Data Retention and Anonymity Services
Introducing a New Class of Realistic Adversary Models

Stefan Berthold, Rainer Böhme, Stefan Köpsell

Technische Universität Dresden
Faculty of Computer Science
01062 Dresden, Germany
{stefan.berthold|rainer.boehme|stefan.koepsell}@tu-dresden.de

Abstract. The recently introduced legislation on data retention to aid prosecuting cyber-related crime in Europe also affects the achievable security of systems for anonymous communication on the Internet. We argue that data retention requires a review of existing security evaluations against a new class of realistic adversary models. In particular, we present theoretical results and first empirical evidence for intersection attacks by law enforcement authorities. The reference architecture for our study is the anonymity service AN.ON, from which we also collect empirical data. Our adversary model reflects an interpretation of the current implementation of the EC Directive on Data Retention in Germany.

1 Introduction

In the absence of anonymising technology, every computer connected to the Internet communicates with a unique address. So online users can be identified by the address of their communication device and the time of activity. The objective of anonymity services is to hide the relation between individual users and addresses from the users' communication partners and, with certain technologies, also from possible eavesdroppers on the communication links. However, by using an anonymity service, Internet users forward (part of) their traffic to the anonymity service, which such obtains all information necessary to re-identify anonymised users. Therefore, in principle, every anonymity service should be constructed in a way to delete this information as soon as possible in order to protect itself from becoming a target for adversaries who are interested in de-anonymising Internet users. This was valid until lately.

Recent legislation in the European Union, and particularly in Germany, requires anonymity services to store this sensitive data for months before it can be deleted. This is understood as a necessary trade-off between the interest of data protection and law enforcement: anonymity services are susceptible to be abused for criminal activities. In such cases, anonymity should indeed be revocable (though other means than data retention have been proposed to achieve this end [1, 2]).

In common terms of the literature on privacy and anonymity, deanonymisation by means of data retention can be considered as kind of attack. In contrast to other typically studied attacks on anonymity services, here the capabilities of the adversary are determined by law. This opens up a remarkably clear insight into what is in and what is out of control of the adversary and thus outlines a pretty specific adversary model. The characteristics of this kind of adversary model are different from common models in the relevant literature. The centre of interest has shifted from whether the adversary is *able to snoop or infer* private data to the extent the anonymity service is *obliged to retain (and provide on request)* the data. Obviously, any party that has (unauthorised) access to the data falls under this adversary model.

Thus, data retention regulations may mark a turning point for the design and the analysis of anonymity services. The question we will face in future is about how much anonymity is legally achievable under attacks that make use of retained data. And, correspondingly, one challenge will be to adapt current anonymity services or construct new ones with appropriate features to resist these attacks or mitigate their (side-) effects.

This introductory paper on the new class of legally defined adversary models can only focus on selected specific aspects, namely mix networks and intersection attacks. More specifically, we study the advantage which the adversary may gain from data that has been retained in line with the legal framework in Germany. The mix reference implementation for our study is AN.ON, an anonymity service which has been developed at TU Dresden and has been running successfully on the Internet for years. To quantify the impact of the attacks, simulations have been conducted based on data which has been gathered from a part of the broad AN.ON user community,[1] who gave their consent to participate in our study. Our main result is that dummy traffic, though not always a strong measure against arbitrary adversaries, is strikingly helpful against the law-abiding adversaries.

The paper is structured as follows. Section 2 recalls very briefly the essential principles of anonymity services. In Section 3, we give our interpretation of the current legislation. In the absence of case law, we take this interpretation as a base for the following sections. It also defines the adversary model for the rest of the paper. In Sections 4 and 5, we describe the cross-section attack and the intersection attack. We expect that these two attacks are most likely to be mounted on retained data in order to compromise or revoke the anonymity of AN.ON users. In Section 6, we describe the setup of our study to measure the potential of intersection attacks using retained data. The results of this study are presented in Section 7. In Section 8, we extend intersection attacks to the case of uncertainty of the adversary about the connection between two observable events (existing versions of the attacks assume full certainty). Finally, Section 9 concludes the paper and points to further generalisations and research topics of interest against the backdrop of the new class of realistic adversary models.

[1] The number of AN.ON users can only be estimated, but not finally determined. AN.ON by default does not store data that allows to distinguish different users.

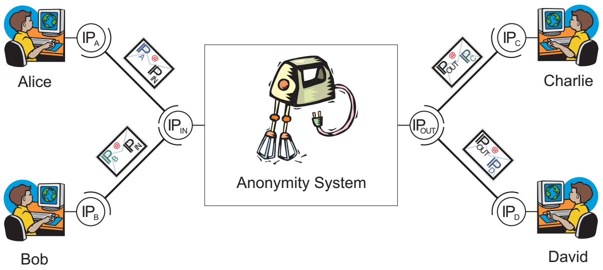

Fig. 1. Anonymity service modelled as "black box" which replaces IP addresses of forwarded messages.

