Traffic Analysis

Scott E. Coull
RedJack, LLC.
Silver Spring, MD USA

Related Concepts and Keywords

Side-channel attack, information theory, cryptanalysis, covert channel analysis

Definition

Traffic analysis is the process of inferring information from the features of communication traffic other than its contents, such as message timing and size.

Background

The notion of traffic analysis originates from military signals intelligence where, as early as World War I, the dynamics of military communications were used to learn information about troop movements, the location of command headquarters, and the enemy chain of command. This information was inferred solely from the traffic patterns without access to the contents of the communication, which were often obfuscated or encrypted. In the context of computer security, traffic analysis typically refers to the examination of data sent over computer networks. As with military signals intelligence, traffic analysis of computer networks examine impact of hidden information on the observable features of the network data, such as packet sizes or their inter-arrival timing. By examining these features, an observer can make inferences on the hidden information based on this causal relationship. While obvious similarities exist between the field of cryptanalysis and traffic analysis, it is important to point out that cryptanalysis attempts to reveal the contents of an encrypted message by directly examining the ciphertext, whereas traffic analysis examines the secondary features of the network traffic and may reveal information beyond simply the contents of the traffic, including the identity of the sender or their location.

Theory

The theoretical underpinnings of traffic analysis are based on an information theoretic model similar to that of covert channel analysis. That is, the observed features of the network traffic can be considered to pass through an information channel from the source process producing the hidden information to the observer. The information gained through traffic analysis is bounded by the capacity of this channel [11]. More formally, let X be a random variable representing the distribution of messages produced by the source process, and Y be a random variable that represents the distribution of values for a given feature of the network traffic. In essence, the relationship between the distributions X and Y quantifies the extent to which messages affect the feature values of the network traffic before security mechanisms are applied. As a concrete example, one might consider X to be the distribution of web pages visited by a user, and Y might be the packet sizes produced when downloading those web pages before the packet has been encrypted. These two random variables can be considered to be the input and output, respectively, for a noisy channel that determines the extent to which knowing the values of Y tells you information about the values of X, which is labeled as the *Feature Channel* in Figure 1.

Of course, when performing traffic analysis, the network traffic features represented by Y may change due to network conditions or security mechanisms, such as padding the packet contents to alter sizes. As such, let Z define a third random variable that represents the distribution of feature values actually observed after these alterations are made. The changes made to the feature values are formalized by considering another noisy information channel with input distribution Y and output distribution Z, labeled as the *Transmission Channel* in Figure 1. The bound on the information that can be learned through traffic analysis, therefore, is given by the channel capacity of this entire system, which is calculated as the mutual information of Y and Z conditioned on X. The mutual information provides a notion of the correlation between two distributions, and is calculated in terms of Shannon entropy as:

$$I(Z;Y|X) = H(Z|X) - H(Z|Y,X) \tag{1}$$

where $H(Z|X)$ is the Shannon information of Z conditioned on X and $H(Z|Y,X)$ is the Shannon entropy of Z conditioned on the joint distribution of Y and X [11].

Intuitively, the information theoretic model states that in order for traffic analysis to work, there must be some causal effect where the unaltered feature values Y depend in some way on the hidden messages X. Moreover, there must also be some relationship between the observed features Z and the unaltered features caused by the hidden messages. Given this requirement, the most effective strategy

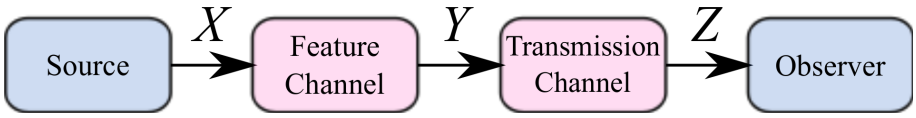

Figure 1: An information theoretic model of traffic analysis.

for preventing traffic analysis is to destroy the correlation among X, Y, and Z by ensuring the distributions are independent (e.g., via randomization), or by quantizing the values in the observable feature distribution. An excellent introduction for those wishing to gain a broader understanding of the use of information theory in defining information leakage can be found in [4].

