Differentially Private Transmission Control Protocol Synchronise Packet Counts

Nenekazi Nokuthala P. Mkuzangwe and Fulufhelo Nelwamondo

(Corresponding author: Nenekazi N. P. Mkuzangwe)

Modelling and Digital Science, Council for Scientific and Industrial Research
P. O. Box 395, Pretoria, 0001, South Africa
Department of Electrical and Electronic Engineering Sciences, University of Johannesburg
P. O. Box 524, Johannesburg, 2006, South Africa
(Email: mmkuza@gmail.com)

(Received Mar. 28, 2018; Revised and Accepted Sept. 22, 2018; First Online July 16, 2019)

Abstract

Real Transmission Control Protocol synchronise (TCP SYN) packet counts availability will be of great benefit in anomaly detection. TCP SYN packet counts can be used for training intrusion detection system to detect a denial of service attack called TCP SYN flooding. However, there are privacy and legal issues that limit the direct release of such data to the public. This work aims at providing differentially private TCP SYN packet counts. Utility evaluation indicates that the differentially private counts can be used to make inferences at certain thresholds of the anomaly based detection algorithm with minimum information loss while preserving privacy.

Keywords: Anomaly Detection Algorithms; Differential Privacy; TCP SYN Packets

1 Introduction

Network research (including intrusion detection) depends crucially on the availability of real world traffic traces of network activities. Unfortunately, real world network traces release is highly restricted by privacy and legal issues. Organisations are not willing to share their traces since raw network traces may consists of sensitive information that should not be publicly shared, for example, information that identifies individuals, patterns of the traffic that can be analysed to determine strategies of organisations, hints to the weaknesses of a system, revealing important resources like identifying the busiest machine as a file server, *etc.* [15, 16]. On the other hand unavailability of raw network traces to researchers poses a risk of developing models that compromise accuracy.

To continue with their activities, researchers end up simulating data or signing non-disclosure agreements and these two ways of obtaining data may compromise accuracy and repeatability of the research [16]. Simulated data may not be a true representation of the real life network traffic, hence, using such data in training models may result in models that compromise accuracy. Signing non-disclosure agreements compromise the repeatability of the research since the non-disclosure agreement may be a once off agreement of the use of the data between the data owner and a particular researcher which means other researchers will have no access to that data if they want to repeat the study of that particular researcher.

Real network trace sharing is commonly done through trace sanitisation which includes removal or anonymisation of privacy sensitive packet fields such as payloads and IP addresses. Anonymisation is vulnerable to attacks that infer sensitive information [3]. Mogul and Arlitt [15] proposed an alternative approach to trace anonymisation where data owners perform the analyses in the place of the researchers to preserve privacy, privacy is preserved in this approach based on human verification which is prone to error. To remove human verification Mirkovic [14] proposed rules that an analyst must adhere to in order to preserve privacy. The privacy property provided by these rules is not clear. The existing proposals like in [14, 15], provide no guarantee in protecting sensitive information and therefore a formal privacy guarantying approach, that will make data owners comfortable to adopt before releasing their data, is needed.

The privacy we consider in this work, for publishing data that preserves privacy, is differential privacy. In differential privacy the released aggregates are perturbed by a randomized algorithm so that the outcome of the algorithm remains approximately the same even if any single record in the input data is arbitrarily modified. For example, Laplace mechanism can be used to provide differential privacy by simply adding Laplace perturbation noise to each aggregate statistic. The drawback of releasing a series of aggregates with differential privacy is high perturbation error [10]. For instance, if private data values are aggregated over a long period of time, say T time stamps, a direct application of Laplace mechanism to the aggregates at each time stamp can result in a high overall perturbation error causing the released aggregates to be unusable especially when T is large [10]. To address this drawback [9] have proposed a real-time system with Filtering and Adaptive Sampling for differentially private Time series monitoring (FAST): A novel solution to sharing time series data with differential privacy. FAST has a filter component that models the series using state space model and estimates the original data from the noisy data using Kalman filter where the resulting estimates are released in the place of the noisy perturbed data.

