

Introducing Traffic Analysis

George Danezis and Richard Clayton

January 21, 2007

1 Introduction

In the Second World War, traffic analysis was used by the British at Bletchley Park to assess the size of Germany's air-force, and Japanese traffic analysis countermeasures contributed to the surprise of their 1941 attack on Pearl Harbour. Nowadays, Google uses the incidence of links to assess the relative importance of web pages, credit card companies examine transactions to spot fraudulent patterns of spending, and amateur plane-spotters revealed the CIA's 'extraordinary rendition' programme. Diffie and Landau, in their book on wiretapping, even went so far as to say that "traffic analysis, not cryptanalysis, is the backbone of communications intelligence" [20]. However, until recently the topic has been neglected by computer science academics. A rich literature discusses how to secure the confidentiality, integrity and availability of communication content, but very little work has considered the information leaked from communications 'traffic data' and how these compromises might be minimised.

Traffic data records the time and duration of a communication, and traffic analysis examines this data to determine the detailed shape of the communication streams, the identities of the parties communicating, and what can be established about their locations. The data may even be sketchy or incomplete – simply knowing what 'typical' communication patterns looks like can be used to infer information about a particular observed communication.

Civilian infrastructures, on which state and economic actors are increasingly reliant, are ever more vulnerable to traffic analysis: wireless and GSM telephony are replacing traditional systems, routing is transparent and protocols are overlaid over others – giving plenty of opportunity to observe, and take advantage of the traffic data. Concretely, an attacker can make use of this

information to gather strategic intelligence, or to penetrate particular security protocols and thus violate traditional security properties.

In this short introduction to the topic, we will highlight the key issues around traffic analysis. We start with its military roots and present the defences the military have developed for protection. We then consider how traffic analysis is being used in modern civilian contexts. We then consider specific ‘computer science’ issues, and provide an overview of the relevant research literature on attacks and defences in contemporary networks. Finally we discuss some contemporary policy issues relating to the retention of traffic data.

2 Military roots

Traffic analysis is a key component of signal intelligence and electronic warfare. In his book ‘Intelligence Power in Peace and War’ [35] Michael Hermann, who has served as chair of the UK Joint Intelligence Committee, describes the value of extracting data from non-textual sources (to be understood as ‘data that is not the message content’):

These non-textual techniques can establish targets’ locations, order-of-battle and movement. Even when messages are not being deciphered, traffic analysis of the target’s C3I system and its patterns of behaviour provides indications of his intentions and states of mind, in rather the same way as a neurologist develops insights about a silent patient by studying EEG traces from the brain.

Traffic analysis was used by the military even before the invention of wireless communications, but it was the broadcast nature of wireless, permitting anyone to listen in, that transformed its usefulness. The first naval action of the First World War, on 5th August 1914, was the cutting of Germany’s trans-Atlantic cables by the British cable ship *Telconia* [38], so that wireless telegraphy would have to be used instead of hard-to-intercept cable communications. Traffic analysis became an extremely potent source of intelligence as wireless communication became more widespread, particularly in naval and air operations. Ships at sea had to balance the value of communicating against the threat of being detected via direction finding if they transmitted. When transmitting, strict standards, governing call-signs and communication procedures, had to be adhered to in order

to minimize the information that traffic analysis could provide.

Another example of traffic analysis providing valuable intelligence (by Herman [35]) is the British reconstruction in 1941 of the structure of the network structure of the German Air Force radio network. This confirmed that a unit was composed of nine and not twelve planes, which led to a more accurate estimate of total strength. Identification of radio equipment was also be used for accurate detection of redeployments: each transmitter can be ‘fingerprinted’ by characteristics such as unintentional frequency modulations, the shape of the transmitter turn-on signal transient, the precise centre of frequency modulation and so on. These fingerprints can be used to track the device even though the messages it is transmitting are in an unbreakable code. Similar techniques can be used today to identify GSM phones [4]. In World War Two radio operators became skilled at recognizing the ‘fist’ of other operators, i.e. the characteristic way in which they typed their Morse code. Indeed, prior to Pearl Harbour, the Japanese transferred their aircraft carrier radio operators ashore, in order to persuade any eavesdropping Americans that the Japanese fleet was still in port. Even in more modern times, as the Desert Storm campaign began in 1991, the British ‘Operation Rhino’ replayed radio traffic from an exercise a few weeks earlier, and thereby misled the Iraqi forces as to where they were attacking.

Intelligence does not necessarily come from radio communications. The recording of aircraft identification numbers, by amateur plane-spotting enthusiasts the world over, permitted the reconstruction of recent CIA activities, and helped to prove the existence of their ‘extraordinary rendition’ programme, which transferred terrorist suspects to third countries, for imprisonment and interrogation [57].

It might be wondered why traffic analysis is so valuable to the military? The technique, although impressive in what it can determine, necessarily provides lower quality information compared with cryptanalysis and recovery of message content. However, it is both easier and cheaper to extract and process traffic data than content. It is easier because ciphers need considerable effort to break (when they break at all). It is cheaper because traffic data can be automatically collected and processed to provide high level intelligence. Computers can collect traffic data and map out structures and locations, whilst a skilled human operator is needed to listen to every radio transmission (often in a foreign language) in order to extract intelligence. For these reasons traffic analysis is often used to perform ‘target selection’ for further intelligence gathering (such as more intensive and

expensive surveillance), jamming or destruction. Given the enormous amount of communication and information on public networks we can expect these ‘economics of surveillance’ to be ever more relevant and applicable.

Signals Intelligence (or Sigint) the military term for techniques that include traffic analysis, is an arms race, and many ‘low probability of intercept and position fix’ communication methods have been devised by the military to minimize exposure to traffic analysis and jamming (a key reference here is Anderson [4]). Their principles of operation are simple: scanning many frequencies can only be done at some maximal rate and a great deal of power is necessary to jam a wide part of the frequency spectrum. Therefore, the first technique used to evade interception, and foil jamming was ‘frequency hopping’, now used in commercial GSM communications to improve reliability in the face of environmental noise. The basic technique is for Alice and Bob to share a key that determines, for each given time period, the frequency at which they will transmit. Eve, on the other hand, does not know the key and has to observe or jam the entirety of the frequency spectrum that may be used. In practice hopping is cheap and easy to implement, it makes it difficult to jam the signal (given that the hop frequency is high enough), but is poor at hiding the fact that communication is taking place. It is mainly used for tactical battlefield communications, where the adversary is unlikely to have very large jammers to hand.