Applications

A variety of hidden information can be inferred through the use of traffic analysis methods in practice. Typically, this hidden information is related in some way to the application-layer contents of encrypted traffic. For these cases, the features of packet size and inter-arrival time are most widely used due to their strong connection to the contents of the traffic. However, there have been some instances in which network-layer information is inferred, such as TCP timestamp information to identify computers [7] or the IPID field to reveal network topology information [2]. The choice of feature(s) used in traffic analysis depends heavily on the type of information to be inferred and the underlying protocol being analyzed. As an example, Voice over IP (VoIP) protocols typically send packets at fixed intervals to ensure call quality, and so timing information is unlikely to yield interesting information about the traffic. On the other hand, the size of SSL packets is often closely related to the size of their plaintext contents, indicating that it would be a useful feature for gaining information about those contents.

To defeat network traffic analysis methods, security protocols often add padding to packets or send packets at fixed intervals. Doing so ensures that the correlation between the observed features and the hidden information is reduced or completely removed. Unfortunately, using these quantization methods often incurs significant overhead in terms of the number of bytes sent or delay in communication, respectively. In order to balance the performance and security needs of network protocols, designers often select several levels of quantization that reduce the information available to the observer while still limiting overhead. It is important to note, however, that this balance can often result in protocols that leak more information about the hidden information than the designer intended.

In practice, traffic analysis can be both a useful security tool for network ad-

ministrators who want to discover potentially malicious activities, and a significant threat to the privacy of network users who simply want to protect their sensitive information. For instance, traffic analysis of packet timing information has been used to trace computer attacks through intermediate "stepping stone" to their origin [20, 5, 15]. However, a similar analysis could also be used to help infer the password of users logging into remote computers via SSH [13]. Other uses of traffic analysis include inferring web pages downloaded over SSL connections [14, 8, 3], revealing the contents of encrypted Voice over IP calls [17, 16], undermining anonymous communications systems [1, 12, 10], and performing network monitoring on encrypted traffic [6, 9, 19]

Open Problems and Future Directions

One exciting new direction of research in the field of traffic analysis lies in the development of mitigation strategies for traffic analysis that, rather than quantizing features, attempts to alter the features to make one type of hidden information look like another with respect to the observable features [18]. Rather than reducing the correlation between observable features and hidden information, these methods attempt to fool the observer into making incorrect inferences based on incorrect correlations. These so-called mimicry, or morphing, methods are helping to shape the future of traffic analysis in two ways. First, the amount of overhead incurred by these techniques is much less than that of quantization, since network data need only be altered to look like other real traffic. This opens the possibility for very efficient network protocols that do not leak information. Second, these methods break the information theoretic model of traffic analysis by purposefully correlating the observable features with the wrong hidden information. This motivates the development of new analytic models that capture not just the correlation, but also the correctness of that correlation.

Recommended Reading

[1] K. Bauer, D. McCoy, D. Grunwald, T. Kohno, and D. Sicker. Low-Resource Routing Attacks Against Tor. In *Proceedings of the 2007 ACM Workshop on Privacy in Electronic Society*, pages 11–20, October 2007.

[2] S. Bellovin. A Technique for Counting NATed Hosts. In *Proceedings of the* 2^{nd} *Annual ACM SIGCOMM Workshop on Internet Measurement*, pages 267–272, November 2002.

[3] S. Coull, M. Collins, C. Wright, F. Monrose, and M. K. Reiter. On Web Browsing Privacy in Anonymized NetFlows. In *Proceedings of the* 16^{th} *Annual USENIX Security Symposium*, pages 339–352, August 2007.

[4] T. Cover, J. Thomas, and M. Burns. *Elements of Information Theory, Vol. 1, (revised edition)*. Wiley Series in Telecommunications and Signal Processing, John Wiley & Sons, Inc., 2006.

