THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

"I am Spartacus"

Citation for published version:

Kwecka, Z, Buchanan, W, Schafer, B & Rauhofer, J 2014, "I am Spartacus": Privacy enhancing technologies, collaborative obfuscation and privacy as a public good', *Artificial Intelligence and Law*, vol. 22, no. 2, pp. 113-39. <https://doi.org/10.1007/s10506-014-9155-5>

Digital Object Identifier (DOI):

[10.1007/s10506-014-9155-5](https://doi.org/10.1007/s10506-014-9155-5)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Artificial Intelligence and Law

Publisher Rights Statement:

© Kwecka, Z., Buchanan, W., Schafer, B., & Rauhofer, J. (2014). "I am Spartacus": Privacy enhancing technologies, collaborative obfuscation and privacy as a public good. *Artificial Intelligence*, 1-27. [10.1007/s10506-014-9155-5](https://doi.org/10.1007/s10506-014-9155-5)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

“I am Spartacus” – Privacy Enhancing Technologies and Privacy as a Public Good.

Introduction

This paper discusses the technology for an obfuscation based method for privacy enhancing tools, which serves the dual function of protecting the reputation and privacy of data subjects, while at the same time protecting legitimate police interests in the confidentiality of an investigation. Unlike most other approaches to PET, in our model for an “Investigative Data Acquisition Platform” the protection of privacy is seen as a communal task, something that we call for reasons that will become apparent below the “Spartacus model” of data protection. This approach requires us however to reconsider

not just the doctrinal legal environment of privacy and data protection law, but also its jurisprudential, ethical and sociological underpinnings. Most approaches to PET reflect the individualistic, libertarian origins of privacy law as an individual right. By contrast, our approach asks how PET can look like in a society that considers privacy a common good and the protection of privacy a communal task, an understanding of privacy law that has recently gained much ground in the academic debate. In the first part of the paper, we therefore describe the motivation for this approach in the form of an extended use case, which allows us to give an informal outline of the solution suggested here. This will prepare the ground for a legal-jurisprudential analysis that is needed for the normative underpinning of the technology. In the second part, we introduce the formal apparatus that supports this type of communal privacy protection. In the third part, we provide a short evaluation of the results, both from a technological and from a legal and ethical perspective, indicating also a number of necessary further research questions, in particular empirical and socio-legal questions regarding common perception of privacy and risks.

1.1 Setting the scene: Obfuscation and privacy protection

Consider as a setting for the discussion in this paper the following example of a traditional, brick and mortar police investigation: The police wants to check the alibi of a suspect, John Doe. They drive in a marked police car to his place of residence, park it in full view on the street next to his house, and then send pairs of uniformed police officers from neighbour to neighbour, asking if they saw Mr Doe at his home during a certain time interval.

This sort of scenario carries two obvious risks: one is a reputation risk for Mr Doe. His neighbours now know at the very least that he is for one reason or other suspected of a wrongdoing. They might also be able to infer from the question some of the information the police holds about Mr Doe – if for instance they are asked for a specific time interval, and it is well known that during that time a robbery happened nearby, it would be obvious to infer that Mr Doe is suspect in a robbery. If the question is: “Have you ever seen very

young girls visiting your neighbour late at night”, another inference would immediately be drawn.

At the same time, this approach also carries risks for the police – the neighbours may inadvertently or intentionally alert Mr Doe that he is subject to a police investigation. One way to protect both the interest of Mr Doe in preventing the disclosure of information about him (here, that he is subject to a police inquiry) and the interest of the police not to alert Mr Doe is to ask questions that are much broader phrased. This could be e.g. asking every person on that street, including Mr Doe, to list everybody whom they saw in the neighbourhood at the relevant time. This way, there is no finger of suspicion that points at one specific person. But this strategy carries obvious costs too. It creates much more information than necessary, most of it noise, which the police then has to process. It also creates a privacy risk for a much larger number of people – the police now knows about the whereabouts of a large number of citizens it has no legitimate interests in. Nonetheless, creating an excess amount of information seems, paradoxically, to be one way of protecting privacy and integrity of the investigation