A second technique is called ‘direct sequence spread spectrum’ (DSSS). This transforms a high-power low-bandwidth signal into a high-bandwidth low-power signal, using a key that is shared between Alice and Bob. It is easy for them to extract the transmitted signal, using their key, but an adversary will have to try to extract the signal from the noise, a difficult task given its low power (that will ideally be under the noise floor). DSSS has also inspired commercial communication systems and is now used in ADSL and cable modems as CDMA. Its most significant implementation problem is synchronization, and the availability of a reference signal (such as GPS) is of great help when implementing a practical system.

The final technique in the arsenal against interception is ‘burst communication’. The key idea is to communicate in a very short burst, so as to minimize the probability the adversary is monitoring the particular frequency being used at the relevant time. A cute variant of this is meteor scatter communications, using the ionization trail of small meteorites hitting the atmosphere to bounce transmissions between special forces troops in the field and a base station. Meteor scatter can

also be used in civilian applications when low bandwidth, high latency, but very low cost and high availability communications are required.

3 Civilian traffic analysis

Contemporary sociology models groups of individuals, not as a mass or a fluid, but in terms of their positions within a ‘social network’. The controversial basis for much of this research is that the position of an agent in the social network is in many ways more characteristic of them than any of their individual attributes. This position determines their status, but also their capacity to mobilize social resources and act (social capital). This position can also be determined via traffic analysis, yielding a map of the social network, and the position of each actor within it.

Social Network Analysis [75], and experimental studies, have recently gained popularity and led to interesting results, that are of use to traffic analysis, but also to network engineering more generally. It was first noted by Milgram [68] that typical social networks present a ‘small world’ property, in that they have a low diameter (experimentally determined to be about six hops between any two members) and are efficiently navigable. In other words there are short paths (i.e. intermediaries) between you and anyone else in the world, and you can find them efficiently: think of using hints from location and profession. This work has been used to build efficient peer-to-peer networks, but remains underused in security and trust analysis. Another key finding is that ‘weak links’ – people you do not know all that well – are instrumental in helping you with activities that are not commonplace but still very important. A well studied example is finding a job, where people using ‘far links’ are on average more successful, than those who limit themselves to their local contacts.

The first mathematical studies [62] of social networks (or ‘power law networks’ as they are often described because of the degree distribution of their edges) tell us a lot about their resilience to failure. It turns out that they are extremely resistant to random node failures, meaning that they stay connected and maintain a low diameter even when many random nodes have been removed. On the other hand such networks are very sensitive to the targeted removal of the nodes with high degree. After a few nodes have been removed the network will become disconnected, and well before that the diameter increases substantially. An equally effective attack is for an adversary

to remove nodes according to their ‘between-ness’, i.e. how many other nodes in the network they are connected to. Traffic analysis can be used to select the appropriate targets to maximize communication degradation and disruption.

Carley et al [11] proposed using network tools to disrupt networks of terrorists, and addressed the issues raised when multiple agents were involved, so that removing a single ‘leader’ would be effective. More recent research by Nagaraja and Anderson [52] tries to find strategies for a peer-to-peer network of nodes to resist node deletion attacks. The intuition behind these defensive strategies is that nodes connect to other random nodes in order to get resilience, while connecting according to a power law strategy to get efficient routing. When under attack the network regenerates links to maximize fault tolerance. When things are calm it reconfigures itself to be efficient. Garfinkel [29] considers the ‘Leaderless Resistance’ model of self-organising independent cells without any central control. He notes that it is “a desperate strategy employed by movements that do not have broad popular support and that fear infiltrators” and makes a number of policy suggestions for combating it.

Social network analysis is now starting to be used for criminal intelligence [66, 42]. Investigators try to map out criminal organisations by the use of traffic analysis techniques on telephone or network traffic and location data. This can be used to select targets for more intensive surveillance, and also to select appropriate targets for arrest and prosecution. Often these arrests are aiming to maximally disrupt the organization targeted. It is not always appropriate to arrest the most central or the most well-connected member – this would merely serve as a promotion opportunity for smaller crooks to take up the position. It was found to be more effective to arrest the ‘specialists’, i.e. those people in the organization that have a unique position or skills, that others would find difficult to fill. Examples include those who can forge papers, or crooked customs officials.

Similar techniques were used by the US military to locate Saddam Hussein in 2003. Tribal and family linkages were used to identify particular individuals with close ties to him, and these were selected for closer surveillance [37]. The latest (December 2006) US Army Counterinsurgency Manual now specifically deals with social network analysis, and discusses the Saddam Hussein operation as an example [69]. The ties between the 9/11 conspirators have also been mapped, and these connections clearly pick out Mohamed Atta as the central figure [44]. Additionally, Dombrowski et al [22] show how it is possible to predict the shape of a social network only some

of whose members and links are known.

Moving away from social links, the German police searched for Baader-Meinhof safe-houses by analysing gas and electricity records, looking for rented apartments with spikes in fuel usage where the bills were paid by transfers from banks in different parts of the country. Thirty-five years later, the UK police search for cannabis farms (where the plants are kept warm in artificial sunlight) by looking for unusually heavy usage of electricity – or, if the meter has been overridden, a mismatch between the power consumed in a locality and that which is billed for: an infra-red scan from a helicopter will then locate the house which is warmer than its neighbours. In more academic work, Fawcett and Provost [24] show how data mining techniques can be used to detect cellular phone fraud, with their automated approach proving better than hand-crafted detection rules.

