COMPUTING DIFFERENTIALLY PRIVATE EVENT LOGS

Given the above definitions, this section translates the guessing advantage parameter δ , into an algorithm that anonymizes the event log. Fig. 2 outlines the proposed approach. First, we construct the DAFSA state annotated event log as mentioned in Sect. III-A. We use the algorithm proposed by Reissner et al. [18] to build the DAFSA. Second, we calculate the ϵ values needed by the approach. Finally, we use the calculated ϵ values to perform oversampling of cases and time noise injection into the event log. Below, we describe the steps in detail.

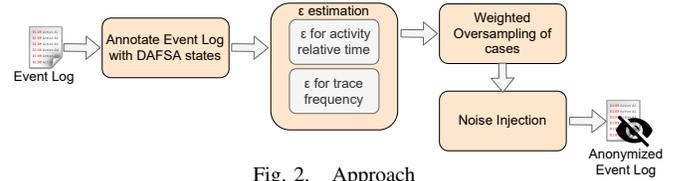


Fig. 2. Approach

A. ϵ Estimation

Given the DAFSA annotated event log, we group the common prefixes and suffixes in traces by the DAFSA transitions. This grouping has two outputs: a distribution of execution time of activity instances that goes through the same DAFSA transitions and the DAFSA transitions contingency table. A case's time attribute comprises two components: the start time of the case and the relative execution time of every activity in the case (a time difference between its execution timestamp and the start timestamp of the case) in a time unit.

To calculate ϵ for the distribution of execution time instances, grouped by a DAFSA, we use Eq (2) and (1), i.e., we use different ϵ values for every event in the event log. We normalize the input values to become in the range [0, 1] as mentioned in Sect. III-C. Eq (1) provides the maximum ϵ to fulfill **R2**.

On the other hand, the DAFSA transition contingency table contains a count histogram of the common prefixes and suffixes. To calculate the ϵ value for the marginals, we use Eq (3) and (1). The ϵ of all the DAFSA transitions of the contingency table is the same, but the noise is drawn for every transition independently. The propositions and the privacy proofs of ϵ estimation are presented in the supplementary material [21].

B. Weighted Oversampling of Cases

Given the above-calculated ϵ for trace frequency, we draw a random noise value from the Laplace distribution $Lap(\Delta f/\epsilon)$, for each DAFSA transition. Adding or removing a prefix/suffix of a trace to the log affects a single frequency count by 1, so $\Delta f = 1$. We take the absolute value of the noise (additive noise) to fulfill **R1**. The noise (quantified by δ) is added as increments to the frequencies of the contingency table. Using an absolute value (additive noise only), we alter the privacy guarantees of the differentially private mechanism, as the attacker can break Def. 2.1 if the sampled noise falls within $[0, \Delta f]$. This happens with probability $1 - \exp(-\epsilon/\Delta f)$. This additional risk is factored into δ .

Next, we increment the counts of the contingency table by means of oversampling of DAFSA transitions as defined in Def. 3.8. Finally, to maintain the consistency of the DAFSA transitions, we oversample the traces in the log using Def. 3.9, as presented in Algorithm 1. This algorithm does not suppress nor add case variants to fulfill **R1**. The supplementary material [21] provides the correctness proofs of Algorithm 1 (cf. Proposition 5, 6, and 7, and Theorem 1).

Algorithm 1: Case Oversampling Algorithm

```

1. Input: Event Log,  $\epsilon$ , DAFSA
2. Output: Differentially Private Event Log
3. DafsaLookup = Build DAFSA transitions to case variant lookup ;
4. DafsaLookup[i].neededNoise =  $|z_i|$ , where  $z_i$  is sampled from
    $Lap(\Delta f/\epsilon)$  independently for every transition  $t_i$ ;
5. cnt = count(DafsaLookup.addedNoise < DafsaLookup.neededNoise) ;
6. while cnt > 0 do
7.   selectedTransition = pick a random transition such that
     DafsaLookup.addedNoise < DafsaLookup.neededNoise;
8.   pickedTraces = pick  $x$  random traces that traverse
     selectedTransition, where  $x$  = selectedTransition.neededNoise;
9.   foreach  $t \in$  pickedTraces do
10.    DafsaLookup[t].addedNoise ++ ;
11.    replicate a random case with a case variant =  $t$  ;
12.   end
13.   cnt = count(DafsaLookup.addedNoise <
     DafsaLookup.neededNoise);
14. end
15. generate new CaseID for every case;
16. shuffle cases ;