We can now transfer this scenario to the internet, for instance a request to an ISP for data that establishes when a suspect was online, or a request to a bank about online transactions carried out by a client. We assume here and in the following that the data was legitimately held by that company, either because it is necessary to fulfil its contractual obligation with Mr Doe, or because there is statutory data retention duty imposed on them. We also assume that the police warrant is legitimate and necessary. At this point, we face the same dilemma as described above – the formal request for information discloses to the data controller that the police has a legitimate interest in one of their clients. This in turn might give the data controller an incentive for action. A bank for instance may decide to disassociate themselves from a client who has been frequently the subject of data discloser requests, on the assumption that he carries a risk for reputation damage should he become subject of a high profile trial. This in turn may alert the client to investigative activities against him. Concerns about cybercrime

and terrorism has resulted in a considerable extension of data retention duties by internet service providers, which add to the already significant data online retailers, banks or social media providers hold about their clients. For the purpose of criminal investigations, channels have been created that allow the police under certain circumstances, defined in law, to demand access. In the aftermath of 2001, public authorities were granted much wider rights to gather operational data (Swire & Steinfeld, 2002; Young, Kathleen, Joshua, & Meredith, 2006). For a number of years public opinion accepted privacy intrusions as the sacrifice everybody must make. However, slowly the public opinion is shifting back to the state where intrusion of privacy is considered as unacceptable. This is shown by different surveys such as the one conducted by the Washington Post in 2006 (Balz & Deane, 2006), where 32% of respondents agreed that they would prefer federal government to ensure that privacy rights are respected rather than to investigate possible terrorist threats. This was 11% increase from the similar survey conducted in 2003. However, while this indicates a general societal willingness “to do something” about privacy during police investigations, the “obfuscation” method described above – asking much wider, less focussed questions – can’t easily be transferred to an online environment. The formal procedure that is required to gain data access requires that the query is sufficiently precise and focussed, to prevent fishing expeditions and unnecessary privacy intrusions of innocent citizens. In Europe, the Data Protection Directive allows national police forces access to data only “in specific cases.” As Bignami (2007) noted, this provision is explicitly designed to prohibit high-tech fishing expeditions, whether done by the police or by market actors. Again Bignami:

“The police cannot make blanket requests for calling information. Rather, they must compile detailed requests for information on specific telephone numbers. The requirement of specificity is a means of guaranteeing that the police have at least some grounds for suspecting those telephone numbers of being involved in a criminal conspiracy.”

Paradoxically therefore, a method that could in principle protect citizens from the misuse of their data prohibits certain privacy enhancing methods.

Nonetheless, using obfuscation is an attractive privacy enhancing tool in principle. In the online scenario, it is the protection of the interests of third parties that prevents us to hide the identity of Mr Doe behind a veil of “excess data”. This however would change if a sufficient number of other clients of the company in question waived their rights, and under the assumption of mutuality and reciprocity provide the “fog” of data that shields the identity of the subject of a data query from the data controller, though not the police. The bank or ISP will in this model only know that the subject of the query is amongst the arbitrarily large number of records they are asked to hand over to the police. The police in turn must only be able to read amongst all the data handed over to them the data of the person they are interested in. We will see below how a combination of a trusted third party approach together with encryption methods can provide just such a set up.

A particularly intuitive example of such a solidarity based protection of Identity against a data query though comes from the film “Spartacus”. In one of the most climatic scenes of the film, a Roman general demands from the captured remains of the former slave army that they turn Spartacus over to him. To protect his friends, Spartacus stands up to say “I am Spartacus.” However, the solidarity of his soldiers is so great that several of them come forward, shouting “I am Spartacus!” until the shouts dissolve into a cacophony of thousands of former slaves each claiming “I am Spartacus!”. This makes it impossible for the general to identify and arrest Spartacus. This story also points to one of the main issues that technology alone cannot tackle – the legal and social environment necessary for such an approach to work. Enlightened self-interest plays a role, and our model will assume reciprocity: I’m willing to accept a (ideally very low) privacy risk to myself when making my data available as “fog”, but I know that should I ever be at the centre of an investigation, others will do the same for me. As the Spartacus example shows, people are sometimes willing to take personal risks for a communal good. This requires us however to reconsider the normative foundations of privacy law. Using obfuscation as a means to protect privacy is by no means new – and other writers have made the connection to the film too. (e.g. Howe and Nissenbaum 2009). However, as Brunton and Nissenbaum note, most of