Traffic analysis inspired techniques can also be used to protect systems and build trust. Ad-vogato [46] is a social network based system, that provides a community for free software developers. The fact that they introduce each other allows the system to establish whether an author is likely to be a spammer, and filter their messages out. Gibson et al [30] observed that the apparently anarchic structure of web page links could be seen as comprising many communities with central ‘authoritative’ pages linked by ‘hub pages’. Google’s PageRank [56] uses techniques that are very similar to web-page and social network profiling – in that it considers pages that are more central in the network (with more links pointing to them) as more authoritative. Techniques have also been devised [41] to automatically detect and extract web communities. Their results can be used both to assist and to attack users.

In a different milieu, Renesys Corporation monitors the Internet’s global routing table and analyses the BGP protocol traffic sent by service providers as they announce which blocks of IP addresses they will carry traffic for. Analysis of this data permits Renesys to generate ‘market intelligence’ indicating when major ISP customers are starting to move to new providers, when ISP market share is changing, or the impact of mergers or acquisitions on customer numbers [14].

4 Contemporary Computer and Communications Security

The Internet is a not battlefield in the usual sense, yet there is considerable potential for conflict within it. We shall see how traffic analysis techniques can be used to attack secured systems,

extract potentially useful information, and be used to censor (the equivalent of jamming) or abuse and spam third parties (the equivalent of deception). We shall also outline the key defence strategies one can use on the Internet to foil these attacks – and draw the parallels and also the differences with the military world.

The key difference to keep in mind when studying civilian traffic analysis research is that the attackers have reduced means. It is not military powers, with large budgets and the ability to intercept most communications that worry us, but commercial entities, local governments, law enforcement, criminal organizations and terrorist networks that have become the adversary. Therefore research has focused on attacks and solutions that can be deployed at low cost, and provide tangible tactical benefits (a pass phrase, a record of web accesses, ...). Beyond this, more strategic work is beginning to be done on the ways in which Internet traffic analysis can be of use to law enforcement, along with practical approaches for ensuring that routine surveillance can be evaded.

So what can we do if we are not allowed to look at the plaintext contents?

4.1 The Traffic Analysis of SSH

The secure shell protocol (SSH) permits users to log in remote terminals in a secure fashion. It does this by performing authentication using a pass-phrase and a public keyring, and subsequently encrypts all information transmitted or received, guaranteeing its confidentiality and integrity. One would think that any subsequent password entry (that might be required to log in to further remote services), over an SSH connection, should be safe. Song et al. [65] show that there is a lot of information still leaking. In interactive mode SSH transmits every key stroke as a packet and hence the password length is trivially available.

However, because keyboard layouts are not random, and passwords are often based upon real words, the exact timing of the keystrokes is related to how quickly one particular character can be typed after another. Hence more advanced techniques, using hidden Markov models, can be used to extract further information from inter-packet timing and lower the effective entropy of the passwords, thereby making brute force guessing far easier.

It turns out that you do not need to measure the typing abilities of the person entering the password and another user can be used to build a profile, because the similarities between users are exploitable. This links in with subtly different results from Monroe and Rubin's [49] research

on identifying and authenticating users using keystroke dynamics. Although their focus was on biometrics and authentication their results have a clear relevance to the traffic analysis of SSH. They show that there can be enough variability in typing patterns between users to be able to identify them, particularly after a long sequence has been observed. As a result not only the content of your communications may be leaked but also your identity – despite all of confidentiality that SSH apparently affords.

4.2 The Traffic Analysis of SSL

The Secure Socket Layer (SSL, and its close friend Transport Layer Security (TLS)) was introduced primarily to provide private web access. HTTP protocol requests and replies are encrypted and authenticated between clients and servers, to prevent information from leaking. Yet there is plenty of research [15, 9, 67, 13, 36] to suggest that information is leaking out of this shell.

The key weaknesses come down to the shape of traffic that is inadequately padded and concealed. Browsers request resources, often HTML pages, that are also associated with additional resources (images, stylesheets, ...). These are downloaded through an encrypted link, yet their size is apparent to an observer, and can be used to infer which pages are accessed (for example, it would be possible to tell which specific company reports were being downloaded by an investment banker – with consequent possibilities for profitable stock trading). There are many variants of this attack: some attempt to build a profile of the web-site pages and guess from that which pages are being accessed while others use these techniques to overcome naive anonymizing SSL proxies. In the latter case the attacker has access to the cleartext input streams and he tries to match them to encrypted connections made to the proxy.

It should be noted that latent structure and contextual knowledge are of great use when extracting information from traffic analysis. Levene and Loizou [45] provided a theoretical basis for computing the entropy of web navigation and demonstrated that this ‘surfing’ should not be seen as just random. Danezis [15] assumed that users will usually follow links between different web resources. By learning just the approximate lengths of the resources that were accessed, he showed that a hidden Markov model can be used to trace the most likely browsing paths a user may have taken,. This approach provides much faster and more reliable results than considering users that browse at random, or web-sites that have no structure at all.

4.3 Web Privacy

Can a remote web server that you are accessing tell if you have also been browsing another site? If you were looking at a competitor's site then maybe giving you a better price might be in order!

Felten et al. [25] show that it is possible to use the caching features of modern web browsers to infer information about the web-sites that they have been previously browsing. The key intuition is that recently accessed resources are cached, and therefore will load much more quickly than if they had to be downloaded from the remote site. Therefore by embedding some foreign resources into a served page, the attacker's web-server can perform some timing measurements, and infer particular previous browsing activity.

Note that this attack can be performed even if the communication medium is anonymous and unlinkable. Most anonymization techniques work at the network layer, making it difficult to observe network identities, but perform only minimal filtering in higher layers. The presence of caches leads to the Felten attack, but doing away with any caching would be a major problem for anonymous communication designers since it is important to use any efficiency improvements possible to make the, already slow, browsing more usable.