2 Anonymity services in a nutshell

For the purpose of our study, we can understand an anonymity service as a "black box" which acts as a proxy (cf. Fig. 1). Users redirect their network traffic through the proxy in order to achieve anonymity. For instance, browsing the web without revealing the identity is a common application of anonymity services. In this context, we understand anonymity as the obfuscation of all relations that let an outsider, the adversary, learn about the links between incoming and outgoing proxy traffic. Consequently, the adversary would not be able to determine the persons who are exchanging messages through the proxy, if anonymity is preserved. This should even hold if the adversary eavesdrops the data on all communication links of the proxy. Anonymity can be achieved by a combination of cryptography and data handling methods, such as padding, reordering, delaying etc. [3]

3 Legal background

The directive 2006/24/EC (data retention directive) "on the retention of data generated or processed in connection with the provision of publicly available electronic communication services or of public communication networks", passed by the European parliament on March 15th, 2006, sets the legal framework of data retention for the European Union member states. According to the directive, the member states have to "bring into force the laws, regulations and administrative provisions necessary to comply with this directive by no later than 15 September 2007" [2006/24/EC]. The goal of the directive is to strengthen the success of law enforcement in the area of Internet-related crime and, more generally, whenever electronic communication is involved in criminal activities. The need for directive emerged from the fact that data about past communication relations is already unavailable when criminal offences come to trial, in which evidence from the communication relations could be helpful. Data on communication relations can provide indications about the person who accessed a specific website or who called a specific telephone, for instance.

Germany reacted to the data retention directive and adapted several laws [4]. With respect to anonymity services on the Internet, the changes of the Telecommunications Act are most significant [5]. This act defines in detail what kind of data has to be stored for various types of communications providers, including telecommunication companies like fixed-line or mobile phone providers and Internet service providers (ISPs). The act defines a retention period of six months. It anticipates services like anonymity services, which are in the first place contradictory to the law enforcement goals. In order to prevent any information gap, the Telecommunications Act declares in §113a 'Retention of Data':

> '(6) Those, who provide telecommunication services and thereby alter data which have to be stored according to this law, have to store the original data and the new data as well as the time of the alteration.'[2]

Anonymity services can be understood as proxy servers. The idea behind such proxies is briefly described in Section 2. In terms of sentence (6) of §113a of the Telecommunications Act the proxy, that is the anonymity service, replaces the IP addresses of senders and receivers with the proxy IP address in order to relay messages (cf. Fig. 1 above) and 'thereby alters data which have to be stored' according to the law. Consequently, anonymity services in principle have to store possibly identifying information about their users.

An urgent question is which data exactly has to be logged by anonymity services such as AN.ON in order to comply with the data retention law. In §113a, the Telecommunications Act distinguishes several types of services and defines for each service the sort of data to be stored. The closest match for AN.ON is 'Internet Service Provider' (ISP). According to the Telecommunications Act, an ISP has to log the IP address of a user, a unique identifier of the connection, and the period of time in which this assignment was valid. In combination with sentence (6), this means that the anonymity service has to log the replacement of IP addresses only, but nothing more, particularly no 'identifiers' of higher layers, such as TCP port numbers etc. Besides, consulted lawyers argue that only the replacement of source IP addresses (but not destination IP addresses) are allowed to be retained. They justify their assessment with sentence (8) of §113a: '...data about retrieved Internet pages must not be retained.' The lawyers also conclude that logging is allowed only for IP packet flows in upstream direction, that is only for packets from the user to the service, for instance a web server, but not for downstream packets.[3] In fact, the effective interpretation of the law remains uncertain until the German Federal Supreme Court makes a final decision. For

[2] Note that the quotations of the Telecommunications Act are unofficial translations of the official law in German. The authors are not aware of any official translation of the current version of the Telecommunications Act. The former version (of 22 June 2004) is available in English (online at: `http://www.bmwi.de/BMWi/Redaktion/PDF/Gesetz/telekommunkationsgesetz-en`).

[3] This is due to the fact that in a bidirectional communication, upstream and downstream are linakble. Thus logging of downstream source addresses implies logging of upstream destination addresses–which is prohibited by law.

this study, we assume that our interpretation is correct.[4] Thus, we can derive that anonymity services have to log the replacement of the original source IP address whenever an IP packet is forwarded from a user to a server. In other words, the anonymity service has to log the time and source IP address of *every* IP packet it receives from a user.

4 Cross-section attack

In this section, we study a very simple attack that could be mounted on retained data. This is at the same time the foundation for the intersection attack which will be introduced in Section 5. Looking from the perspective of law enforcement, the reply to the typical law enforcement request, "To which person was IP address $\mathrm{IP}_{\mathrm{out}}$ assigned at time t?" (Q_1), would include all retained source IP addresses for time t. We will refer to these IP addresses by the symbol $\mathcal{S}(t)$, that is, we consider $\mathcal{S}(t)$ denoting the set of retained IP addresses at time t.

The number of elements in $\mathcal{S}(t)$ can be understood as a measure of anonymity, as for instance in [7].[5] Critics of this way of measuring anonymity mention (rightly so) that the probability distribution of all elements in $\mathcal{S}(t)$ is not necessarily uniform [8, 9]. However, we know that the best case from the perspective of law enforcement would be, if $\mathcal{S}(t)$ contains one element only. This is also the worst case for anonymity by any measure. The size of $\mathcal{S}(t)$ depends on two parameters: (a) on the extent of use of the anonymity service and (b) on the resolution of the timestamp t. Note that the timestamp is not specified in greater detail by law.