these approaches still put the burden to produce the excess data on the individual who wants to protect herself. This is a situation very different from the one encountered in *Spartacus*, and indeed our proposed solution. The few examples of collective obfuscation that they identify are typically “low tech”, e.g. swapping of loyalty cards, do not involve any risk for the collaborators and are directed against illegitimate privacy intrusions by private companies. Our problem, and hence our solution, differs in all these aspects. First, existing methods of collaborative obfuscation are low technology approaches by grassroots activists trying to undermine corporate data mining in the long run, which would not work in the type of scenario we discuss, a sophisticated online investigation for a single, specific event. Furthermore, in our scenario there is a legitimate police investigation, and whatever method we chose to protect Doe’s privacy interests, they must not interfere with the legitimate exercise of police functions. Indeed, as we indicated above, protecting Doe’s privacy is in the interest of *both* him and the police – an approach which we hope will help to revise the often overly simplistic concept of privacy as an irreconcilable conflict between police and individuals. Finally, it is worth remembering the outcome of *Spartacus*, the movie. Unable to identify Spartacus, General Crassus crucifies all of the slaves. In our approach too, and in marked difference to previous approaches to collective obfuscation, people will be asked to expose themselves to a – very limited – risk, something necessitated by the specific characteristics of our scenario. Because of this not inconsiderable demand we make on other users, it is necessary to spend a bit more time on the philosophical and jurisprudential underpinnings of our approach, and generally the nature of privacy, to legitimate and put into context this demand.

1.2 Privacy as a public good and a public responsibility

Privacy has traditionally been framed in law as a paradigmatic case of an individual right that pitches the self-interest of individual against the communal interest of the state. This is a feature it shares with a traditional understanding of human rights law in general, as individual rights that protect against state action only. To a degree, we can read this even from the etymology of the

word. “Privacy” is derived from the Latin “privare” – to rob or to deprive. Our private time was for the ever egregious Romans the time when we robbed our friends from of pleasure of our company, and the time we deprived the state of our service. This perception remains to a degree with us today. Privacy is in public discourse often portrayed as ultimately selfish, and in the age of social media if not anti-social, then at least a-social. Only recently, an alternative discourse in human rights scholarship has emerged, which portrays privacy itself as a social or public value on which other important public goods, in particular democracy and public participation rests. Privacy enables individuals to criticise and resist measures or acts of government that are of an undemocratic or even totalitarian nature. It has therefore been suggested that privacy is necessary to protect individuals from the pressure to conform to societal expectations in a way that poses a threat not only to human dignity and a person’s individuality, but also to the liberty that flows from it. Equally, Simitis (1984 p.399) argued forcefully that even though privacy has often been misunderstood as conflicting with transparency, free speech and other democracy enhancing concepts, its role in fostering participation must not be overlooked. Even earlier, Bloustein (1964 p.1003) argued that

“[t]he man who is compelled to live every minute of his life among others and whose every need, thought, desire, fancy or gratification is subject to public scrutiny, has been deprived of his individuality and human dignity. Such an individual merges with the mass. His opinions, being public, tend never to be different; his aspirations, being known, tend always to be conventionally accepted ones; his feelings, being openly exhibited, tend to lose their quality of unique personal warmth and to become the feelings of every man.

And indeed, the experience in many totalitarian regimes has shown that an absence of privacy has the potential for creating a “society of followers”.

This interdependency between the protection of privacy and the protection of other essential features of a democratic society is also highlighted by Raab (2011) who argues that values like personal autonomy and self-determination

“are important not primarily because individuals may wish to live in isolation (for they do not, mostly), but so that they can participate in social and political relationships at various levels of scale, and so that they can undertake projects and pursue their own goals”.