4.4 Network Device Identification and Mapping

Can you tell if two different addresses on the Internet are in fact the same physical computer? Kohno et al. at CAIDA [43] have devised a technique that allows an attacker to determine if two apparently different machines are the same device. They note that the clock skew (the amount by which the clock drifts per unit of time) is characteristic of the hardware, and the physical conditions in which the crystal is maintained (heat, light, etc). Therefore, if the clock drift of the remote machines seems to match for a long time, it is very likely that the machine is in fact the same. The technique they use is resistant to latency, and can be applied remotely even if the target machine synchronises its clock with NTP or SNMP, and even if it is a web server that echoes the current time at a one second granularity.

The technique can be used in forensics to link visiting machine identities together, and to determine if two web-sites are hosted on the same consolidated server. Equally, it can be used by hackers to detect if the multiple machines they are accessing are merely different versions of a

virtualized honey-pot machine.

Murdoch [50] has extended this work by observing that the clock skew will change as the temperature changes. He has shown that by modulating the amount of traffic sent to a machine he can affect the amount of work it must do, and he can detect the resultant changes in system temperature by examining variations in the clock skew. Hence, if he accesses a ‘hidden’ machine via an anonymizing overlay network (such as Tor [21]) and varies how much traffic he sends to it, then it will heat up and cool down as the workload changes. If he can observe a corresponding pattern of clock skew change on a candidate machine to which direct access is possible, this is sufficient to link that machine to the hidden identity – and the anonymization scheme is overcome.

The opposite question of is often of interest – are machines physically different? Given two connections originating from the same network address, have they actually been initiated by one or multiple machines? It can be of particular relevance to count the number of machines behind NAT (Network Address Translation) gateways and firewalls. Bellovin [7] noted that the TCP/IP stack of many operating systems provides a host specific signature that can be detected, and used to estimate the number of hosts behind a gateway. To be exact the IPID field, used as a unique number for each IP packet, is in the Windows operating system a simple counter that is incremented every time a packet is transmitted. By plotting the IPID packets over time, and fitting lines through the graph, one can estimate the number of unique Windows hosts. However, this technique is becoming less effective because many systems now scramble the IPID field to prevent ‘idle scanning’ and so more complex analysis would now be necessary.

In IPv6 (the latest Internet protocol version) device addresses consist of a 64bit network prefix and a 64bit network identifier. This identifier needs to be unique, and initial proposals were for it to be constructed from the 48bit Ethernet MAC address for the interface. However, this provides a way for remote systems to link visits from the same mobile machine, despite them coming from different network locations. Narten and Draves (RFC3041) [53] developed a ‘privacy extension’ for allocating identifiers randomly, and Aura (RFC3972) [5] documented a method of creating IPv6 addresses that are bound to a public key, so that machines could formally demonstrate address ownership without disclosing their identity to remote systems. However, Escudero Pascual [59] criticises these schemes, particularly because it is possible for remote machines to determine that visitors are using privacy preserving addresses – which may of itself be sufficient to make their

traffic stand out.

Finally, many network mapping techniques have been introduced in the applied security world, and included in tools such as `nmap` [27]. The key operations that such tools perform are scanning for network hosts, scanning for open network ports on hosts, and identifying the operating systems and services running on them – this info being used to assess whether they might be vulnerable to attack. The degree of sophistication of these tools has increased with the deployment of network intrusion detection (IDS) tools, such as the open source `snort` [64], that can detect the scanning activities. `nmap` now can be configured to detect hosts and open ports using a variety of techniques including straightforward `ping`, TCP connect, TCP SYN packet, but also indirect scans. For example, ‘idle scanning’ involves forging a TCP open (SYN) packet claiming to be from a cut-out machine and destined to the target. It is possible to determine whether the target was prepared to accept the connection (it will send SYN/ACK) or if the port is ‘closed’ (it will send RST or nothing). This is done by determining if the IPID value of the cut-out machine has been altered by the sending of a RST in response to the unexpected (to it) SYN/ACK. The obvious advantage is that any IDS system at the target will believe that the cut-out machine is to the instigator of the scan. The full `nmap` documentation is well worth a read [28].

4.5 Detecting Stepping Stones

Much work has been done by the intrusion detection community to establish if a host is being used as an attack platform [72, 10]. The usual scenario involves a firewall that sees incoming and outgoing connections, and tries to establish if a pair of them may be carrying the same stream. This might mean that the internal machine is compromised and used to attack another host, i.e. it is a ‘stepping stone’ for the attacker to hide their identity.

The two main classes of techniques for detecting stepping stones are *passive*, where the firewall only observes the streams, and *active*, where the stream of data is modulated (often called ‘water-marked’). Since an adversary is controlling the content of the stream, and maybe encrypting it, both types of detection rely on traffic data – usually the correlation between packet inter arrival times – to match incoming and outgoing streams. The family of traffic analysis techniques that arise are similar to those that are used to attack anonymous communication channels.

The key result in this area [73, 60] is that if the maximum latency of the communication is

bounded there is no way of escaping detection in the long run. This result is of course tied to a particular model (the adversary can match packet for packet, which is not obvious if the streams are encrypted under different keys or mixed with other streams), and ‘covert’ channels out of its scope may prove it wrong and escape detection. It is worth observing that an arbitrary set of active detectors is extremely difficult (maybe even impossible) to defeat.

5 Exploiting Location Data

Wireless communication equipment often leaks location data to third parties, or wireless operators. The extent to which this data can be used to degrade security properties is still to be seen, but some experiments have already been performed, and their results are a precursor of a much richer set of attacks to come.

Escudero Pascual [58] describes an experiment he set up at the ‘Hacker’s at Large’ (HAL) summer camp. The camp had multiple wireless LAN access points, which recorded the wireless MAC address of the users whose traffic they handled. This provided a time-map of users’ movements throughout the event, including clues about which talks they attended (the access points were related to the venues). Even more striking were the inferences that could be drawn about the relationship between users: random pairs of users could be expected to have a low probability of using the same access point at any time and access point usage between them should be uncorrelated over time. As a result, any above average correlation between two users, is indicative of a social relationship between the users, i.e. they are consistently moving together at the same time around the camp.

