

Linking Correlated Network Flows through Packet Timing: a Game-Theoretic Approach

Juan A. Elices
University of New Mexico
Email: jelices@ece.unm.edu

Fernando Pérez-González
Universidad de Vigo
Email: fperez@gts.uvigo.es

Abstract—Deciding that two network flows are essentially the same is an important problem in intrusion detection or in tracing anonymous connections. A stepping stone or an anonymity network may try to prevent flow correlation by delaying the packets, introducing chaff traffic, or even splitting the flow in several subflows.

We introduce a game-theoretic framework for this problem. The framework is used to derive the Nash equilibrium under two different adversary models: the first one, when the adversary is limited to delaying packets, and the second, when the adversary also adds dummy packets and removes packets from the flow. As the optimal decoder is not computationally feasible, we restrict the possible decoder to one that estimates and compensates the attack. Our analysis can be used for understanding the limits of flow correlation based on packet timings under an active attacker.

Index Terms—traffic analysis, game theory, flow watermark, network security

I. INTRODUCTION

Network attackers intentionally hide their identity to avoid prosecution. A broadly-used way of achieving this anonymity is relaying the traffic through a chain of compromised hosts called stepping stones [1]. Intrusion detection and tracing back an attack require deciding that two flows are correlated. Linking network flows can also be used to compromise low-latency anonymous networks, such as Tor.

Two approaches exist for finding correlated flows: passive analysis and active watermarks. They differ in whether the flow is modified or not. In general, an active watermark needs shorter sequences but at the expense of being detectable [2].

An adversary (AD), for instance a stepping stone or an anonymous network, may take countermeasures to prevent the correlation such as introducing delays to packets or adding dummy packets to the flow, or even more drastic measures, e.g. dividing the flow into different subflows each taking different paths. To the best of our knowledge, in previous works, only [3] and [4] consider an AD, and in both the AD is limited to delaying packets.

This paper studies the limits of traffic analysis, passive or active, in an adversarial environment. The most natural solution to avoid the loop of proposing an attack and creating an *ad-hoc* solution is to cast the problem into a game-theoretic framework and look for the optimum strategies that the players, traffic analyst (TA) and AD, should adopt. A game-theoretic framework similar to the proposed one has

been used in other contexts such as Information Hiding [5] or Source Identification [6].

The rest of the paper is organized as follows: in Section II we introduce the notation, together with some basic concepts of game theory. Section III presents a rigorous definition of the traffic analysis game. Section IV solves the problem under an AD model that only introduces delays. Section V deals with a stronger AD problem that also adds chaff traffic and divides the flow. Conclusions are presented in Section VI.

II. NOTATION AND BASIC CONCEPTS

We use the following notation. Random variables are denoted by capital letters (e.g., X), and their individual realizations by lower case letters (e.g., x). The domains over which random variables are defined are denoted by script letters (e.g., \mathcal{X}). Sequences of n random variables are denoted with X^n if they have random nature or by x^n if they are deterministic. X_i or x_i indicate the i th element of X^n or x^n , respectively. The probability distribution function (pdf) of a random variable X is denoted by $f_X(x)$, $x \in \mathcal{X}$. We use the same notation to refer to pdf of sequences, i.e. $f_{X^n}(x^n)$, $x^n \in \mathcal{X}^n$. When no confusion is possible, we drop the subscript in order to simplify the notation. We denote with Δ the difference operation of a sequence, i.e. $\Delta x^n = \{x_2 - x_1, \dots, x_n - x_{n-1}\}$ and with $\Delta x_i = x_{i+1} - x_i$ the i th element of this sequence.

A. Game Theory

Game theory is the mathematical study of interaction among independent, self-interested agents. Formally, a two-player game is defined as a quadruple $G(A_1, A_2, u_1, u_2)$, where $A_i = \{a_{i,1}, \dots, a_{i,n_i}\}$ are the actions available to the i th player, and $u_i : A_1 \times A_2 \mapsto \mathbb{R}$, $i = 1, 2$ is the utility function or payoff of the game for player i . An action profile is the double $a \in A_1 \times A_2$. We are interested in zero-sum games, where $u_1(a) + u_2(a) = 0, \forall a \in A_1 \times A_2$, which means that the gain (or loss) of utility of player 1 is exactly balanced by the losses (or gains) of the utility of player 2. In this case, we can simplify the game notation to a triplet $G(A_1, A_2, u)$, where $u = u_1 = -u_2$.