We have quantified the activity of users[6] of our AN.ON system in order to get a better idea of $\mathcal{S}(t)$ and its size. To keep the task manageable, we decided to log the start and end time of *anonymous channels* only. The alternative would have been to log all incoming IP packets, but that would be rather expensive. In AN.ON, anonymous channels are the basic end-to-end communication vehicle, similar to a TCP/IP connection.

We found that nearly half of all channels lasted no more than one second, so we assume that analysing the channel activities leads to a good approximation of the actual size of $\mathcal{S}(t)$. Fig. 2 shows the results of the quantification at the 'Dresden–Dresden' cascade of our AN.ON system.[7] The red dots depict the total number of users logged in, regardless if they were active or idle. The black dots show the number of users with at least one open channel. For both aggregations, a time resolution of *one minute* was used. For comparison, the blue dots depict

[4] Other interpretations of the law can be found in the literature, e.g. in [6] the authors assume that: "the German legislation requires operators of anonymisers to link all incoming and outgoing messages and store this relation."

[5] The literature often refers to $\mathcal{S}(t)$ as the "anonymity set".

[6] When we speak about 'users' (e.g. number of users, activity of users etc.), we mean established connections to the AN.ON system. It is not possible to tell how many different human beings are behind them.

[7] Our AN.ON system is based on mix cascades. A cascade is a fixed chain of anonymity service servers (called mixes). Users may freely choose the cascade they want to use.

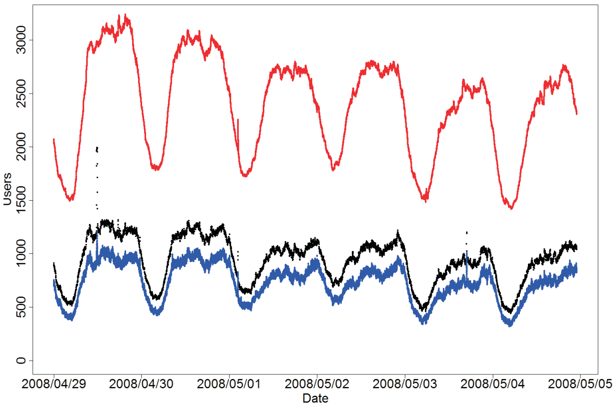

Fig. 2. Curves showing the total number of users logged in (red), the number of users which have an open channel within a given minute (black) and the number of users with an open channel within a given second (blue).

a setting which is similar to the black dots, but with a time resolution of one second.

Observe that the size of $\mathcal{S}(t)$, cf. Question Q_1, has never fallen below 400 between 29th of April 2008 and 5th of May 2008. That is, even when resolving a Q_1 request, a law enforcement agency would still have to investigate at least 400 users to identify the person they are looking for. As we see in Fig. 2, the accuracy of the time resolution (seconds vs. minutes, that is blue vs. black dots) is less important in practice than the overall usage rate of the anonymity service.

5 Intersection attack

A single request for $\mathcal{S}(t)$, that is the set of all retained users at a single point in time, might not be sufficient to narrow down the number of suspects to a reasonable small set, as we have seen in the previous section. Thus, a law enforcement agency could request the sets of online users for *several* points in time. With these sets, the agency would be able to mount an intersection attack [10] which, in theory, drastically narrows down the set size. Intersection attacks, however, require that the requested points in time are related to events that are linkable to one and the same target person. For the sake of simplicity, we assume that each event is related to exactly one point in time. Thus, a newly formulated request of a law enforcement agency would be “To which person was IP address $\mathrm{IP}_{\mathrm{out}}$ assigned at times t_1, t_2, and t_3?” (Q_2). If the law enforcement agency possesses a priori knowledge that one and the same target person is responsible for the events of interest observed at t_1, t_2, and t_3, then this person (or rather her identifier) belongs to the intersection of $\mathcal{S}(t_1) \cap \mathcal{S}(t_2) \cap \mathcal{S}(t_3)$.[8]

Note that events may basically occur on various layers, the application layer or the network layer, for instance. On the application layer, a law enforcement agency may observe that the same e-mail account was accessed several times. On the network layer, a law enforcement agency may run a honeypot and therefore obtain the exact timing of incoming IP packets which belong to one and the same TCP/IP connection.

6 Setup of our study on intersection attacks

6.1 Preparation of the AN.ON client software

In our study, we quantify the size of an anonymity set that remains after intersection attacks. The main problem with a study of intersection attacks on AN.ON user data is that (due to the very nature of anonymity services) there is no way to link the anonymised sessions of one and the same user. In order to get useful data for our study, additional identifiers were needed to be submitted by users to the AN.ON service.

We adapted the AN.ON client software such that users can decide whether they want to take part in the study. In the adapted software, a random number of 117 Bits has been generated as identifier for those users who take part in the study. The identifier has been transmitted to AN.ON each time the user logs in to the Dresden–Dresden cascade. Thus, sessions of users who voluntarily participated in the study became linkable over the time of the study.

The identifiers have been recorded between 21th of May and 20th of July, 2008. On 21th of May, the adapted client has been released and older clients reacted by requesting the update. Thus, we expect that in the following days, the vast majority of AN.ON users installed the new client and was therefore asked whether to participate in the study or not. In total, we recorded 70,591 replies, 38,738 (54.88%) of which agreed to support the study. The remaining 45.12% of the users continued to use AN.ON without any linkability of their sessions.