Intel Research at Cambridge designed a similar experiment. Members of staff were issued with Bluetooth devices that would record when another transmitting Bluetooth device was in range. The idea was to measure the ambient Bluetooth activity, not only to tune ad-hoc routing protocols for real world conditions, but also to establish how often a random pair of devices meet – thereby establishing how effective the ad-hoc communication infrastructure would be for two way communications. To the surprise of the researchers analyzing the data, the devices of two members of staff were found to be meeting each other rather often at night – which led them to draw conclusions about their, otherwise undisclosed, relationship.

This is completely in line with evidence gathered by the MIT ‘Reality Mining’ project [47]. The project distributed about a hundred mobile phones to students and staff of the Media Lab, under the condition that all their traffic data (GSM, Bluetooth and location data) could be used for analysis. The users were also asked to fill in forms about themselves and who they considered to be their friends or colleagues. The traffic data and questionnaires were then used to build classifiers: it turned out that calling or being with someone at 8pm on a Saturday night is a very good indicator of friendship.

They also uncovered location signatures that could differentiate a student from a member of staff. What is even more impressive is that they did not use the physical locations to draw inferences, but instead the frequency at which they were found to be at places designated as ‘work’ or ‘home’. Students tended to have a more uncertain schedule, while members of staff were much more predictable in their habits. This of course led to research about the amount of entropy that location data provides, and as might be expected, for some individuals if one is given a set of locations and time, it is possible to predict with high probability their next move and new location.

So the evidence from these preliminary studies is highly suggestive that whatever the wireless medium used, mobile phone, wireless LAN or bluetooth, sensitive information about your identity, your relations to others and your intentions can be inferred merely through traffic analysis.

6 Resisting Traffic Analysis on the Internet

A relatively old, but only recently mainstream, sub-area of computer security research is concerned with ‘anonymous communications’ and more generally communications that do not leak any residual information from their meta data. The field was started by Chaum [12], introducing the ‘mix’ as a basic building block for anonymity, and has continued since, adapting the techniques to provide private email communications and more recently web-browsing. A thorough overview of the field and key results is available in two recent PhD theses by Danezis and Serjantov [16, 63].

Fielded anonymous communication systems, that are the direct products of twenty years of research, include Mixmaster [48] and Mixminion [17] for email, and JAP [8] and Tor [21] for web-browsing. They all increase the latency of communication and its cost in terms of traffic volumes.

A range of traffic analysis attacks have been used to degrade the security of anonymous commu-

nications networks. Long term intersection attacks (also referred to as disclosure attacks) exploit long term observations of input and output messages to detect communicating parties. These attacks [39, 40, 3, 18] consider the anonymity network as a black box, and only observe parties sending and receiving messages. The key observation is that for anonymous communications to be usable, the latency of messages has to be bounded. As a result the act of sending a message is correlated in time, albeit not exactly, with observing the corresponding message being received. An adversary can therefore observe the anonymity system for a sufficiently long period to obviate the lack of exactness, and infer the communication relationships between different users, and in turn de-anonymize the messages. Since this family of attacks is not concerned with the internals of the anonymity network, it is considered to represent a fundamental limit on how well any such technology can protect users against traffic analysis.

Stream traffic analysis has been used to trace web requests and replies through low-latency networks. Such attacks make use of the timing of the packet streams transferred by each anonymizing relay to follow the connection between the communicating parties. Packet counting is the simplest variant – an adversary simply counts the number of packets in a certain time interval and tries to match it with the number of packets on another network link [63]. Low-latency anonymity systems are required to transport packets so quickly that this attack is often possible. A slightly more sophisticated method involves creating a template (a probabilistic model) of the stream to be traced, and matching it with other streams [16]. Unless a very strict traffic regime is imposed, with the side effect of slowing down data transfer or adding large amounts of dummy traffic, such attacks will always be successful in the long run. As a result, stream tracing attacks also represent a fundamental limit on the anonymity of low latency systems.

Finally the attacker can infiltrate the network or try to influence the way in which honest nodes chose paths to anonymize their traffic. An important study of the effect of insiders on the security of anonymity systems is presented by Wright *et al* [76], along with the predecessor attack on the crowds anonymity system. Crowds implements a simple pass-the-parcel algorithm to anonymize traffic: messages are passed from one node to the other, until one of them – with some preset probability – sends it out onto the network. Only link encryption is used, and the intention is that anonymity will be achieved because although nodes will know the content of messages, they will be unable to tell who the initial sender of the message was. The predecessor attack relies nodes

having persistent patterns of communications. This means that the actual initiator will appear as the predecessor of a particular message or request rather more often than other random nodes (that merely relay the communications).

Lately attacks have focused on weaker adversaries, such as those considered by the Tor system, and it has been shown that some forms of traffic analysis can be performed without even having any access at all to the actual data streams to be traced. In particular, remote network monitoring techniques have been used to lower the anonymity of Tor [51]. Streams travelling over the same infrastructure influence each others timing, and therefore can be used by an adversary to perform traffic analysis on remote hosts. Similarly, as already mentioned, covert channels based on the effects of heat on clock drift can be used to de-anonymize servers [50]. The fact that even such minuscule phenomena can be used to perform traffic analysis against hardened systems illustrates how difficult the task of securing systems against traffic analysis is. It also illustrates that so little importance has been paid to securing public networks against traffic analysis that the information leaked can be detected and abused far, far away from its source.

Source and destination network addresses are not the only raw material for traffic analysis: the timing characteristics of encrypted traffic on a link, such as its frequency or particular bursts, may also reveal information to a third party (as seen with the examples of SSL and SSH). Military and diplomatic circles have long been avoiding this problem by using line encryptors that fill a leased line with ciphertext, no matter if any information is being transmitted. This prevents an enemy noticing that traffic has either increased (or indeed decreased) as the result of an event (as, apocryphally, it is said that volume of late-night Pentagon pizza orders change when hostilities are imminent [31, 74]).