This research aims at the use of differential privacy as a means of providing privacy to network trace. Specifically, the number of Transmission Control Protocol Synchronise packets associated with HTTP requests made to a web server(s) by employees of an organisation on an eight hour working day are released with differential privacy. Differential privacy is used in this work to protect the identity of web browsing employees from being inferred by an adversary from the original number of TCP SYN packets using possible background knowledge about the employees' web browsing patterns. The differential privacy randomisation mechanism called Laplace mechanism is utilised. Laplace mechanism adds noise to the aggregated statistics of the data (the number of TCP SYN packets or TCP SYN counts in this study). Releasing a series of aggregates with differential privacy tends to lead to high perturbation error more especially if the data values are aggregated over a long period [10] and the released perturbed aggregates may end up having less research utility or none. Therefore to improve the accuracy (the closeness to the original aggregates) of the perturbed aggregates in this study, the added noise is reduced (filtered) using the filtering component of [9]. The noise filtered aggregates become the released differentially private aggregates. The research utility of the released aggregates is tested using two utility metrics and by comparing the performance of two anomaly based intrusion detection algorithms on the original aggregates and the released aggregates. The utility measures are used to establish if the inferences made using the released aggregates are close to those reached using the original aggregates.

The contribution of this work is that we are providing privacy preserving network trace called TCP SYN packet counts that are research useful as indicated by the research utility tests conducted in the study.

2 Related Work

This section presents work done in releasing network trace in a privacy preserving manner. Mogul and Arlitt [15] proposed an alternative approach to trace anonymization where the owners of the data perform the analyses in the place of the researchers, *i.e.* researchers ship their code to the owners of the data to preserve privacy. One of the potential drawbacks of the proposed approach, as pointed out by the authors, is that debugging the analysis software will be difficult since the code would have been trained on a different dataset. To remove human verification, Mirkovic [14] proposed rules that an analyst must adhere to in order to preserve privacy. The privacy property provided by these rules is not clear.

Dijkhuizen and Ham [5] conducted a literature survey over the period of 1998-2017 on network traffic anonymisation techniques and their implementation. In the survey

- A brief description of currently available anonymisation techniques and a rough indication of their effectiveness is provided,
- Fields containing privacy sensitive information in the link, internet and transport layers are discussed,
- Existing anonymisation tools and frameworks are described and compared against each other ,
- Future research directions to enable easier sharing of network traffic are provided.

McSherry and Mahajan [12] investigated the potential for network trace analysis while providing the guarantees of differential privacy. Their results show that differential privacy has the potential of being the basis for analysing mediated network trace. Fan *et al.* [8] proposed algorithms that use the rich correlation of the time series of aggregates and estimated the original aggregates from the noisy aggregates (values that are perturbed by a differential privacy mechanism) using the state space approach. They have shown that differentially private aggregates of web browsing activities can be released in real time while preserving the utility of the released data. Blocki *et al.* [1] presented a new mechanism for releasing perturbed password frequency list and the released password list is close to the original list. Deng and Mirkovic [4] proposed a mechanism that achieves commoner privacy-interactive k-anonymity. Commoner privacy fuzzes, by omitting or aggregating or adding noise, only those output points where individuals contribution is an outlier. They also discussed query composition and showed how they can guarantee privacy via pre-sampling step or query introspection. They implemented their privacy mechanism and query introspection on network traces using a system called Patrol. They compared the performance of their privacy preserving mechanism against differential privacy and crowd blending privacy. The results indicate that their proposed mechanism release outputs that have a higher research utility as compared to the two privacy preserving techniques. However, differential privacy guarantees high privacy than the other two techniques [4] and can protect against both all-but-one and interactive adversaries. The other two techniques can protect an individual from interactive adversary only. Several approaches to improve the utility of release aggregates using differential privacy exists. Therefore, releasing aggregates using differential privacy is still of benefit.

The works presented by [1, 4, 8, 12] indicate that differentially privacy can be adopted to preserve privacy in publishing network traces. In this work we are attempting to use differential privacy to release TCP SYN packets counts whereas [1, 8] released differentially private password list and number of sessions in the database browsing page i at time k respectively. Deng and Mirkovic [4] released differentially private, commoner private and crowd blending private packet counts sent per source port, packet counts received per destination service port, connection count in the trace and traffic volume in the trace.