6.2 Formal notation

In the style of the symbol $\mathcal{S}(t)$, which we informally introduced in a previous section, we define the symbol $\mathcal{S}_\cap(T)$ as the AN.ON users which were retained at each of the times $t_1, \ldots, t_n \in T$. This requires the understanding of AN.ON *sessions*. The AN.ON client opens a session when it connects to the anonymity service. The session is closed when the client quits. Thus, a session is always related to a user and can be described by a login and a logout time. The formal definitions of $\mathcal{S}(t)$ and $\mathcal{S}_\cap(T)$ are reflected in Equation (2) and (3).

[8] In Section 8 we discuss the case where the linkability between two events is not possibilistic but probabilistic.

Let I_u be the set of all user IDs, I_s be the set of all session IDs, and $\mathcal{P}(I_s)$ the power set of all session IDs I_s. Then $X : I_u \to \mathcal{P}(I_s)$ would be the mapping of user IDs to all related sessions:

$$X(uid) = \big\{ sid \in I_s \;\big|\; sid \text{ related to } uid \big\}. \tag{1}$$

The login and logout time of a AN.ON session $sid \in I_s$ can be reflected in the two symbols $t_{\text{in}}(sid)$ and $t_{\text{out}}(sid)$. Then $\mathcal{S}(t)$ would be the set of users which have been logged in to AN.ON between the times t and $t + t_{\text{res}}$ where t_{res} is the time resolution (a second or a minute in our study):

$$\mathcal{S}(t) = \big\{ uid \in I_u \;\big|\; sid \in X(uid), t_{\text{in}}(sid) < t + t_{\text{res}}, t_{\text{out}}(sid) \geq t \big\} \tag{2}$$

With $\mathcal{S}(t)$, we can define $\mathcal{S}_\cap(T)$, the anonymity set after an intersection attack with the times $T = \{t_1, \ldots, t_n\}$. We suppose that all elements in T are pairwise different, that is for $T = \{t_1, \ldots, t_n\}$ holds $|T| = n$:

$$\mathcal{S}_\cap(T) = \bigcap_{t \in T} \mathcal{S}(t) \tag{3}$$

In our study, we focus on intersections between user sets of *two* points in time only. That is, we explore $\mathcal{S}_\cap(T)$ with the samples T where $|T| = 2$ and the elements of T are chosen by random. This setting can be understood as the case that law enforcement agencies request the set of persons which have been logged in at t_1 and at t_2 as well, or in $T = \{t_1, t_2\}$, respectively.[9]

7 Results of our study on intersection attacks

Table 1 shows characteristics of distributions of $|\mathcal{S}(t_1)|$, $|\mathcal{S}(t_2)|$, and $|\mathcal{S}_\cap(\{t_1, t_2\})|$ from our study with 5 million samples of two points in time t_1 and t_2 with variations in the time resolution (seconds vs. minutes) and the user data (login/logout vs. activity).[10]

Fig. 3(a) depicts two frequency density diagrams that show immediate results from our study with 5 million samples of two points in time t_1 and t_2. On the horizontal axis, we plot the size of $\mathcal{S}(t_1)$ or $S_\cap(\{t_1, t_2\}$. On the vertical axis, we see the frequency densities of these set sizes with regard to our samples. The red line marks the frequency densities of set sizes of $\mathcal{S}(t_1)$ (which are nearly the same as for $\mathcal{S}(t_2)$[11]). The blue line mark the frequency densities of set sizes of $\mathcal{S}_\cap(\{t_1, t_2\})$. The parameters and summary measures of the distributions are reported in Table 1. Similar results are shown in Fig. 3(b), 3(c), and 3(d) with variations in the time resolution and the user data (login/logout vs. activity).

[9] In practice, law enforcement agencies will rather request who used the IP address of the last mix of a cascade in T than requesting the login state of AN.ON users.

[10] With activity, we refer to channel activity as described in Section 4.

[11] Note that $\mathcal{S}(t_1)$ and $\mathcal{S}(t_2)$ are drawn from the same distribution.

Table 1. Summary statistics of the distributions of $|\mathcal{S}(t_1)|$, $|\mathcal{S}(t_2)|$, and $|\mathcal{S}_\cap(\{t_1,t_2\})|$ over 5 million random draws.

	user behaviour	time res.	min	1st quartile	mean	median	3rd quartile	max
$\mathcal{S}(t_1)$	login/out	*sec*	148	529	661	665.6	814	1210
		min	207	534	667	672	821	1220
	activity	*sec*	23	153	202	199.2	245	576
		min	77	181	238	235.9	288	802
$\mathcal{S}(t_2)$	login/out	*sec*	147	529	660	665.5	814	1210
		min	207	534	667	671.9	821	1220
	activity	*sec*	23	153	202	199.2	245	576
		min	77	181	238	235.9	288	802
$\mathcal{S}_\cap(\{t_1,t_2\})$	login/out	*sec*	9	91	120	133.1	159	1118
		min	11	92	121	134.1	161	1125
	activity	*sec*	0	7	11	13.39	17	307
		min	0	10	15	18.42	23	562