While this shift towards the recognition of privacy as a public good is welcome, for our purpose it has the problem that much of the reassessment also resulted in questioning the role of consent and privacy waivers. As long as privacy was seen merely as an individual right, governments found it easy to convince individuals of the legitimacy of a privacy-security trade-off. Similarly, free social media services such as Facebook offer essentially a “trade in” between privacy and free use of services, paid for by advertising revenue. This turned privacy into a tradable object under the control of the rights holder, and marginalised the concept of “privacy risk”. How valuable is privacy if so many people are willing to trade it in for mere pennies in discounts when shopping with a loyalty card, or hit points in an online game? Just as privacy became in this model private property, so did privacy risks, which were conceptualised as only one of a number of competing risks and benefits such as fear of crime, loss of convenience¹ or loss of material gain. Theories that emphasise the value of privacy as a common good therefore also became sceptical of the notion of free alienation of privacy in market places, and with that the role of individual consent. As Regan (1995 p. 233) argues, there is a risk that

“[i]f one individual or a group of individuals waives privacy rights, the level of privacy for all individuals decreases because the value of privacy [in the collective view of society] decreases”.

Or put differently, in a society where “Big Brother” is daytime television and everybody shares their feelings on Facebook or Twitter, refusing to participate in the sharing of data is at best mildly odd, at worst in itself suspicious. In an administrative or law enforcement context, this means that an already existing

“information imbalance” between citizens and the state is further shifted in the state’s favour. We have to be careful therefore that our own approach and its use of “consent”, does not inadvertently undermine this very notion of solidarity which is crucially dependent on the legal conceptualisation of privacy as a common good. Technology, institutional arrangements and law all play an important role in this balancing act.

These preliminary jurisprudential reflections provide us with an abstract normative framework for the technological solution to protect the privacy of people caught for whatever reasons on the police radar during an investigation. It assumes that the protection of privacy is not just a task for the individual, but a communal concern, based on solidarity and not (just) self-interest. The aim is a solution where through solidarity in a community, the identity of a suspect is protected, without interfering with legitimate police interests. This requires reassurances, technological, institutional and legal, for those people who are willing to assist in the protection of each other’s privacy. In the next section, we introduce our proposal for a “Data Acquisition Platform” (IDAP), focussing mainly on the first aspect, how the necessary trust can be created that allows actions of solidarity.

2 Introducing IDAP.²

2.1 BACKGROUND AND RELATED WORK

Leaving the investigative context aside, the retrieval of information from a third-party in a private manner is a generic problem that has been researched for use in a variety of different scenarios such as cooperative scientific computation (Du and Atallah 2001); and on-line auctions (Cachin, 1999). The things people search for disclose potentially a lot about them. This is after all the central part of Google’s business model – online behavioural profiling based on search queries allows the targeting of advertising with a high degree of accuracy (Tene 2008). More and more often, analysing search queries by suspects also play a role in criminal investigations, establishing motives, methods and state of mind (Lawless 2007). Initially, Private Information

² A more detailed description of the technical aspects of IDAP can be found in the 2011 PhD thesis of one of the authors, Kwecka, Cryptographic privacy-preserving enhancement method for investigative data acquisition. <http://researchrepository.napier.ac.uk/4437/>

Retrieval (PIR) protocols were designed with a basic requirement of acquiring an interesting data record, or just a specific data bit, from a dataholder, *sender*, in a way that this dataholder is unable to judge which record is of interest to the requestor, *chooser*. These protocols were not concerned with the secrecy of the records stored in the database, thus in its least optimised state a PIR could have been achieved by transferring the whole database from the *sender* to the *chooser*, as this would allow the *chooser* to retrieve a record in a private manner. To use a very simple analogy, if an individual wants to browse the offerings of an online retailer of medical self-help books, but does not want to leave a trail that indicates to the retailer unnecessarily which illness he may suffer from, downloading the catalogue in pdf and searching it in the privacy of his own home has advantages over online browsing. There are no privacy concerns on the side of the retailer in this case, as all the information is public anyway. Consequently, the main motivation behind the PIR schemes is achieving minimal communicational and computational complexity (Ostrovsky & William E. Skeith III, 2007). A stronger notion than PIR is *1-out-of-n* Oblivious Transfer (OT) primitive that allows the retrieval of a randomly selected record from the dataset of n elements held by the *sender* in a way that the *sender* cannot learn which record has been transferred, and the *chooser* cannot learn anything about other records in the dataset (Schneier, 1995). *1-out-of-n* OT protocols that allow *chooser* to actively select a record to be retrieved, and that have linear or sub-linear complexity, can be referred to as symmetric PIR (SPIR) protocols, since they protect the records of both parties during the information retrieval. These useful privacy-preserving data retrieval protocols can be employed in a variety of systems: electronic watch-lists of suspects (Frikken & Atallah, 2003); cooperative scientific computation (Du & Atallah, 2001; Goldwasser & Lindell, 2002); and on-line auctions (Cachin, 1999). Frikken's and Atallah's approach deserves some further comments, as it shares some of the technological solutions with our proposal, but due to a very different legal-ethical approach to privacy advocates an implementation thatacerbates rather than reduces the specific issue we want to address. A typical application for their solution is the following: the police have received