3 Problem Statement

This section formally defines the problem of monitoring, using differential privacy, the new connections to the web server initiated by employees of an organisation that are browsing the web in a given working day (eight hours). Specifically, the number of TCP SYN packets sent to the webserver(s) during each 10s interval of a given working day resulting from the new connection request to the web server(s) by employees of an organisation that are browsing the web are released using differential privacy to protect the identity of web browsing employees from being inferred by an adversary from the original number of TCP SYN packets using possible background knowledge about the employees' web browsing patterns. That is, if the adversary knows the surfing behaviours of employees in an organisation releasing original HTTP associated TCP SYN packet counts can result to an adversary identifying the presence or absence of at least one employee in the organisation's database of HTTP associated TCP SYN packets. For an example, if the adversary knows that employee A surfs the net noticeably more (more HTTP associated TCP SYN packets generated for this employee) than the other employees and this employees surfs the net at a particular time interval during the day then the presence or the absence of that employee can be determined by the adversary since if employee A is present in the database the TCP SYN packet counts in that period will be noticeably higher than the TCP SYN packet counts in that period in a database that has the same records as the first database except that employee A has been removed. Therefore that noticeable difference in the TCP SYN counts in that period between the two databases has to be masked and differential privacy is capable of doing so. Furthermore, according to Yurcik *et al.* [18] TCP flags can be used to fingerprint different operating systems. Therefore releasing raw TCP SYN packets can expose the different operating systems of the machines in use. In this work, the TCP SYN packets that initiate new TCP connections between HTTP clients (web browsing employees) and the webserver(s) are monitored with differential privacy. Specifically, the number of TCP SYN packets sent to the webserver(s) during each 10s interval of a given working day resulting from the new connection request to the web server(s) by employees of an organisation that are browsing the web are released using differential privacy. The availability of such aggregated TCP SYN packet counts will assist the intrusion detection researchers in training their intrusion detection system in order to be able to detect attacks such as TCP SYN flooding attack. The goal of this work is to provide the number of TCP SYN packets sent during each 10s interval of a given working day without disclosing the presence or absence of a particular web browsing employee. Formally the problem statement is stated below as:

Private TCP SYN packet counts monitoring: Let x_t denote the number of TCP SYN packets sent to the web server at time interval t, $1 \leq t \leq T$ where T is the length of the monitoring period. For every time interval t, a private count s_t is to be released such that the released series $s_t, t = 1, ..., T$ is ε-differential private.

Furthermore, similarly to [8], we decided to have a limit on the number of webpage requests initiated by an individual employee to the webserver in the 8 hours, hence we set a limit on the number TCP SYN packets sent to the webserver(s) by an individual employee on a given 8 hour working day, since

1) An employee should not be browsing the web the whole 8 hours (except it is their job description, in which this work excludes those types of employees or organisations or cases),

2) Any web browser can only browse a limited number of webpages in a given 8 hours,

3) From a privacy point of view, if an employee requests an unlimited number of webpages in the 8 hours then large amount of noise will be required in order to account for such influence on the aggregate. The limit to the TCP SYN packets sent by an individual employee to the web server(s) on a given eight hour working day is denoted by C_{max} and we assume $C_{max} < T$.

4 Differential Privacy

In this work we aim to provide differentially private TCP SYN packet counts. A mechanism is said to be differentially private if its output is not significantly affected by the removal or addition of any record. Therefore at the release of the outcome, an adversary learns almost the same information about any individual record, regardless of its presence or absence in the original database.

Definition 1. *(ε-differential privacy [2]). A privacy mechanism A satisfies ε-differential privacy if for any dataset D_1 and D_2 differing on at most one record, and for any possible anonymised dataset $D \in Range(A)$,*

$$\Pr[A(D_1) = D] \leq e^{\varepsilon} \Pr[A(D_2) = D]. \qquad (1)$$

where the probability is taken over the randomness of A.

The privacy parameter ε, also called the privacy budget [13], specifies the degree of privacy offered. Intuitively, a lower value of ε implies stronger privacy guarantee and a larger perturbation noise, and a higher value of ε implies a weaker guarantee while possibly achieving higher accuracy. Two databases D_1 and D_2 that differ on at most one record are called neighbouring databases. In our problem definition, a database "record" represents a new connection request to the webserver, *i.e.* the record is associated with the sending of the TCP SYN packet to the webserver by the client (web browsing employee) and therefore our work is designed to protect the presence or absence of every web browsing employee.

Laplace Mechanism. Dwork *et al.* [7] show that ε-differential privacy can be achieved by adding independent and identically distributed noise to query result $q(D)$:

$$\begin{aligned} q(D) &= q(D) + (N_1, ..., N_m), \\ N_i &= Lap(0, \frac{GS(q)}{\varepsilon}) \text{ for } i = 1, ..., m. \end{aligned}$$

where m represents the dimension of $q(D)$. The magnitude of N conforms to a Laplace distribution with 0 mean and $GS(q)/\varepsilon$ scale, where $GS(q)$ represents the global sensitivity [7] of the query q.

Global sensitivity. The global sensitivity [7] is the maximum L1 distance between the results of q from any two neighbouring databases D_1 and D_2. Formally, it is defined as follows:

$$GS(q) = \max ||q(D_1) - q(D_2)||.$$

Composition. The composition properties of differential privacy provide privacy guarantees for a sequence of computations as outlined in theorem 1 below.