Fixed rate encryption equipment is expensive to purchase (and operate) so there is a temptation to move to off the shelf routers, software encryption and the use of general purpose wide-area network links. Very little research has been done on protecting encrypted Internet Protocol links against traffic analysis, despite warnings about the threat posed against standard protocols like IPSec [6] and TLS. Venkatraman and Newman-Wolfe [55, 70] have looked at imposing traffic schedules to minimize information leaked as well as covert channels. Ways to analyse the cost and anonymity provided by such systems is presented in [54]. The earliest mention of this problem can be found in 1983 [71], with the conclusion that “beyond the host level, further limitation on in-

formation release becomes increasingly expensive and are probably not necessary in a non-military environment”.

A related problem, of concern in military circles, is that an enemy could observe a network and even though all the traffic was encrypted, determine the function of each node through traffic analysis. A weather station would generate reports on an hourly basis, but the more interesting target of the military headquarters could be distinguished by the multiple flows in and out of its node. The US DARPA Agency set this problem as one of their Challenges in 1998 [19] and it has been addressed, albeit only for fairly limited network topologies, in a number of papers from Guan *et al* [33, 32, 34] that consider adding extra traffic (padding) and rerouting some of the traffic along alternative network paths.

7 Data Retention

For some time, Law Enforcement officers (the police, secret services etc) have been using telephone call traffic data to identify criminals. Initially very simple enquiries were made: determining who made the last call that the murder victim received, tracking the source of a ransom demand, and so on. However, there has been a growing use of genuine traffic analysis techniques to develop ‘friendship trees’ and thereby identify the roles of individuals within a conspiracy. However, the denationalisation of incumbent fixed line telephone companies has broken their close ties with the police, and the growth of mobile telephone usage has led to a fragmentation of the market and fierce price competition, so that collection and storage of traffic data is now seen as an expensive burden. At the same time, new flat-rate business models have seen the business justification for call traffic data disappear. This has led to considerable anxiety within Law Enforcement that a valuable source of information will cease to be available.

In parallel, criminals have started to use the Internet for their communications and Law Enforcement has found that within this open system, with an extremely disparate set of service providers, the ‘traceability’ of communications can be problematic, and traffic analysis almost impossible. In particular, there has been concern that voice traffic will migrate from the closed and ordered telephony world to ‘Voice over IP’ (VoIP) running on the open and anarchic Internet.

In response, particularly after the terrorist attacks in Madrid (2004) and London (2005) interest

grew in mandatory ‘Data Retention’, requiring communications service providers to retain their traffic data logs for a fixed period, often far longer than their business needs would require. The term Data Retention should be contrasted with a ‘Data Preservation’ regime, where data is preserved specially in response to a specific request from Law Enforcement.

The United States has long had a Data Preservation regime, but in 2006 Congress started being pressured to consider moving to a Data Retention regime, with online child exploitation being cited as unnecessarily hard to investigate [61]. Meanwhile, in Europe, the EU adopted the Data Retention Directive (2006/24/EC) in March 2006 [23]. This provides for telephone companies to implement Data Retention by September 2007 and Internet companies by March 2009 at the latest. There is some doubt over the legal status of the Directive, which is being challenged (early 2007) by Ireland on the basis that it should have been implemented under ‘Third Pillar’ procedures for ‘Police and Judicial Co-operation in Criminal Matters’ rather than a ‘First Pillar’ Directive for ‘Market Harmonisation’. In practice, even though it is a Directive there is little harmonisation with EU member states free to decide on retention periods of anything between six months and two years, and with such technically incompetent definitions having been chosen that it they could refer to every point-to-point connection made over the Internet, or merely to records of emails passing through major servers. It looks like being several years before any clarity emerges, and it is very likely indeed that retention regimes will differ markedly in different countries.

Notwithstanding all this technical confusion, there has been very little informed debate on the types of information that will be capable of being extracted from the retained data. As should be apparent from even the limited survey we have presented in this chapter, there is significant scope for drilling down to reveal the most private of information about activities, habits, interests and even opinions. Storing this data, in an easily accessible manner, represents a systemic vulnerability that cannot be overstated enough.

In order to make balanced judgments between the needs of Law Enforcement and the entitlement of law-abiding citizens to privacy, policy makers must become far more aware of the wealth of information that could be extracted from such data about every aspect of the networked society. Even the extraction of apparently anonymous profiles from traffic databases would greatly facilitate privacy violations and routine surveillance. We believe that resistance to Traffic Analysis must be perceived of as a public good – the more that any attacker knows about the habits of your neighbours

the more they can tell about you!

8 And finally...

Our understanding of the threat that traffic analysis attacks represent on public networks is still fragmented, and research in this growing field is still very active. The results we have presented should act as a warning against ignoring this threat: traffic analysis not only can be used to reveal what is going on, but can also be used to bypass apparently robust security mechanisms.

References

- [1] *2005 IEEE Symposium on Security and Privacy (S&P 2005)*, 8–11 May 2005, Oakland, CA, USA. IEEE Computer Society, 2005.
- [2] *2006 IEEE Symposium on Security and Privacy (S&P 2006)*, 21–24 May 2006, Berkeley, California, USA. IEEE Computer Society, 2006.
- [3] Dakshi Agrawal and Dogan Kesdogan. Measuring anonymity: The disclosure attack. *IEEE Security & Privacy*, 1(6):27–34, 2003.
- [4] Ross Anderson. *Security engineering*. Wiley, 2001.
- [5] Thomas Aura. Cryptographically generated addresses (CGA), March 2005.
- [6] Steven M. Bellovin. Probable plaintext cryptanalysis of the ip security protocols. In *NDSS*. IEEE Computer Society, 1997.
- [7] Steven M. Bellovin. A technique for counting natted hosts. In *Internet Measurement Workshop*, pages 267–272. ACM, 2002.
- [8] Oliver Berthold, Hannes Federrath, and Stefan Köpsell. Web MIXes: A system for anonymous and unobservable Internet access. In H. Federrath, editor, *Designing Privacy Enhancing Technologies*, volume 2009 of *LNCS*, pages 115–129. Springer-Verlag, July 2000.
- [9] George Dean Bissias, Marc Liberatore, David Jensen, and Brian Neil Levine. Privacy vulnerabilities in encrypted HTTP streams. In George Danezis and David Martin, editors, *Privacy Enhancing Technologies*, volume 3856 of *Lecture Notes in Computer Science*, pages 1–11. Springer, 2005.