information that some known suspects are planning a bomb attack, possibly using fertiliser. They want to query the database of a fertiliser retailer, ideally without alerting the retailer on the identity of their suspects. This can have several reasons, including suspicions against the retailer himself. As indicated above, one obvious solution would be to simply request the entire database, or data about everybody who bought fertiliser, and analyse it on a police server. However, this would mean that the police also gets hold of data about a large number of innocent citizens – with the fear that they might e.g. re-use the data to check it against tax records or other types of investigative activity which is perceived by the population as marginal in comparison to terrorism. Frikken's and Atallah's solution is to provide the police not with the entire database of the retailer, but a segment of it that is sufficiently large to hide their interest in a specific person from the retailer. To protect the wider public though, the selection of data is determined by an objective criterion such as a list of people with previous criminal records, possibly for related offences. This minimises the privacy risk for innocent citizens. It does however potentially increase the privacy risk for people on the lists from which selection takes place substantially. The retailer could in this case learn that a number of his customers have previous records, or have come to the attention of the police in some other way. We can see now the different jurisprudential assumptions behind their model and ours: In Frikken and Atallah, privacy is a conditional right that can be lost through misbehaviour. This does not just apply to the suspect in an investigation – who can reasonably suffer restrictions in his privacy to further the aims of the criminal justice system. Rather, once convicted of a criminal offence, the offender suffers in perpetuity reduced privacy rights, even in cases that have nothing to do with him and only accidentally concern him. Where our model is based on the voluntary solidarity between all citizens (whether or not they have a previous record, or are on a police watch list), in their model a subset of the citizenry, those who for one reason or the other have already become subject to police interest, are forced to provide the cover for the investigators.

With the use of the protocols described above, a *chooser* would be capable of privately retrieving a record from the *sender's* database, by secretly referring to its index in this database. In SPIR such index is expected to be publically available in an electronic catalogue or a directory (Aiello, Ishai, & Reingold, 2001; Bao & Deng, 2001). However, ISPs and other dataholders with large databases of private data cannot be expected to maintain such freely available indexes. Also, it is expected that an investigator would normally refer to a suspect by name, ID or phone number, etc. For this reason before the data can be received using SPIR, a search would need to be performed by the *chooser* against the records in the *sender's* database. Such a private search operation requires a protocol that allows two parties to compare the values of their data in a private manner. The protocols that are optimised to make comparisons for equality are referred to as Private Equality Test (PEqT) protocols. PEqT protocols are often based on commutative (Frikkien & Atallah, 2003; Kwecka et. al. 2008) or homomorphic cryptosystems (Bao & Deng, 2001).

An interesting record can be located in a database using a *1-out-of-n* PEqT protocol and then retrieved with help of SPIR. Often each of these protocols would have a separate computationally expensive preparation phases, such solution would not be optimal for IDAP. The exception to this rule is a range of protocols including: private intersection; private intersection size; and Private Equijoin (PE) defined in (Agrawal, Evfimievski, & Srikant, 2003). These protocols are based on commutative encryption and thanks to the use of different properties of the underlying commutative algorithms are capable of allowing for both private matching and private data retrieval.

2.2 BUILDING BLOCKS

This section describes the PE protocol that is the basis for the creation of the privacy preserving investigative platform - IDAP. The PE protocol relies on commutative cryptography, thus some background for this is provided first.