Theorem 1. *Sequential composition [13]. Let each A_i provide ε_i-differential privacy. A sequence of $A_i(D)$ over the dataset D provides $\sum_i \varepsilon_i$-differential privacy.*

Given Theorem 1, the Laplace perturbation is applied at every time series time stamp to guarantee (ε/T)-differential privacy, where T is the series length.

5 Differentially Private TCP SYN

In this section the application of differential privacy to the TCP SYN packet counts is outlined.

5.1 Privacy Mechanism

The Laplace Mechanism is suitable for numerical queries [19] and is adopted in this work as the privacy mechanism since we are monitoring a numerical aggregate statistic.

5.2 Global Sensitivity

In this section the global sensitivity for monitoring the TCP SYN packet counts per 10s interval in a given eight hour working day is analysed. Let D be the database that consists of employees HTTP requests to the web server in a given 8 hour working day, $q(D) = x_1, ..., x_T$ be the sequence of outputs from the count queries , where x_t denotes the number of TCP SYN packets sent during *t-th* 10s interval and T be the series length (number of 10s intervals in an 8 hour working day). To determine the global sensitivity $GS(q)$, we studied the HTTP related TCP SYN packets in the DARPA 1999 dataset and noticed that an individual can request more than one webpage in a given time interval t and can appear in more than one time intervals. This means more than one TCP SYN packets can originate from the same source in a given time interval t. The effect of this is that the removal or addition of an individual to database D would change the output by at least 1. As we have observed also that the individual can appear in more than one time interval, the global sensitivity of the count query will be affected since global sensitivity defines the maximum contribution of an individual to the function output [10]. From the DARPA 1999 dataset we found $C_{max} = 712$, where C_{max} value is the maximum HTTP related TCP SYN packets originating from the same source over the eight hours. We therefore set $GS(q) = 712$ since this is the highest maximum contribution by an individual in D.

5.3 Filtering

As we have mentioned in the introduction that direct application of Laplace mechanism to the original aggregates may lead to high perturbation error and leaving the released aggregates to be of no useful value, we adopted the filtering component of [9] in order to improve the accuracy (closeness to the original aggregates) of the released aggregates. Their filtering component utilizes time series modelling and estimation algorithm. In their context, filtering, refers to the derivation of the posterior estimates of the original time series from the noisy measurements with the hope of removing background noise from the signal. They estimated the original time series from the noisy measurements using a Kalman filter [11] based estimation algorithm and used a state space model to describe the underlying dynamics of a time series as well as how an observation is derived from a hidden state [9]. In this work we modelled the time series and noisy measurements and estimated the original series from the noisy estimates to obtain the posterior estimates referred to as Kalman count estimates as follows:

Time series modelling: For the TCP SYN packet count series *i.e.* $\{x_t, t = 1, ..., T\}$, we defined the following models; process model:

$$x_t = x_{t-1} + \omega_t, \text{ where } \omega_t \sim N(0, Q),$$

where ω_t denotes the process noise at time interval t,

which is assumed to be a white Gaussian noise with variance Q.
Similarly, the measurement model for the noisy observations that are obtained from the Laplace perturbation mechanism is:

$$z_t = x_t + \nu_t, \text{ where } \nu_t \sim Laplace(0, GS(q)/\varepsilon),$$

where ν_t is the measurement noise at time interval t. Fan and Xiong [9] have established that the posterior distribution cannot be analytically determined if the distribution of the measurement noise is not Gaussian and reported that it is sufficient to approximate the distribution of the measurement noise to a Gaussian distribution. Thus, the following Gaussian distribution was proposed:

$$\nu_t \sim N(0, R), withR \propto (GS(q))^2/\varepsilon^2. \qquad (2)$$

In this work, we adopted the same approximation in Equation (2).

Estimation algorithm. We adopted the estimation algorithm of [8] which is based on the Kalman filter and approximated Laplace noise with Gaussian noise as suggested by [9]. Kalman filter [11] is a recursive method that provides an efficient means to estimate the state of a linear Gaussian process, by minimizing the variance of the posterior error. It consists of two steps, namely, prediction and correction steps. In the prediction step the state is predicted with the dynamic model. In the correction step the state is corrected with the observation model such that the error covariance of the estimator is minimised. The prediction and correction algorithms adopted in this work can be found in [8].

Privacy guarantee. The estimation algorithm provides ε-differential privacy since by definition of Laplace mechanism and sensitivity analysis in section 4, the Laplace perturbed values $\{z_t, t = 1, ..., T\}$ satisfy ε-differential privacy and similarly to [8], neither the Prediction nor Correction interacts with the raw data so there is no extra privacy leakage incurred by those two procedures.