- [10] Avrim Blum, Dawn Xiaodong Song, and Shobha Venkataraman. Detection of interactive stepping stones: Algorithms and confidence bounds. In Erland Jonsson, Alfonso Valdes, and Magnus Almgren, editors, *RAID*, volume 3224 of *Lecture Notes in Computer Science*, pages 258–277. Springer, 2004.
- [11] Kathleen M Carley, Ju-Sung Lee, and David Krackhardt. Destabilizing networks. *Connections*, 22(3):79–92, 2002.
- [12] David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, February 1981.
- [13] Heyning Cheng and Ron Avnur. Traffic analysis of SSL encrypted web browsing, 1998.
- [14] Renesys Corporation. Market intelligence provides objective analysis of the service provider market. http://www.renesys.com/products_services/market_intel/.
- [15] George Danezis. Traffic analysis of the HTTP protocol over TLS. <http://www.cl.cam.ac.uk/~gd216/TLSanon.pdf>, 2003.
- [16] George Danezis. *Better Anonymous Communications*. PhD thesis, University of Cambridge, 2004.
- [17] George Danezis, Roger Dingledine, and Nick Mathewson. Mixminion: Design of a type III anonymous remailer protocol. In *IEEE Symposium on Security and Privacy*, Berkeley, CA, 11–14 May 2003. IEEE.
- [18] George Danezis and Andrei Serjantov. Statistical disclosure or intersection attacks on anonymity systems. In Fridrich [26], pages 293–308.
- [19] Defense Advanced Research Projects Agency. Research challenges in high confidence networking, July 1998.
- [20] Whitfield Diffie and Susan Landau. *Privacy on the Line: The Politics of Wiretapping and Encryption*. MIT Press, 1998.
- [21] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th USENIX Security Symposium*. USENIX, August 2004.
- [22] Matthew Dombroski, Paul Fischbeck, and Kathleen Carley. Estimating the shape of covert networks. In *Proceedings of the 8th International Command and Control Research and Technology Symposium. Conference held at the National Defense War College, Washington DC. Evidence Based Research, Track 3, Electronic Publication, Vienna, VA. CCRP*, 2003.
- [23] European Union. Directive 2006/24/EC of the European Parliament and of the Council of 15 March 2006 on the retention of data generated or processed in connection with the provision of publicly available electronic communications services or of public communications networks and amending Directive 2002/58/EC. *Official Journal of the European Union*, L 105:54–63, 2006.

- [24] Tom Fawcett and Foster Provost. Adaptive fraud detection. *Journal Data Mining and Knowledge Discovery*, 1(3):291–316, July 1997.
- [25] Edward W. Felten and Michael A. Schneider. Timing attacks on web privacy. In *ACM Conference on Computer and Communications Security*, pages 25–32. ACM, 2000.
- [26] Jessica J. Fridrich, editor. *Information Hiding, 6th International Workshop, IH 2004, Toronto, Canada, May 23–25, 2004, Revised Selected Papers*, volume 3200 of *Lecture Notes in Computer Science*. Springer, 2004.
- [27] Fyodor. Nmap – free security scanner for network exploitation and security audit. <http://www.insecure.org/nmap/>.
- [28] Fyodor. Nmap manual, 2006. <http://www.insecure.org/nmap/man/>.
- [29] Simson L. Garfinkel. Leaderless resistance today. *First Monday*, 3(8), 2003.
- [30] David Gibson, Jon Kleinberg, and Prabhakar Raghavan. Inferring web communities from link topology. In *Proceedings of the 9th ACM Conference on Hypertext and Hypermedia*. ACM, 1998.
- [31] Paul Gray. And bomb the anchovies. *Time*, 13 August, 1990.
- [32] Y. Guan, X. Fu, R. Bettati, and W. Zhao. Efficient traffic camouflaging in mission-critical QoS-guaranteed networks. In *Proceedings of IEEE Information Assurance and Security Workshop, West Point, June 2000*, pages 143–149. IEEE, 2000.
- [33] Y. Guan, C. Li, D. Xuan, R. Bettati, and W. Zhao. Preventing traffic analysis for real-time communication networks. In *Proceedings of The IEEE Military Communication Conference (MILCOM) '99, November 1999*. IEEE, 1999.
- [34] Yong Guan, Xinwen Fu, Dong Xuan, P. U. Shenoy, Riccardo Bettati, and Wei Zhao. Netcamo: camouflaging network traffic for qos-guaranteed mission critical applications. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 31:253–265, 2001.
- [35] Michael Herman. *Intelligence Power in Peace and War*. Cambridge University Press, 1996.
- [36] Andrew Hintz. Fingerprinting websites using traffic analysis. In Roger Dingledine and Paul F. Syverson, editors, *Privacy Enhancing Technologies*, volume 2482 of *Lecture Notes in Computer Science*, pages 171–178. Springer, 2002.
- [37] Victoria Hougham. Sociological skills used in the capture of Saddam Hussein. *Footnotes, Newsletter of the American Sociological Association*, 33(6), July/August 2005.
- [38] David Kahn. *The Codebreakers*. Scribner, 1967.