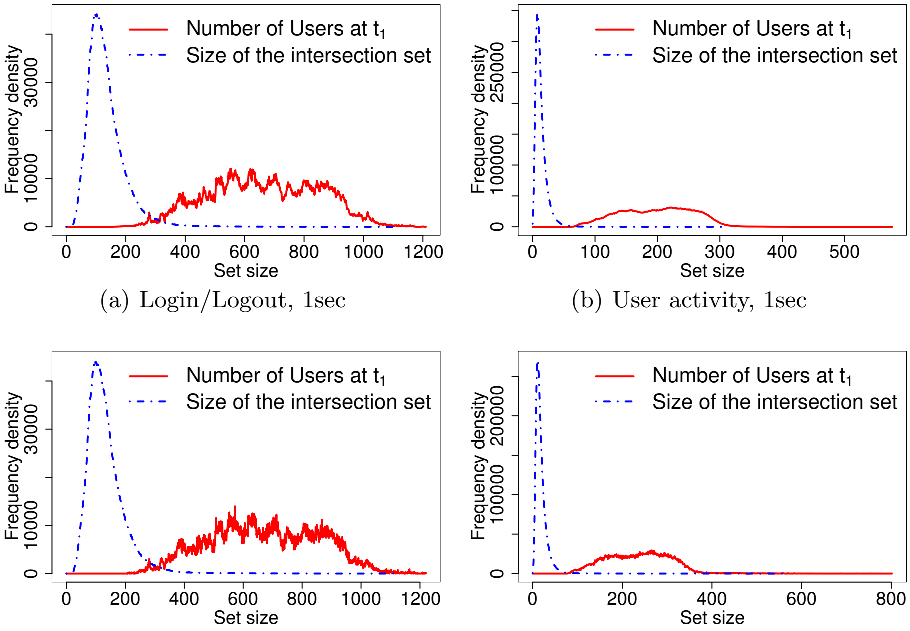

(a) Login/Logout, 1sec

(b) User activity, 1sec

(c) Login/Logout, 1min

(d) User activity, 1min

Fig. 3. Frequency density diagrams of the anonymity set sizes $|\mathcal{S}(t_1)|$ and $|\mathcal{S}_\cap(\{t_1,t_2\})|$.

Table 2. Parameters of different linear regression models with dependent variable $\ln\left|\mathcal{S}_\cap(\{t_1,t_2\})\right|$; $N = 5$ million data points; std. errors in brackets; all coefficient significant at 0.001 level.

	Model			
	1	2	3	4
Predictors				
intercept	-2.05 (0.04)	-3.13 (0.05)	1.53 (0.03)	1.56 (0.02)
minimum set size $\ln\left(\min\left(\lvert\mathcal{S}(t_1)\rvert, \lvert\mathcal{S}(t_2)\rvert\right)\right)$	1.09 (0.01)	1.00 (0.01)	1.00 (0.00)	0.98 (0.00)
maximum set size $\ln\left(\max\left(\lvert\mathcal{S}(t_1)\rvert, \lvert\mathcal{S}(t_2)\rvert\right)\right)$		0.25 (0.01)		
time interval $\ln\lvert\delta\rvert$			-0.22 (0.00)	-0.22 (0.00)
periodicity indicator $f_\triangle(\delta)$				0.24 (0.00)
Summary				
adjusted R^2	0.48	0.49	0.79	0.82

Observe that the anonymity set size is greater with a coarser time resolution and, even more significantly, if the adversary has no access to the activity data of AN.ON users, but only to their login/logout behaviour.

We fitted four regression models (using ordinary least squares) to analyse the multivariate relationship between the intersection size and various explanatory variables, cf. Fig. 3(a) and Table 1. The parameter estimates of these models are compiled in Table 2.

In Model 1, we try to explain the size of $\mathcal{S}_\cap(\{t_1,t_2\})$ by the minimum size of $\mathcal{S}(t_1)$ and $\mathcal{S}(t_2)$. The results are log transformed to reasonably normalise the residuals. Additionally, in Model 2, we add the maximum size of $\mathcal{S}(t_1)$ and $\mathcal{S}(t_2)$ as a second explanatory variable. We see that the gain of explained variance, cf. adjusted R^2 in Table 2, is small and the coefficient lower – albeit positive and statistically significant. This is what we expect, since the intersection set $\mathcal{S}_\cap(\{t_1,t_2\}) = \mathcal{S}(t_1) \cap \mathcal{S}(t_2)$ is at most as great as the smallest set of $\mathcal{S}(t_1)$ and $\mathcal{S}(t_2)$. As the set size fluctuates heavily over time, the size of the intersection is strongly related to the minimum size of $\mathcal{S}(t_1)$ and $\mathcal{S}(t_2)$.

In all following models, we drop the less-influentical maximum and use solely the minimum set size to control for a varying number of users over time. Model 3 includes the time interval between both events as predictor. The negative sign of the coefficient for term $\ln|\delta|$ indicates that there is an inverse relation between the time interval and the intersection size. That means, smaller time intervals lead to greater intersection sets, since the smaller the time interval between t_1 and t_2, the higher is the likelihood that a user who is logged in at t_1 is still logged in at t_2. The considerable gain in R^2 of 31 percentage points reveals that time between events matters.