6 Experimental Work

This section presents the dataset, parameter values and utility evaluation methods used in this work. We also describe how counts perturbation and filtering were done.

6.1 Data Set

DARPA 1999 dataset was used in this study. We used attack free data taken on a Monday. TCP SYN packets associated with HTTP requests to seven webservers were collected between 08:00 to 16:00 *i.e.* TCP SYN packets collected over 8 hours. Seven servers were used in order to limit the number of times an individual (web browsing employee) appears in the dataset so that the restrictions set in Section 3 for individuals browsing the net in a given eight hour working day are met. The number of TCP SYN packets in 10 second intervals were determined.

6.2 Parameters

Parameter values are as follows: The experiments were conducted at the interval privacy budget of, $\varepsilon_t = 0.01$, *i.e.* for each 10s interval we used a Laplace mechanism that provides ε_t-differential privacy, since it provides the lowest overall privacy budget(that can be obtained by using Theorem 1) of the recommended privacy budgets (0.01 and 0.1) [6]. For the utility evaluation using the average relative error and utility loss metrics, interval privacy budgets, $\varepsilon_t = 0.01, 0.1$ and 1 were used for comparison purposes. Process noise, $Q = 10000$ was empirically determined as the value that yields better estimates of the original TCP SYN packet counts given the interval privacy budget.

- Measurement noise, $R = (GS(q))^2/\varepsilon_t^2$;
- Global sensitivity, $GS(q) = 712$.

6.3 Laplace Perturbation and Filtering

The number of TCP SYN packets in 10 second intervals were determined and the Laplace noise was added to each count in each interval. The Kalman filter based estimation algorithm was used to estimate the original counts from the Gaussian perturbed counts (estimates of the Laplace perturbed counts as suggested by [9]). The estimates of the original counts are the ones that are released instead of the noisy counts resulting from Laplace perturbation. Figures 1, 2 and 3 present the original counts, noisy counts resulting from Laplace perturbation and estimates of the original counts, referred to as Kalman count estimates for the first 500 10s intervals respectively.

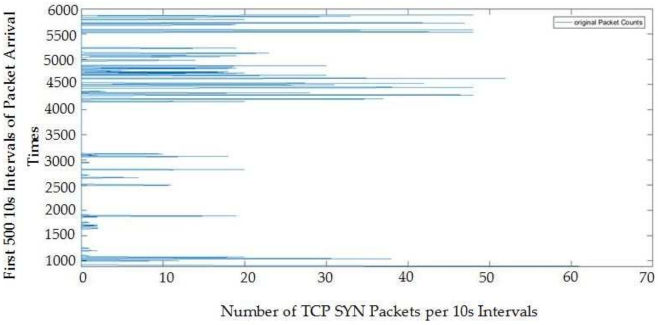

Figure 1: Original packet counts

6.4 Utility Evaluation

To measure the quality of released time series $s_t, t = 1, ..., T$:

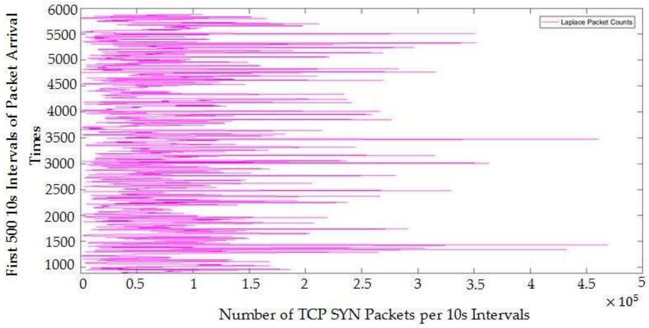

Figure 2: Laplace perturbed packet counts

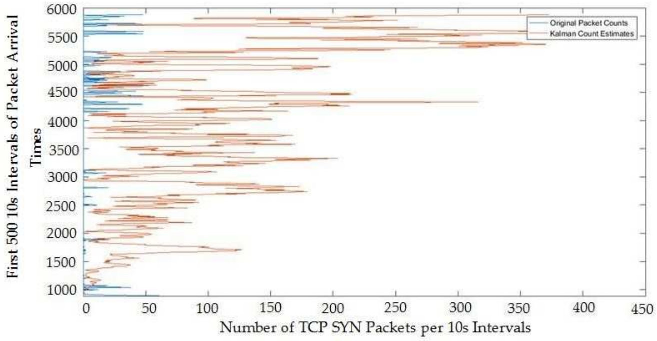

Figure 3: Original packet counts v.s. Kalman count estimates

- Two utility metrics called average relative error (E) and utility loss (U) were used,
- The performances of the Cumulative Sum (CUSUM) and Adaptive Threshold algorithms on the original aggregates were compared to the their performances on the released aggregates.