```

Algorithm 1 starts by constructing a correspondence (lookup) table between the DAFSA transitions and the case variants (line 3). This table maps every DAFSA transition to the case variants that traverse it. We use this lookup table to track the updates over transitions. Second, we draw a random noise from the Laplace distribution $Lap(\Delta f/\epsilon)$ (line 4) independently for every transition. Next, we count the DAFSA transitions that need noise injection and their needed noise (lines 5 and 6). Next, we pick a random transition that needs noise (with sampling weights of their occurrence frequency) (line 7). Then, we randomly choose a case variant that goes through the chosen transition (with sampling weights of their number of instances) (line 8). We replicate the chosen case variant by a number of times equals the needed noise (lines 9-11). For every replication, we choose a random case variant instance from the log to replicate. Next, we update the DAFSA lookup with the injected noise (line 10). We repeat this process until

all the transitions have the minimum required noise. We keep ϵ values of the time attribute the same as the original cases to draw noise as in Sect. IV-C. Then, we generate new case IDs for the cases (line 15), and we change their order (line 16).

C. Noise Injection

At this step, we have the DAFSA annotated event log, with the oversampled case instances and an ϵ value of each event. First, we divide the ϵ value of the replicated cases by the number of replications, as oversampling is considered repeating the same query more than once [4]. Then, we draw a random noise from the Laplace distribution $Lap(\Delta f/\epsilon)$ to anonymize the relative time for every activity instance. Next, we transform the amount of noise from the normalized range ($[0, 1]$) to the original range. Finally, we add the time noise to the original execution relative time. After the noise injection, we transform the relative execution time of activities to timestamps. By the end of this step, the event log is anonymized by ϵ values calculated for the input maximum guessing advantage δ .

V. EVALUATION

To address the problem stated in Sect. I under requirements **R1** and **R2**, the proposed method injects differentially private noise in two ways: (i) by oversampling some traces in the log; and (ii) by altering the event timestamps. The effect of the oversampling can be measured by means of an oversampling ratio, as defined below. Meanwhile, the alteration of timestamps leads to the anonymized log having a longer timespan than the original one, i.e., a time dilation. Below, we evaluate the method by studying the effect of its parameters (δ and p) on the time dilation and the oversampling ratio.

Accuracy Measure. We measure the time dilation by comparing the anonymized logs with the original event log as ground truth. We use the *symmetric mean absolute percentage error* (SMAPE). SMAPE is the absolute difference between true value t_i and the anonymized value a_i divided by the sum of the absolute values of t_i and a_i . $SMAPE = \frac{1}{n} \sum_{i=1}^n \frac{|t_i - a_i|}{|t_i| + |a_i|}$, where n is the number of events in the event log. We measure the effect on case variant frequencies by measuring the *oversampling ratio*. To calculate the oversampling ratio, we divide the number of cases in the anonymized event log m_a by the number of cases in the original event log m . $oversampling\ ratio = \frac{m_a}{m}$.

Event log Selection. For our experiment, we use the real-life event logs publicly available at 4TU Center for Research Data³ as of February 2021. The considered event log files, their characteristics, and descriptive statistics are available in the supplementary material [21]. The selected logs contain the process execution of different domains, e.g., government and healthcare. From the set of available logs, we excluded the event logs that are not business processes (e.g., “Apache Commons”, “BPIC 2016” logs, “Junit 4.12”). Also, we exclude the set of event logs “coSeLog WABO” as they are a pre-processed version of BPI challenge 15. Finally, we select a single event log for each set of logs in BPI challenges 13, 14, 15, 17, 20.

³<https://data.4tu.nl/>

Running Environment. We implement the proposed model as part of a prototype, namely Amun⁴. We use python 3.8.5 and PM4PY⁵ for parsing the XES files. We run the experiment on a single machine with AMD Opteron(TM) Processor 6276 and 32 GB memory. Also, we consider only the end timestamp to calculate the relative time of an event for simplicity, and the same approach is still valid to apply differential privacy.

In the following experiment, we keep the precision $p = 0.1$. We measure the activity execution time in hours. In a practical application, the user decides the time unit. We run the experiments ten times and report the averages. The supplementary material contains an experiment for different values of p .

Results. The results are shown in Table III. We use 14 real-life event logs to study the method’s parameters. The anonymized logs are available in the supplementary material [21] with the estimated ϵ ’s for events. In this table, the oversampling ratio represents the amount of replicated cases to prevent singling out cases using the set of prefixes and suffixes of an activity. The larger the oversampling ratio, the more the cases will be replicated. For instance, the oversampling ratio for the event log CCC19 is large, because every case follows a different case variant. Similarly, BPIC15 has 1199 cases and 1170 case variants, which results in a large oversampling ratio. Meanwhile, all cases in the log Credit Requirement follow the same case variant, resulting in an oversampling ratio of just above 1. Likewise, the Traffic log has a small oversampling ratio as it contains 150 370 cases and only 231 case variants.