[5] D. Donoho, A. 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* 5^{th} *Annual International Symposium on Recent Advances in Intrusion Detection*, pages 17–35, October 2002.

[6] T. Karagiannis, K. Papagiannaki, and M. Faloutsos. BLINC: Multilevel Traffic Classification in the Dark. In *Proceedings of the Annual ACM SIGCOMM Conference*, pages 229–240, August 2005.

[7] T. Kohno, A. Broido, and K. Claffy. Remote Physical Device Fingerprinting. In *Proceedings of the* 26^{th} *Annual IEEE Symposium on Security and Privacy*, pages 93–108, May 2005.

[8] M. Liberatore and B. Levine. Inferring the Source of Encrypted HTTP Connections. In *Proceedings of the* 13^{th} *ACM Conference on Computer and Communications Security*, pages 255–263, October 2006.

[9] P. McDaniel, S. Sen, O. Spatscheck, J. Van der Merwe, W. Aiello, and C. Kalmanek. Enterprise Security: A Community of Interest Based Approach. In *Proceedings of the* 13^{th} *Annual Network and Distributed Systems Security Symposium*, February 2006.

[10] S. J. Murdoch and G. Danezis. Low-Cost Traffic Analysis of Tor. In *Proceedings of the* 26^{th} *Annual IEEE Symposium on Security and Privacy*, pages 183–195, May 2005.

[11] C. E. Shannon. A Mathematical Theory of Communication. *Bell System Technical Journal*, 27:379–423, 623–656, July/October 1948.

[12] V. Shmatikov and M. H. Wang. Timing Analysis in Low-Latency Mix Networks: Attacks and Defenses. In *Proceedings of the* 11^{th} *Annual European Symposium on Research in Computer Security*, pages 18–33, September 2006.

[13] D. Song, D. Wagner, and X. Tian. Timing analysis of keystrokes and SSH timing attacks. In *Proceedings of the* 10^{th} *Annual USENIX Security Symposium*, August 2001.

[14] Q. Sun, D. R. Simon, Y. Wang, W. Russell, V. N. Padmanabhan, and L. Qiu. Statistical Identification of Encrypted Web Browsing Traffic. In *Proceedings of the* 23^{rd} *Annual IEEE Symposium on Security and Privacy*, pages 19–31, May 2002.

[15] X. Wang, D. Reeves, and S. F. Wu. Inter-Packet Delay Based Correlation for Tracing Encrypted Connections Through Stepping Stones. In *Proceedings of the* 7^{th} *Annual European Symposium on Research in Computer Security*, pages 244–263, October 2002.

[16] C. Wright, L. Ballard, S. Coull, F. Monrose, and G. Masson. Spot Me if You Can: Uncovering Spoken Phrases in Encrypted VoIP Conversations. In *Proceedings of the* 29^{th} *Annual IEEE Symposium on Security and Privacy*, pages 35–49, May 2008.

[17] C. Wright, L. Ballard, F. Monrose, and G. Masson. Language Identification of Encrypted VoIP Traffic: Alejandra y Roberto or Alice and Bob? In *Proceedings of the* 16^{th} *Annual USENIX Security Symposium*, pages 43–54, August 2008.

[18] C. Wright, S. Coull, and F. Monrose. Traffic Morphing: An Efficient Defense Against Statistical Traffic Analysis. In *Proceedings of the* 16^{th} *Annual Network and Distributed Systems Security Symposium*, pages 237–250, February 2009.

[19] K. Xu, Z. Zhang, and S. Bhattacharyya. Profiling Internet Backbone Traffic: Behavior Models and Applications. In *Proceedings of the Annual ACM SIGCOMM Conference*, pages 169–180, August 2005.

[20] Y. Zhang and V. Paxson. Detecting stepping stones. In *Proceedings of the* 9^{th} *Annual USENIX Security Symposium*, pages 171–184, August 2000.