In Model 4, we explore the influence of user behaviour on the set size of $\mathcal{S}_\cap(\{t_1, t_2\})$. We expect that the user behaviour follows regular pattern, for instance a periodicity of 24h.[12] This is so because we expect that users pursue similar tasks at similar times of the day. Users who log in to AN.ON during the working hours may regularly use AN.ON in their profession, for instance journalists. Those users who use AN.ON for their leisure time activities may regularly log in after the working hours. In order to check the support of our expectation in the sample data, we estimate the coefficient of an indicator variable computed from a periodic triangular function $f_\Delta(\delta)$ which generates an indicator variable that yields a value between 0 and 1, where a value of 0 marks the smallest match with the 24h pattern and a value of 1 denotes the best match.

$$\delta = |t_1 - t_2| \tag{4}$$

$$f_\Delta(\delta) = \left| 1 - \frac{\delta \bmod (24 \cdot 60^2)}{12 \cdot 60^2} \right| \tag{5}$$

The positive coefficient indicates that the sample data in fact shows periodicity, although the additional explanatory power of this simple linear function is rather small (3 percentage points).

8 Probabilistic intersection attacks

So far, we have assumed that the law enforcement agency has no uncertainty with respect to the linkability of a set of events. This might be true in some cases (e.g., if the IP packets per TCP/IP connection example above) – but not in general. In the above example of the e-mail account, the login information (username and password) could be shared among a group of persons. Therefore the intersection could lead to false negatives. Consequently, for practical cases it is necessary to consider the uncertainty about the likability of two or more events. This will turn possibilistic intersection attacks into probabilistic ones. Note that previous work in this field (cf. [11, 12]) has focused on calculating anonymity sets resulting from intersection attacks using *probabilistic algorithms*, whereas our contribution is to model linkability in a probabilistic sense.

In the following, we formally study the case of two events, E_1 and E_2. Let p be the probability that both events are caused by the same user.[13]

Further, as illustrated in Fig. 4, let $\mathcal{A}$ be the set of suspect senders (i.e., the result of a cross-section attack) at the time of E_1 and $\mathcal{B}$ the set of suspects for E_2. The intersection $\mathcal{A} \cap \mathcal{B}$ is the set of senders connected to the anonymity service at the time of both events (not necessarily in one session). We write the cardinality of set $\mathcal{X}$ as $|\mathcal{X}|$ with $\mathcal{X} = \mathcal{A}, \mathcal{B}$.

[12] It was not possible to explore patterns on a weekly or longer basis, since our study period was too short.

[13] Note that, in this model, there is no uncertainty about the value of p. Of course in an extended model one could also consider that p is not observable and can only estimated with a parametric model.

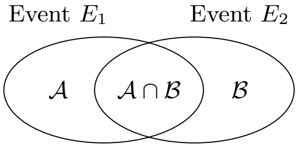

Fig. 4. Terminology for probabilistic intersection attack with $\mathcal{A}$ being the set of users that were logged in when event E_1 occurred and $\mathcal{B}$ being the set of users that were logged in when event E_2 occurred.

Problem statement: Given the quantities $|\mathcal{A}|$, $|\mathcal{B}|$, $|\mathcal{A} \cap \mathcal{B}|$ and p (the probability that both events were caused by the same sender, i.e., $p = \Pr(\exists S \in \mathcal{A} \cap \mathcal{B} \mid S \sim E_1 \wedge S \sim E_2)$, where '$\sim$' denotes a causal relationship), what is the probability for individual senders *having caused at least one event*, dependent on which set they belong to. More precisely, we want to calculate

1. $P_{\mathcal{A} \cap \mathcal{B}} = \Pr(S \sim E_1 \vee S \sim E_2 \mid S \in \mathcal{A} \cap \mathcal{B})$, the probability that a specific sender S who has used the anonymity service at the time of both events is responsible for at least one of the events;
2. $P_{\mathcal{A} \setminus \mathcal{B}} = \Pr(S \sim E_1 \vee S \sim E_2 \mid S \in \mathcal{A} \setminus \mathcal{B})$, the probability that a specific sender S who has used the anonymity service at the time of event E_1 *but not* at the time of event E_2 is responsible for at least one of the events; and vice versa,
3. $P_{\mathcal{B} \setminus \mathcal{A}} = \Pr(S \sim E_1 \vee S \sim E_2 \mid S \in \mathcal{B} \setminus \mathcal{A})$, the probability that a specific sender S who has not used the anonymity service at the time of event E_1 *but* at the time of event E_2 is responsible for at least one of the events.

The problem can be solved by evaluating a decision tree (Fig. 5). The root branches distinguish whether both events were actually caused by the same user or not. Hence, the probabilities for the branches are p and $1 - p$. If both events in fact origin from the same sender, then the only possibility is that the actual sender belongs to set $\mathcal{A} \cap \mathcal{B}$. Otherwise, four solutions for the assignment of sets of senders to events are possible, and their probabilities depend on the relative set sizes. Note that members of $\mathcal{A} \cap \mathcal{B}$ may have caused one or both (if $|\mathcal{A} \cap \mathcal{B}| > 1$) events even when both events were caused by different users. We further make the convention that $S_1 = S \in \mathcal{A} \iff S \sim E_1$ and $S_2 = S \in \mathcal{B} \iff S \sim E_2$.