We say that an action profile $(a_{1,i^*}; a_{2,j^*})$ represents a Nash equilibrium (NE) if $u(a_{1,i^*}; a_{2,j^*}) \geq u(a_{1,i}; a_{2,j^*}) \quad \forall a_{1,i} \in A_1$ and $u(a_{1,i^*}; a_{2,j^*}) \leq u(a_{1,i^*}; a_{2,j}) \quad \forall a_{2,j} \in A_2$, intuitively, this means that none of the players can improve its

As the AD must decide its action in real time and (9) is computationally expensive, we approximate it by

$$A_{AD} \approx \arg \min_{a^n} \sum_{i=1}^{n-1} \log f_{\Delta Y}(\Delta x_i + \Delta a_i). \quad (10)$$

which is a good approximation when ΔD has zero-mean and its variance is much smaller than ΔY (as it is the case in practice). Note that under this approximation AD is maximizing the likelihood of $x^n + a^n$ coming from y^n , i.e., making the sequence as typical as possible.

A. Performance

In this section we construct a simulator and present the scenarios we use in the remaining of the paper. Afterwards, we compare the detector performance between the AD choosing its action optimally and an AD that selects its attack randomly.

1) *Simulator and Scenarios*: Simulations are carried out in the following way. First, we randomly generate one sequence, X^n . Then we generate 10^6 sequences Y^n and we calculate ϵ such that (7) holds. To deal with the probabilistic nature of the delay, we repeat the following 50 times: we apply D_1^n , next, we implement the adversary action according to (10). Subsequently, we introduce another delay D_2^n . Finally, we calculate the utility for this X^n using (6). We repeat the whole process for 1000 different X^n , the plots show the average utility for the 1000 sequences, \bar{u} . Recall that the utility function is the probability of detection. In order to show that AD is acting rationally we compare it with an AD that introduces delays randomly, i.e., a^n is an independent and identically distributed (i.i.d.) sequence uniformly distributed between 0 and A_{max} .

Note that (3) needs an estimation of $f_{\Delta D}$ and $f_{\Delta Y}$, to this end we apply kernel smoothing techniques [7]. As it is customary, we separate the data into two subsets: training, to obtain the pdfs, and test, used in the simulator, using 50% of the samples for each.

Scenario A represents a stepping stone that forwards SSH traffic inside the Amazon Web Services network. TA-1, AD and TA-2 (cf. Figure 1) are EC2 instances located in Virginia, Oregon and California, respectively. We use the IPDs from 8746 replayed SSH connection captures with 64 million packets from [8] and [9]. The simulated delays correspond to Scenario 10 from [10].

Scenario B simulates a web page accessed from Tor network whose real origin is to be found. In it, TA-1 corresponds to the web server, AD to the Tor entry relay, and TA-2 to the client. For instance, this case can correspond to a company in whose forum an anonymous insulting post has been placed using Tor and it is to be known whether the source comes from an employee within the company. We use the IPDs of 113690 replayed HTTP connections that sum around 139 million packets taken from the same repositories. The delays correspond to the measurements of Scenario 11 from [10].

Results are depicted in Figures 2a and 2b for Scenario A and B, respectively. We see that if AD chooses the delays

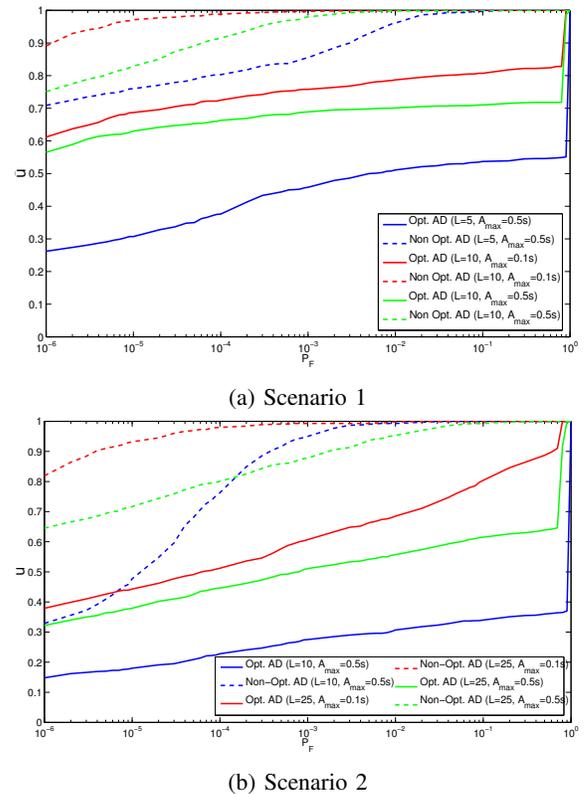


Fig. 2: Performance under a delaying AD model.

optimally, the impact of the attack is much larger than using random delays.