- [39] Dogan Kesdogan, Dakshi Agrawal, Vinh Pham, and Dieter Rautenbach. Fundamental limits on the anonymity provided by the mix technique. In *Security & Privacy* [2], pages 86–99.
- [40] Dogan Kesdogan and Lexi Pimenidis. The hitting set attack on anonymity protocols. In Fridrich [26], pages 326–339.
- [41] Jon M. Kleinberg. Hubs, authorities, and communities. *ACM Computing Surveys*, 31(4es):5, 1999.
- [42] Peter Klerks. The network paradigm applied to criminal organisations. *Connections*, 24(3):53–65, 2001.
- [43] Tadayoshi Kohno, Andre Broido, and Kimberly C. Claffy. Remote physical device fingerprinting. In *IEEE Symposium on Security and Privacy* [1], pages 211–225.
- [44] Valdis E. Krebs. Uncloaking terrorist networks. *First Monday*, 7(4), April 2002.
- [45] Mark Levene and George Loizou. Computing the entropy of user navigation in the web. *International Journal of Information Technology and Decision Making*, 2:459–476, 2003.
- [46] Raphael L. Levien. *Attack resistant trust metrics*. PhD thesis, University of California at Berkeley, 1995. Draft Thesis.
- [47] MIT Media Lab Human Dynamics Group. Reality mining. <http://reality.media.mit.edu/>.
- [48] U. Moeller, L. Cottrell, P. Palfrader, and L. Sassaman. Mixmaster protocol version 2. Technical report, Network Working Group, May 25 2004. Internet-Draft.
- [49] Fabian Monrose, Michael K. Reiter, and Susanne Wetzels. Password hardening based on keystroke dynamics. In *ACM Conference on Computer and Communications Security*, pages 73–82. ACM, 1999.
- [50] Steven J. Murdoch. Hot or not: revealing hidden services by their clock skew. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *13th ACM Conference on Computer and Communications Security (CCS), Alexandria, VA, USA*, pages 27–36. ACM, 2006.
- [51] Steven J. Murdoch and George Danezis. Low-cost traffic analysis of Tor. In *IEEE Symposium on Security and Privacy* [1], pages 183–195.
- [52] Shishir Nagaraja and Ross Anderson. The topology of covert conflict. Technical Report UCAM-CL-TR-637, University of Cambridge, Computer Laboratory, 2005.
- [53] T. Narten and R. Draves. Privacy extensions for stateless address autoconfiguration in IPv6, January 2001.
- [54] Richard E. Newman, Ira S. Moskowitz, Paul F. Syverson, and Andrei Serjantov. Metrics for traffic analysis prevention. In Roger Dingledine, editor, *Privacy Enhancing Technologies*, volume 2760 of *Lecture Notes in Computer Science*, pages 48–65. Springer, 2003.

- [55] R.E. Newman-Wolfe and B.R. Venkatraman. High level prevention of traffic analysis. *Seventh Annual Computer Security Applications Conference*, pages 102–109, 1991.
- [56] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford Digital Library Technologies Project, 1998.
- [57] Trevor Paglen and A. C. Thompson. Planespotting: Nerds with binoculars bust the CIA’s torture taxis. *The Village Voice*, 15 October, 2006.
- [58] Alberto Escudero Pascual. *Anonymous Untraceable Communications: Location privacy in mobile internetworking*. PhD thesis, Royal Institute of Technology – KTH / IMIT, 2001.
- [59] Alberto Escudero Pascual. Privacy extensions for stateless address autoconfiguration in IPv6 – “requirements for unobservability”. In *RVP02, Stockholm*. Stockholm, June 2002.
- [60] Pai Peng, Peng Ning, and Douglas S. Reeves. On the secrecy of timing-based active watermarking trace-back techniques. In *SEIP [2]*, pages 334–349.
- [61] Rodney Petersen. Towards a u.s. data-retention standard for ISPs. *Educause Review*, 41(6):78–79, 2006.
- [62] William J. Reed. A brief introduction to scale-free networks. Technical report, Department of Mathematics and Statistics, University of Victoria, 2004.
- [63] Andrei Serjantov. On the anonymity of anonymity systems. Technical Report UCAM-CL-TR-604, University of Cambridge, Computer Laboratory, 2004.
- [64] Snort team. Snort. <http://www.snort.org/>.
- [65] Dawn Xiaodong Song, David Wagner, and Xuqing Tian. Timing analysis of keystrokes and timing attacks on SSH. In *Tenth USENIX Security Symposium*. USENIX, 2001.
- [66] Malcolm K. Sparrow. The application of network analysis to criminal intelligence: An assessment of the prospects. *Social Networks*, 13:251–274, 1991.
- [67] Qixiang Sun, Daniel R. Simon, Yi-Min Wang, Wilf Russell, Venkata N. Padmanabhan, and Lili Qiu. Statistical identification of encrypted web browsing traffic. In *IEEE Symposium on Security and Privacy*, pages 19–30. IEEE, 2002.
- [68] J. Travers and S. Milgram. An experimental study of the small world problem. *Sociometry*, 32(425), 1969.
- [69] US Army Headquarters Department. *Field Manual 3-24: Counterinsurgency*. US Army, 2006.
- [70] BR Venkatraman and RE Newman-Wolfe. Transmission schedules to prevent traffic analysis. *Proceedings, Ninth Annual Computer Security Applications Conference*, pages 108–115, 1993.

- [71] V.L. Voydock and S.T. Kent. Security Mechanisms in High-Level Network Protocols. *ACM Computing Surveys (CSUR)*, 15(2):135–171, 1983.
- [72] Xinyuan Wang and Douglas S. Reeves. Robust correlation of encrypted attack traffic through stepping stones by manipulation of interpacket delays. In Sushil Jajodia, Vijayalakshmi Atluri, and Trent Jaeger, editors, *ACM Conference on Computer and Communications Security*, pages 20–29. ACM, 2003.
- [73] Xinyuan Wang, Douglas S. Reeves, and Shyhtsun Felix Wu. Inter-packet delay based correlation for tracing encrypted connections through stepping stones. In Dieter Gollmann, Günter Karjoth, and Michael Waidner, editors, *ESORICS*, volume 2502 of *Lecture Notes in Computer Science*, pages 244–263. Springer, 2002.
- [74] Andrew Warinner. Security clearances required for Domino's, 1996. <http://home.xnet.com/~warinner/pizza.html>.
- [75] Stanley Wasserman, Katherine Faust, Dawn Iacobucci, and Mark Granovetter. *Social Network Analysis : Methods and Applications (Structural Analysis in the Social Sciences)*. Cambridge University Press, 1994.
- [76] Matthew Wright, Micah Adler, Brian Neil Levine, and Clay Shields. An analysis of the degradation of anonymous protocols. In *NDSS*. The Internet Society, 2002.