Obviously, probabilities $\pi_1, \ldots, \pi_4$ can be calculated from the number of possibilities in relation to its total per sub-tree. The cardinalities of $\mathcal{A} \setminus \mathcal{B}$ and $\mathcal{B} \setminus \mathcal{A}$ are given implicitly as $|\mathcal{A} \setminus \mathcal{B}| = |\mathcal{A}| - |\mathcal{A} \cap \mathcal{B}|$ and $|\mathcal{B} \setminus \mathcal{A}| = |\mathcal{B}| - |\mathcal{A} \cap \mathcal{B}|$, and $\binom{a}{b}$ denotes the binomial coefficient. The probabilities of interest can be obtained by evaluating the decision tree in Fig. 5. As can be seen in Equation (6)–(8), these probabilities are linear in p.

$$P_{\mathcal{A} \cap \mathcal{B}} = p + (1 - p)(\pi_2 + \pi_3 + \pi_4) \tag{6}$$

$$P_{\mathcal{A} \setminus \mathcal{B}} = (1 - p)(\pi_1 + \pi_2) \tag{7}$$

$$P_{\mathcal{B} \setminus \mathcal{A}} = (1 - p)(\pi_1 + \pi_3) \tag{8}$$

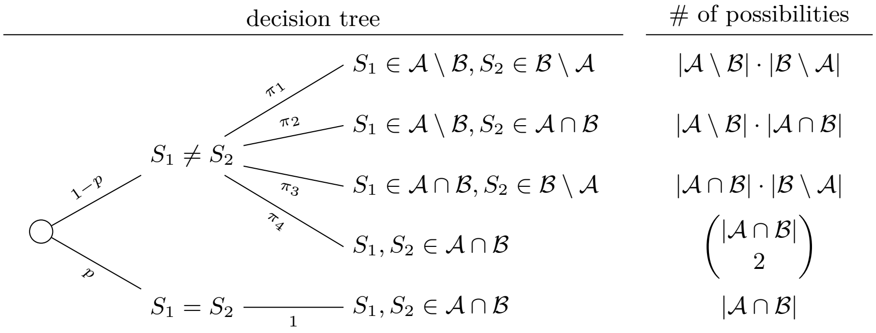
decision tree
# of possibilities
$1-p$
p
$S_1 \neq S_2$
$S_1 = S_2$
π_1
π_2
π_3
π_4
1
$S_1 \in \mathcal{A} \setminus \mathcal{B}, S_2 \in \mathcal{B} \setminus \mathcal{A}$
$S_1 \in \mathcal{A} \setminus \mathcal{B}, S_2 \in \mathcal{A} \cap \mathcal{B}$
$S_1 \in \mathcal{A} \cap \mathcal{B}, S_2 \in \mathcal{B} \setminus \mathcal{A}$
$S_1, S_2 \in \mathcal{A} \cap \mathcal{B}$
$S_1, S_2 \in \mathcal{A} \cap \mathcal{B}$
$|\mathcal{A} \setminus \mathcal{B}| \cdot |\mathcal{B} \setminus \mathcal{A}|$
$|\mathcal{A} \setminus \mathcal{B}| \cdot |\mathcal{A} \cap \mathcal{B}|$
$|\mathcal{A} \cap \mathcal{B}| \cdot |\mathcal{B} \setminus \mathcal{A}|$
$\binom{|\mathcal{A} \cap \mathcal{B}|}{2}$
$|\mathcal{A} \cap \mathcal{B}|$

Fig. 5. Decision tree for probabilistic intersection attack with two events.

Generalisations of the probabilistic intersection attack to more than two events are up to further research.

9 Conclusions

Due to recent changes in the legislation that now require the retention of identifying information, anonymity services face the challenge to resist a new kind of adversary. Such adversaries can force the anonymity service to collaborate to a certain extent, which is defined by law. The intention of our study is to assess the risk which arises from adversaries that are mounting intersection attacks on retained data of anonymity systems. We have measured the remaining anonymity from real data on user behaviour, which we believe is representative for information that can be requested by a law enforcement agency. The results indicate that hiding in an anonymity set works well as long as adversaries pose single request at distinct points in time without relating several requests to each other. However, the results also show that an adversary who combines the results of different requests, and therefore requests several anonymity sets in order to intersect them, has much more success in narrowing down individuals in the anonymity set. Compared to a single request, the intersection of only two requests reduces the size by far more than 50%. Though this is hardly sufficient for law enforcement agencies that seek to reduce anonymity sets to single persons, the results can be further refined, presumably with similar success, by intersecting more anonymity sets that are known to contain the target person.

Our results show that there is a remarkable difference with regard to the size of the remaining anonymity set between different ways of requesting data. The anonymity sets are larger if the set of those users is requested who were *logged in* in a distinct moment in time. The anonymity sets decrease if only *active* users are requested. Presumably the anonymity sets are even smaller, if the requests do not concern the application layer of the anonymity service, but the underlying network layer. Our study, however, is limited to the application layer.