In our experiments, the reason for the significant dilation, measured in percentage SMAPE, is outliers in the event logs. For instance, the log BPIC3 contains only one case with a case duration of 2.1 years, and another case duration equals 1.3 years. This unique behavior causes the proposed approach to estimate smaller ϵ values to inject more noise to prevent case identification using their duration. Similarly, the log Sepsis cases contains outliers. Nevertheless, the histogram of case durations of the Hospital log does not show outliers, where the minimum number of cases with a unique case duration is 6. The same case happens with Unrineweginfectie and BPIC19 logs. Furthermore, as per the definition of Differential Privacy, we divide the ϵ by the number of replication. That results in a smaller ϵ value, which means more noise. For instance, the Sepsis cases event log has a large SMAPE percentage because the ϵ values are divided by the oversampling ratio of the replicated case, which results in more significant noise.

The variant analysis of both the Traffic and Sepsis cases event logs, and their anonymized version, with $\delta = 0.2$ is shown in Fig. 3. As reported in Table III, the oversampling ratio of the Traffic log is almost 1, and hence the active cases over time of both the Traffic log and its anonymized version are almost the same. However, there are time shifts due to the noise injected into the event execution time. Fig. 3(b) shows the active case over time, and Fig. 3(d) shows the frequency of case variants. Contrarily, the Sepsis cases log has a large sampling

ratio; hence the anonymized version has more significant active cases over time, as shown in Fig. 3(a) and Fig. 3(c).

We conduct a wall-to-wall run time experiment to assess the efficiency of the method. We measure the time between reading the input XES file and the generation of its anonymized version. The results are reported in of Table III, and the values are in minutes. The run time grows with the growth of case variant (as it contains more DAFSA transition groups) and oversampling ratio (as it needs more oversampling iterations). The reported execution times for logs with a large number of events and for low δ values are in the order of hours, e.g. 9 hours for BPIC18 (2.5 million events) with $\delta=0.2$. This is because the noise injection algorithm iterates multiple times over each transition (lines 6-14 in Algorithm 1) and the number of DAFSA states for this log is high (638,242 states). This shortcoming can be tackled via parallelization, as the privacy quantification over each DAFSA transition is independent of others. The above experiments were all done using one single thread in order to avoid bringing additional variables (number of computing nodes and cores) into the experiments.

We acknowledge that the above observations are based on a limited population of logs (14). However, these logs were selected out of a broader population of close to 50 real-life logs based on justified selection criteria.

VI. CONCLUSION AND FUTURE WORK

This paper proposed a concept of differentially private event log and a mechanism to compute such logs. A differentially private event log limits the increase in the probability that an attacker may learn a suffix of an individual’s trace given a prefix (or vice-versa), or the timestamp of an activity in an individual’s trace. To this end, we inject differentially private noise by oversampling the traces in the log. This approach neither suppresses nor adds case variants (cf. Sect. IV-B) and hence fulfills **R1**. To fulfill **R2**, we quantify ϵ based on a technique that finds the maximum ϵ (minimum noise) that keeps the guessing advantage below δ (cf. Sect.III-C and Prop. 1).

A limitation of the proposed method is that it anonymizes the relative time of each event with respect to the start of the case, but it does not anonymize the start times of the cases. We plan to address this limitation by applying methods for differentially private disclosure of time series. A second limitation is that the input log is assumed to have only three columns: case ID, activity label, and timestamp. Real-world event logs contain other columns, e.g., resources. To address this limitation, we need to extend the approach used to represent event logs, e.g., via multi-dimensional data structures as opposed to DAFSAs.

The proposed method introduces high levels of noise in the presence of unique traces or temporal outliers. To address this limitation, we plan to investigate an approach where high-risk traces are suppressed so that the amount of injected noise into the remaining traces is lower. Although suppression can significantly reduce the required noise level and strengthen the privacy guarantees, it breaks the property that the differentially private log has the same case variants as the original one. To mitigate this drawback, we will seek to define an approach