6.4.1 Average Relative Error

Average relative error (E) is a widely used metric to evaluate the accuracy of the data. It measures how well the released time series $s_t, t = 1, ..., T$ follows the original series $x_t, t = 1, ..., T$. It is defined as follows:

$$E = \frac{1}{T}\sum_{t=1}^{T}\frac{|s_t - x_t|}{\max\{x_t, \delta\}}$$

where $\delta = 1$ in order to handle cases where $x_t = 0$. Smaller values of E indicate high similarity between the released and the original series. We computed E values for the Laplace perturbed series and the Kalman count estimates corresponding to the three interval privacy budget values and are plotted in Figure 4. As indicated in Figure 4, the average relative errors for the Laplace perturbed counts were 67701, 6770 and 677 for $\varepsilon_t = 0.01, 0.1$ and 1 respectively while the Kalman count estimates resulted to average relative errors of 984, 433 and 135 for $\varepsilon_t = 0.01, 0.1$ and 1 respectively. These results indicate that the Kalman counts estimates which are the released counts are closer to the original counts.

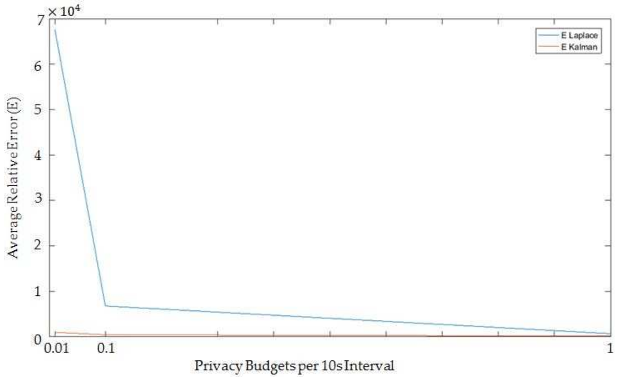

Figure 4: Average relative error comparison

6.4.2 Utility Loss

Utility loss is a relative cumulative difference between the true data points $x_t, t = 1, ..., T$ and the fuzzed data points $s_t, t = 1, ..., T$ [4]. It is defined as follows:

$$U = \frac{\sum_{i=1}^{N}|s_i - x_i|}{\sum_{i=1}^{N}|x_i|} \tag{3}$$

Small values of this measure indicate higher research utility [4]. We computed U values for the Laplace perturbed series and the Kalman count estimates corresponding to the three interval privacy budget values $\varepsilon_t = 0.01, 0.1$ and 1 and are plotted in Figure 5. As indicated in Figure 5, the utility loss values for the Laplace perturbed counts were 47262, 4725 and 472 for $\varepsilon_t = 0.01, 0.1$ and 1 respectively. The Kalman count estimates resulted to utility loss values of 679, 300 and 94 for $\varepsilon_t = 0.01, 0.1$ and 1 respectively. These results indicate that the Kalman count estimates have higher research utility than the Laplace perturbed counts.

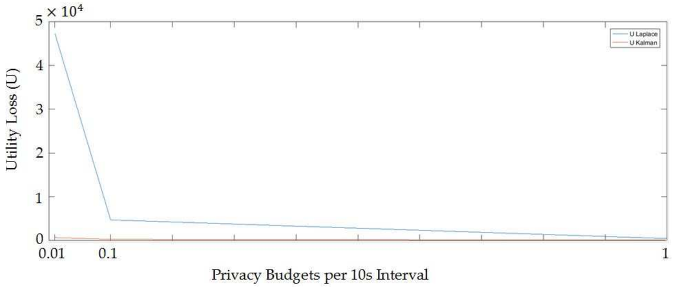

Figure 5: Utility loss comparison

6.4.3 CUSUM Algorithm

The CUSUM algorithm comes from the family of change point detection algorithms that are based on hypothesis testing and was developed for independent and identically distributed random variables. The detailed description of the CUSUM algorithm is not given in this work, it can be found in [17]. The CUSUM algorithm is used in this work to determine if inferences made using the released data are close to the ones made using the original data. Specifically, in this work the false positive rates obtained from CUSUM algorithm detection thresholds using the released data are compared to those obtained using the original data. Figure 6 presents these false positive rates. If we look at the overall pattern of the curves in Figure 6, the Kalman count estimates (which are the released differentially private counts) curve for $h \leq 8$ tend to follow the pattern of the original counts curve for $h \leq 6$ with a lag effect, which means inferences made using the Kalman count estimates for $h \leq 8$ will not be too far from the inferences made using the original counts for $h \leq 6$.