Even though this discussion may lead to the conclusion that it is necessarily desirable for an adversary or a law enforcement agency to request user sets of active users only, this idea may be misguiding for anonymity systems such as AN.ON: users may send dummy traffic as a countermeasure. The idea behind dummy traffic of users is to make themselves appear active, even though they are actually idle. This can be achieved by regularly sending data packages from the user to the service without any content of interest. It is indeed crucial that besides the user (and, in certain constructions, the anonymity service), nobody is able to distinguish dummy traffic from ordinary traffic. Thus, if users send dummy traffic, a law enforcement agency which is able to obtain the set of all *active* users would not learn more than an agency which is limited such that it can only observe the set of all users that are *logged in.*

Dummy traffic has been discussed with regard to several attack schemes [13–15]. In general, it has been found to be a rather weak countermeasure in packet-switched networks. However, due to the specific limitations of the "adversary" defined by the data retention act, a continuous connection to the anonymity service together with weak dummy traffic seems a strikingly good solution. The economical aspects of dummy traffic have been mentioned in literature, but might be of decreasing significance in a world with complete network coverage and flat rates.

Acknowledgements

This paper has benefitted from thoughtful comments by Andreas Pfitzmann.

The research leading to these results has received funding from the European Community's Sixth Framework Programme (FP6/2002–2006) under grant agreement no. 507512 and the Seventh Framework Programme (FP7/2007–2013) under grant agreement no. 216483.

References

1. Claessens, J., Díaz, C., Goemans, C., Dumortier, J., Preneel, B., Vandewalle, J., Dumotier, J.: Revocable anonymous access to the Internet? Internet Research: Electronic Networking Applications and Policy **13**(4) (2003) 242–258
2. Köpsell, S., Wendolsky, R., Federrath, H.: Revocable anonymity. In Müller, G., ed.: Proc. of ETRICS. LNCS 3995, Berlin Heidelberg, Springer Verlag (2006) 206–220
3. Danezis, G., Diaz, C.: A survey of anonymous communication channels. Technical Report MSR-TR-2008-35, Microsoft Research (2008)
4. Bundestag: Gesetz zur Neuregelung der Telekommunikationsüberwachung und anderer verdeckter Ermittlungsmaßnahmen sowie zur Umsetzung der Richtlinie 2006/24/EG vom 21. Dezember 2007 [Amending act to the German Telecommunications Act]. Bundesgesetzblatt (Teil I, Nr. 70) (December 2007) 3198–3211 ausgegeben zu Bonn.
5. Bundestag: Telekommunikationsgesetz vom 22. Juni 2004 (2007) BGBl. I S. 1190, zuletzt geändert durch Artikel 2 des Gesetzes vom 21. Dezember 2007 (BGBl. I S. 3198).

6. Pimenidis, L., Kosta, E.: The impact of the retention of traffic and location data on the internet user. DuD Datenschutz und Datensicherheit **32**(2) (February 2008) 92–97
7. Pfitzmann, A., Hansen, M.: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management – A consolidated proposal for terminology. `http://dud.inf.tu-dresden.de/Anon_Terminology.shtml` (2008) (Version 0.31e).
8. Díaz, C., Seys, S., Claessens, J., Preneel, B.: Towards measuring anonymity. In: Designing Privacy Enhancing Technologies, Proceedings of PET'02, Springer-Verlag, LNCS 2482 (2002) 54–68
9. Serjantov, A., Danezis, G.: Towards an information theoretic metric for anonymity. In: Designing Privacy Enhancing Technologies, Proceedings of PET'02, Springer-Verlag, LNCS 2482 (2002) 41–53
10. Berthold, O.: Effiziente Realisierung von Dummy Traffic zur Gewährleistung von Unbeobachtbarkeit im Internet [An efficient implementation of dummy traffic to ensure unobservability on the Internet]. Diploma thesis, Technische Universität Dresden, Faculty of Computer Science, Institute for Theoretical Computer Science (1999) in German.
11. Danezis, G.: Statistical disclosure attacks: Traffic confirmation in open environments. In Gritzalis, S., di Vimercati, S.D.C., Samarati, P., Katsikas, G., eds.: Proceedings of Security and Privacy in the Age of Uncertainty. (May 2003)
12. Mathewson, N., Dingledine, R.: Practical traffic analysis: Extending and resisting statistical disclosure. In Martin, D., Serjantov, A., eds.: Proceedings of the 4th International Workshop on Privacy-Enhancing Technologies. Volume 3424 of LNCS., Springer Berlin/Heidelberg (May 2004) 17–34
13. Berthold, O., Langos, H.: Dummy traffic against long term intersection attacks. In Dingledine, R., Syverson, P., eds.: Proceedings of Privacy Enhancing Technologies workshop (PET 2002), Springer-Verlag, LNCS 2482 (April 2002)
14. Díaz, C., Preneel, B.: Reasoning about the anonymity provided by pool mixes that generate dummy traffic. In: Proceedings of 6th Information Hiding Workshop (IH 2004). LNCS, Toronto (May 2004)
15. Díaz, C., Preneel, B.: Taxonomy of mixes and dummy traffic. In: Proceedings of I-NetSec04: 3rd Working Conference on Privacy and Anonymity in Networked and Distributed Systems, Toulouse, France (August 2004)