V. CHAFF TRAFFIC AND FLOW SPLITTING

Hitherto, we have assumed that there exists a one-to-one relation between the flows at the creator and detector; i.e., no packets are added or removed. In this section, we make a robust test to repacketization and study the game when, in addition to random delays, AD can also insert chaff traffic and divide the flow, and TA-2 observes only one subflow.

To deal with packet addition and removal, TA matches each packet from x^n with the most likely from w^m , $m \leq (1 + P_A)n$, denoting this as $i \rightarrow j$, as follows: $|x_i - (w_j - \rho)| < |x_i - (w_k - \rho)|$, $\forall k \neq j$, and to avoid considering it lost $|x_i - (w_j - \rho)| > \gamma$, where ρ is a synchronization constant obtained through an exhaustive search (auto-synchronization property), and γ is a threshold for which a packet is considered lost. To prevent that two packets i_1 and i_2 are matched to the same j packet, the set of possible matching candidates for the later packet is reduced to the non-matched ones. Also γ should be big enough so that the probability that a packet is considered lost when it actually is not is very small, for instance our simulation uses $\gamma = Q_{|\Delta D|,0.999} + A_{max}$ where $Q_{|\Delta D|,0.999}$ means the 0.999-quantile of the absolute value of ΔD . This procedure outputs two subsequences, w^{n_2} and x^{n_2} , where $n_2 \leq n$.

After this matching we apply a modification to Λ_{AD} in (3)

that takes into account the new sequence length as follows:

$$\Lambda_{AD}(w^{n_2}, x^{n_2}, \hat{a}^{n_2}) = (n - n_2) \log(P_L) + \sum_{i=1}^{n_2-1} \log(P_L) + (1 - P_L) \frac{f_{\Delta D}(\Delta w_i - \Delta x_i - \Delta \hat{a}_i)}{f_{\Delta Y}(\Delta w_i)}, \quad (11)$$

We propose a new AD model that apart from adding delays as previously, it can remove up to $P_L \cdot n$ packets and add at most $P_A \cdot n$ dummy packets. Formally, the available actions for the adversary are:

$$A_{AD} = \{a^n \times l^n \times c^{n_A} : 0 \leq a_i \leq A_{max}, l_i = \{0, 1\} \sum_{i=1}^n l_i \leq P_L, \frac{n_A}{n} \leq P_A, m = n + n_A - \sum_{i=1}^n l_i\}, \quad (12)$$

where l^n is a binary sequence that represents which packets are removed from the flow, and c^{n_A} is the sequence with the timing of the n_A introduced dummy packets.

A. Results

To evaluate the proposed robust algorithm, we modify the simulator using the matching process as follows: longer sequences, i.e. $Y^{n+[P_A n]}$, are needed to calculate ϵ . Afterwards, AD chooses its action such that (11) is minimized. We compare the optimal AD with two non-optimal ADs: a) NO1 AD chooses its attack randomly, i.e. a^n is an i.i.d. sequence uniformly distributed between 0 and A_{max} , l^n is a binary random sequence that contains $\lfloor nP_L \rfloor$ '1's, and c^{n_A} is an i.i.d. sequence uniformly distributed between the timing of the first packet and the timing of the last packet. The b) NO2 AD selects a^n to minimize (11) (optimally) but chooses l^n and c^{n_A} identically as the previous AD.

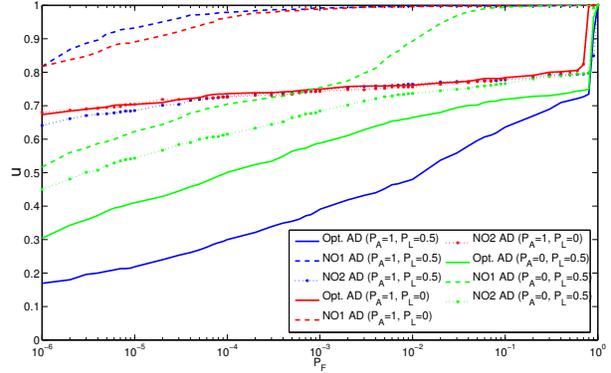
Results are depicted in Figures 3a and 3b. We see that the chaff traffic location when there are no losses does not make a big impact (the Opt-AD and NO-2 AD lines difference is very small). The rest of the data confirm that AD can harm the TA considerably more when deciding its actions logically than acting randomly.