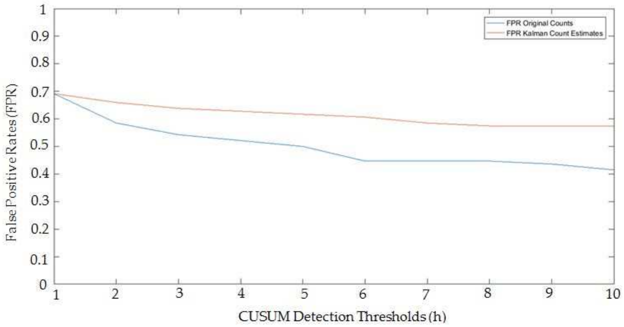

Figure 6: CUSUM false positive rates for the original counts vs Kalman estimates

6.4.4 Adaptive Threshold Algorithm

This algorithm tests whether the traffic measurement, number of Transmission Control Protocol (TCP) Synchronise (SYN) packets in a given time interval, exceeds a certain threshold. To address seasonality (daily and weekly variations) and trends, the threshold value is adaptively set from an estimate of the mean of the traffic measurements. A full description of this algorithm can be obtained in [17]. The Adaptive Threshold algorithm was similarly used as the CUSUM algorithm, the false positive rates obtained from the Adaptive Threshold algorithm detection thresholds using the released counts are compared to those obtained from the original counts. Figure 7 depicts these false positive rates. From Figure 7, the Kalman count estimates curve for $3 \leq k \leq 5$ tends to follow the pattern of the original counts curve for $3 \leq k \leq 4$ with a lag effect, which means inferences made using the Kalman count estimates for $3 \leq k \leq 5$ will not be too different from the inferences made using the original counts for $3 \leq k \leq 4$.

7 Discussion

The utility measure, average relative error at $\varepsilon_t = 0.01$, indicate that the Kalman count estimates are closer to the original counts as compared to the Laplace perturbed counts. The Utility loss measure at $\varepsilon_t = 0.01$ shows that the released counts have higher research utility as compared to the Laplace counts while preserving privacy.

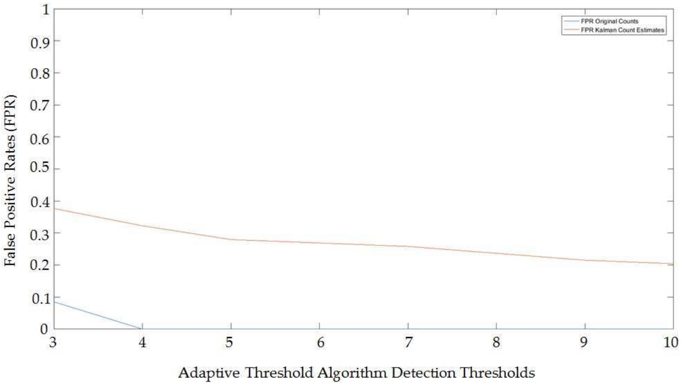

Figure 7: Adaptive threshold algorithm false positive rates for the original counts vs Kalman estimates.

Figures 6 and 7 also show that the false positive rates obtained from the CUSUM algorithm detection thresholds using the released counts are closer to the original counts as compared to those obtained from the Adaptive Threshold algorithm. Furthermore almost all the detection thresholds of the CUSUM algorithm ($h \leq 8$) lead to useful research inferences as compared to Adaptive Threshold algorithm thresholds with only $3 \leq k \leq 5$ thresholds leading to useful research inferences. Where useful research inferences means that inferences made using the released counts will be not that different from inferences made using the original counts. This means the released counts will work well for some algorithms and not work so well for others.

8 Conclusion

We proposed the use of differential privacy as a means of providing privacy to TCP SYN packets counts, adopted the filtering component of [9] in order to improve the accuracy of the released counts and test the utility of the released data by using two utility metrics and comparing the performance of two anomaly based intrusion detection algorithm on the original counts and the released counts. The results indicate that the inferences reached using the released counts are not that different from those reached using the original counts, with an added advantage of privacy.

Acknowledgments

I would like to thank my promoter Professor Fulufhelo Nelwamondo for his guidance and support and the CSIR: Modelling and Digital Science Unit for supporting my studies.

References

[1] J. Blocki, A. Datta, and J. Bonneau, "Differentially private password frequency lists," *IACR Cryptology ePrint Archive*, vol. 2016, pp. 153, 2016.