⁴<https://github.com/Elkoumy/amun>

⁵<https://pm4py.fit.fraunhofer.de/>

TABLE III
SMAPE DILATION, OVERSAMPLING RATIO, AND RUNTIME EXPERIMENTS

Exp. Typ	δ	BPIC12	BPIC13	BPIC14	BPIC15	BPIC17	BPIC18	BPIC19	BPIC20	CCC19	Credited	Hospital	Sepsis	Traffic	Urnine.
SMAPE	0.1	72.5	74.7	62.4	63.2	74.5	64.7	58.7	70.0	47.5	62.6	37.0	70.2	57.9	55.2
	0.2	67.7	70.4	59.5	60.6	69.9	61.8	52.7	66.4	46.3	62.6	36.6	66.4	57.8	51.8
	0.3	64.2	67.6	57.1	58.2	67.1	59.3	48.2	64.3	43.8	62.6	36.2	63.2	57.8	50.9
	0.4	61.3	65.5	55.1	56.2	64.9	57.3	44.5	62.7	44.2	62.6	35.8	60.6	57.8	50.1
	0.5	59.1	63.9	54.0	56.1	63.1	56.2	41.4	61.4	45.0	62.6	35.7	59.4	57.8	49.9
Oversample R	0.1	2.5	1.9	3.8	13.9	1.6	7.1	1.2	2.0	14.7	1.0	14.4	5.2	1.0	1.1
	0.2	1.8	1.5	2.5	7.6	1.4	4.1	1.1	1.5	8.0	1.0	7.7	3.2	1.0	1.1
	0.3	1.6	1.3	2.1	5.4	1.3	3.2	1.1	1.4	5.7	1.0	5.5	2.6	1.0	1.0
	0.4	1.5	1.3	1.9	4.3	1.2	2.6	1.1	1.3	4.4	1.0	4.4	2.3	1.0	1.0
	0.5	1.4	1.2	1.7	3.7	1.2	2.3	1.1	1.3	3.9	1.0	3.6	2.1	1.0	1.0
Time	0.2	12.5	3.9	101.4	8.4	17.4	542.5	87.5	4.1	0.1	1.2	34.1	1.9	8.2	0.2

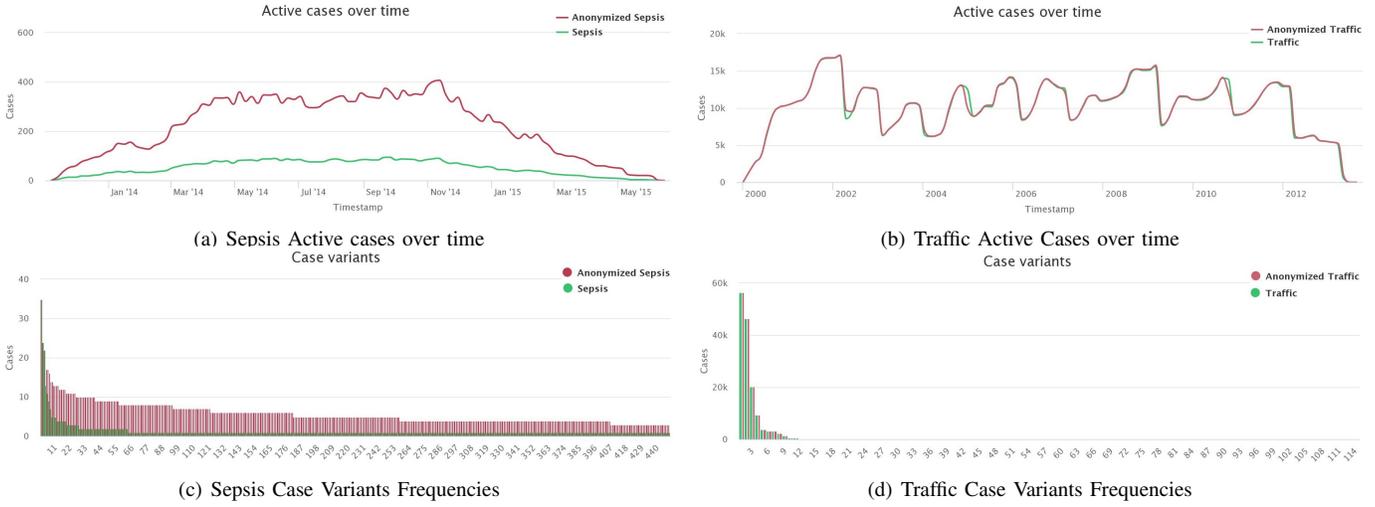


Fig. 3. Variant Analysis comparison between Traffic and Sepsis event logs and their anonymized versions, with $\delta = 0.2$. The anonymized log is in red, and the original log is in green. The case variants are zoomed by 50%.

to minimize the number of case variants that need to be suppressed, given the desired level of guessing advantage.

Acknowledgments Work funded by European Research Council (PIX project) and by EU H2020-SU-ICT-03-2018 Project No. 830929 CyberSec4Europe (<http://cybersec4europe.eu>).