VI. CONCLUSIONS

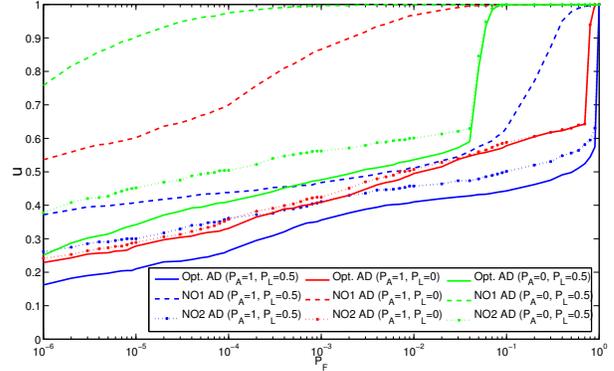
Using a game theoretic framework we have analyzed the flow linking problem. This problem consists in deciding if two flows are linked, having an adversary in the middle that tries to impair the correlation. Using this framework we obtain the optimal choices if both players act rationally. However, in a real implementation finding these optimal choices is not computationally feasible. This made us restrict the decoder to those that estimate the attack. We apply it to two different adversaries, one that only delays packets and another that can also add and drop packets from the flow. While in the present work the original flow is given, a future work will allow this flow to be modified through a watermark.

ACKNOWLEDGMENT

Research supported by Iberdrola Foundation through the Prince of Asturias Endowed Chair in Information Science and



(a) Scenario 1, $L=20$, $A_{max} = 500ms$



(b) Scenario 2, $L=30$, $A_{max} = 500ms$

Fig. 3: Performance under chaff packets and removed packets.

Related Technologies.

REFERENCES

- [1] S. Staniford-Chen and L. Heberlein, "Holding intruders accountable on the internet," in *Security and Privacy, 1995. Proceedings., 1995 IEEE Symposium on*, may 1995, pp. 39–49.
- [2] X. Luo, P. Zhou, J. Zhang, R. Perdisci, W. Lee, and R. K. C. Chang, "Exposing invisible timing-based traffic watermarks with BACKLIT," in *Proceedings of the 27th Annual Computer Security Applications Conference*. ACM, 2011, pp. 197–206.
- [3] D. L. Donoho, A. G. Flesia, U. Shankar, V. Paxson, J. Coit, and S. Staniford, "Multiscale stepping-stone detection: detecting pairs of jittered interactive streams by exploiting maximum tolerable delay," in *Proceedings of the 5th international conference on Recent advances in intrusion detection*, ser. RAID'02. Berlin, Heidelberg: Springer-Verlag, 2002, pp. 17–35.
- [4] A. Blum, D. Song, and S. Venkataraman, "Detection of interactive stepping stones: Algorithms and confidence bounds," in *Recent Advances in Intrusion Detection*, ser. Lecture Notes in Computer Science. Springer Berlin / Heidelberg, 2004, vol. 3224, pp. 258–277.
- [5] P. Moulin and J. O'Sullivan, "Information-theoretic analysis of information hiding," *Information Theory, IEEE Transactions on*, vol. 49, no. 3, pp. 563–593, 2003.
- [6] M. Barni and B. Tondi, "The source identification game: An information-theoretic perspective," *IEEE Transactions on Information Forensics and Security*, vol. 8, no. 3, pp. 450–463, 2013.
- [7] A. W. Bowman and A. Azzalini, *Applied Smoothing Techniques for Data Analysis: The Kernel Approach with S-Plus Illustrations (Oxford Statistical Science Series)*. Oxford University Press, USA, Nov. 1997.
- [8] D. Kotz, T. Henderson, I. Abyzov, and J. Yeo, "CRAWDAD trace set dartmouth/campus/tcpdump (v. 2004-11-09)," <http://crawdada.cs.dartmouth.edu/dartmouth/campus/tcpdump>, Nov. 2004.
- [9] R. R. R. Barbosa, R. Sadre, A. Pras, and R. van de Meent, "Simpleweb/university of twente traffic traces data repository," Centre for Telematics and Information Technology University of Twente, Enschede, Technical Report, 2010.
- [10] J. A. Elices and F. Pérez-González, "Measures to model delays on internet," <http://www.unm.edu/~elices/captures.html>, Jan. 2013.