[2] A. Blum, K. Ligett, and A. Roth, "A learning theory approach to noninteractive database privacy," *Journal of the ACM (JACM'13)*, vol. 60, no. 2, pp. 12, 2013.

[3] S. E. Coull, C. V. Wright, F. Monrose, M. P. Collins, and M. K. Reiter, "Playing devil's advocate: Inferring sensitive information from anonymized network traces," in *NDSS*, vol. 7, pp. 35–47, 2007.

[4] X. Deng and J. Mirkovic, "Commoner privacy and a study on network traces," in *Proceedings of the 33rd Annual Computer Security Applications Conference*, pp. 566–576, 2017.

[5] N. V. Dijkhuizen and J. V. D. Ham, "A survey of network traffic anonymisation techniques and implementations," *ACM Computing Surveys (CSUR'18)*, vol. 51, no. 3, pp. 52, 2018.

[6] C. Dwork, "Differential privacy: A survey of results," in *International Conference on Theory and Applications of Models of Computation*, pp. 1–19, 2008.

[7] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *TCC*, vol. 3876, pp. 265–284, 2006.

[8] L. Fan, L. Bonomi, L. Xiong, and V. Sunderam, "Monitoring web browsing behavior with differential privacy," in *Proceedings of the 23rd International Conference on World Wide Web*, pp. 177–188, 2014.

[9] L. Fan and L. Xiong, "Real-time aggregate monitoring with differential privacy," in *Proceedings of the 21st ACM International Conference on Information and Knowledge Management*, pp. 2169–2173, 2012.

[10] L. Fan and L. Xiong, "Differentially private anomaly detection with a case study on epidemic outbreak detection," in *IEEE 13th International Conference on Data Mining Workshops (ICDMW'13)*, pp. 833–840, 2013.

[11] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.

[12] F. McSherry and R. Mahajan, "Differentially-private network trace analysis," *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 4, pp. 123–134, 2011.

[13] F. D. McSherry, "Privacy integrated queries: an extensible platform for privacy-preserving data analysis," in *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pp. 19–30, 2009.

[14] J. Mirkovic, "Privacy-safe network trace sharing via secure queries," in *Proceedings of the 1st ACM Workshop on Network Data Anonymization*, pp. 3–10, 2008.

[15] J. C. Mogul and M. Arlitt, "Sc2d: an alternative to trace anonymization," in *Proceedings of the SIGCOMM Workshop on Mining Network Data*, pp. 323–328, 2006.

[16] R. Paul, V. C. Valgenti, and M. S. Kim. "Obfuscating and anonymizing network traffic - A new dimension to network research," *School of Electrical Engineering and Computer Science*, 2010. (`https://research.libraries.wsu.edu/xmlui/bitstream/handle/2376/2655/Paul%2C%20R%20Obfuscating%20and%20anonymizing%20.pdf?sequence=1&isAllowed=y`)

[17] V. A. Siris and F. Papagalou, "Application of anomaly detection algorithms for detecting syn flooding attacks," *Computer Communications*, vol. 29, no. 9, pp. 1433–1442, 2006.

[18] W. Yurcik, C. Woolam, G. Hellings, L. Khan, and B. Thuraisingham, "Toward trusted sharing of network packet traces using anonymization: Single-field privacy/analysis tradeoffs," *Computer Science*, 2007. (`https://arxiv.org/abs/0710.3979`)

[19] T. Zhu, G. Li, W. Zhou, and S. Y. Philip, "Preliminary of differential privacy," in *Differential Privacy and Applications*, pp. 7–16, 2017.

Biography

Nenekazi N. P. Mkuzangwe holds a Bachelor of Science and a MSc in Mathematical Statistics, both from Rhodes University, in South Africa. She has published 3 research papers in reviewed conferences.

Fulufhelo Nelwamondo is an electrical engineer by training, and holds a Bachelor of Science and a PhD in Electrical Engineering, in the area of Computational Intelligence, both from the University of the Witwatersrand, in South Africa. Prof Nelwamondo is a registered Professional Engineer, and is the Executive Director at the Council for Scientific and Industrial Research Modelling and Digital Science, South Africa. He is a senior member of the IEEE, and a visiting professor of Electrical Engineering at the University of Johannesburg. He was a post-doctoral fellow at the Graduate School of Arts and Sciences, of Harvard University. Nelwamondo has successfully supervised a number of Masters and PhD degrees in electrical engineering, and continues to do so. He has published over 100 research papers in journals, reviewed conferences and book chapters.