REFERENCES

- [1] A. Cohen and K. Nissim, "Towards formalizing the gdpr's notion of singling out," *Proc. Natl. Acad. Sci.*, vol. 117, pp. 8344–8352, 2020.
- [2] M. Rafiei, M. Wagner, and W. M. van der Aalst, "Tlkc-privacy model for process mining," in *RCIS*. Springer, 2020, pp. 398–416.
- [3] G. Elkoumy, S. A. Fahrenkrog-Petersen, M. F. Sani, A. Koschmider, F. Mannhardt, S. N. von Voigt, M. Rafiei, and L. von Waldhausen, "Privacy and confidentiality in process mining—threats and research challenges," *arXiv preprint arXiv:2106.00388*, 2021.
- [4] C. Dwork, A. Roth *et al.*, "The algorithmic foundations of differential privacy," *Found. Trends Theor. Comput. Sci.*, vol. 9, pp. 211–407, 2014.
- [5] J. Lee and C. Clifton, "How much is enough? choosing ϵ for differential privacy," in *Proc. ISC*. Springer, 2011, pp. 325–340.
- [6] C. Dwork, N. Kohli, and D. Mulligan, "Differential privacy in practice: Expose your epsilons!" *J. Priv. Confidentiality*, vol. 9, no. 2, 2019.
- [7] M. Dumas, M. L. Rosa, J. Mendling, and H. A. Reijers, *Fundamentals of Business Process Management, Second Edition*. Springer, 2018.
- [8] J. Daciuk, S. Mihov, B. W. Watson, and R. E. Watson, "Incremental construction of minimal acyclic finite-state automata," *Computational linguistics*, vol. 26, no. 1, pp. 3–16, 2000.
- [9] P. Laud, A. Pankova, and M. Pettai, "A framework of metrics for differential privacy from local sensitivity," *PoPETs*, pp. 175–208, 2020.
- [10] S. A. Fahrenkrog-Petersen, H. van der Aa, and M. Weidlich, "Pretsa: event log sanitization for privacy-aware process discovery," in *ICPM*. IEEE, 2019, pp. 1–8.
- [11] F. Mannhardt, A. Koschmider, N. Baracaldo, M. Weidlich, and J. Michael, "Privacy-preserving process mining," *BISE*, vol. 61, pp. 595–614, 2019.
- [12] S. A. Fahrenkrog-Petersen, H. van der Aa, and M. Weidlich, "Pripel: Privacy-preserving event log publishing including contextual information," in *BPM*. Springer, 2020, pp. 111–128.
- [13] S. N. von Voigt, S. A. Fahrenkrog-Petersen, D. Janssen, A. Koschmider, F. Tschorsch, F. Mannhardt, O. Landsiedel, and M. Weidlich, "Quantifying the re-identification risk of event logs for process mining," in *CAISE*, 2020, pp. 252–267.
- [14] M. Rafiei and W. M. van der Aalst, "Towards quantifying privacy in process mining," *arXiv preprint arXiv:2012.12031*, 2020.
- [15] E. Batista and A. Solanas, "A uniformization-based approach to preserve individuals' privacy during process mining analyses," *Peer-to-Peer Networking and Applications*, pp. 1–20, 2021.
- [16] G. Elkoumy, S. A. Fahrenkrog-Petersen, M. Dumas, P. Laud, A. Pankova, and M. Weidlich, "Secure multi-party computation for inter-organizational process mining," in *BPMDS*. Springer, 2020, pp. 166–181.
- [17] G. Elkoumy, S. A. Fahrenkrog-Petersen, M. Dumas, P. Laud, A. Pankova, and M. Weidlich, "Shareprom: A tool for privacy-preserving inter-organizational process mining," in *BPM Demos*, 2020, pp. 72–76.
- [18] D. Reißner, R. Conforti, M. Dumas, M. La Rosa, and A. Armas-Cervantes, "Scalable conformance checking of business processes," in *OTM to Meaningful Int. Syst.* Springer, 2017, pp. 607–627.
- [19] B. Barak, K. Chaudhuri, C. Dwork, S. Kale, F. McSherry, and K. Talwar, "Privacy, accuracy, and consistency too: a holistic solution to contingency table release," in *PODS*, 2007, pp. 273–282.
- [20] D. Kifer and A. Machanavajjhala, "No free lunch in data privacy," in *Proc. of ACM SIGMOD*, 2011, pp. 193–204.
- [21] G. Elkoumy, A. Pankova, and M. Dumas, "Mine Me but Don't Single Me Out: Supplementary Material," Mar. 2021. [Online]. Available: <https://doi.org/10.5281/zenodo.5011364>