2.2.1 Commutative Cryptosystems

Many cryptographic applications employ sequential encryption and decryption operations under one or more underlying cryptosystems. The reasons to sequence (cascade) different cryptographic schemes together include strengthening the resulting ciphertext and achieving additional functionality which is impossible under any given encryption scheme on its own (Shannon, 1949; Weis, 2006). A basic cascable cryptosystem can consist of a number of encryption stages, where the output from one stage is treated as an input to another. In such a basic cascable cryptosystem it is necessary to decrypt in the reverse order of encryption operations. However, a special class of sequential cryptosystems - commutative cryptosystems – allows for the decryption of a ciphertext in an arbitrary order. Thus, a ciphertext $c = e_b e_a(m)$ (c – ciphertext, m – plaintext, e – encryption operation under keys a and b), could be decrypted as either $m = d_b d_a(c)$ or as $m = d_a d_b(c)$. The advantages of such cryptosystems were widely promoted by Shamir (1980) as used in his, Rivest's and Aldman's classic game of *mental poker*, employing the Three-Pass (3Pass) secret exchange protocol.

The most commonly used commutative cryptosystem is based on the Pohlig-Hellman (PH), asymmetric private key scheme (1978). This scheme first published in 1978 has never become popular since it is asymmetric, and therefore slow in comparison to other private key systems. While the PH protocol influenced the design of Rivest-Shamir-Adleman (RSA) public key scheme (1978), the main strength of PH is that it is commutative for keys based on the same prime number and that it allows for comparing the encrypted ciphertexts. Consequently, under PH the two ciphertext $c_{ba} = e_b e_a(m)$ and $c_{ab} = e_a e_b(m)$ hiding the same plaintext m are equal (1), while this is not the case with ordinary encryption protocols, that satisfy (2).

$$e_a e_b(m) = e_b e_a(m) \quad (1)$$

$$e_a e_b(m) \neq e_b e_a(m) \quad (2)$$

Thanks to those properties PH can be used in the 3Pass primitive that allows two parties to exchange data without exchange of keys, as well as to perform PEqT that permits private matching of data records.

2.2.2 Three Pass Protocol (3Pass)

The 3Pass protocol, shown in Fig.1, was intended to allow two parties to share a secret without exchanging any private or public key.

Fig. 1 Three-Pass Secret Exchange Protocol.

The protocol was aimed at providing an alternative to public-key encryption and DH-like key negotiation protocols.

The operation of the protocol can be described using the following physical analogy:

1. Alice places a secret message m in a box and locks it with a padlock E_A .
2. The box is sent to Bob, who adds his padlock E_B to the latch, and sends the box back.
3. Alice removes her padlock and passes the box back to Bob.
4. Bob removes his padlock, and this enables him to read the message from inside the box.

There could be more parties, or encryption stages, involved in a 3Pass-like protocol, and this property makes it ideal for locking a plaintext multiple times and then unlocking it in an arbitrary order, as long as the parties are cooperating until the execution of the protocol is completed. Such functionality is required by IDAP as described later in this paper.

2.3 Private Equality Test (PEqT)

PEqT protocols can be used to privately verify whether two secret inputs to the protocol are equal or not. Agrawal, Evfimievski and Srikant (2003)

proposed one of the most scalable and flexible PEqT protocols for operations on datasets. The scheme is illustrated in Fig. 2 and can be described in the following steps:

1. Alice encrypts her input and sends it to Bob.
2. Bob encrypts the ciphertext received from Alice and sends it back.
3. Bob encrypts his secret input and sends it to Alice.
4. Alice encrypts the ciphertext containing Bob's input.
5. Alice compares the two resulting ciphertexts, if they are equal then her and Bob's inputs are equal.
6. Alice may inform Bob about the result.

Fig. 2 Private Equality Test.

This protocol allows two parties to compare their secret inputs.

The following section describes a scheme that extends both the PEqT and 3Pass primitives to form the PE protocol that is the blueprint for our IDAP.

2.4 Private Equijoin Protocol

A PE protocol can enable two parties, the *chooser* and the *sender*, to privately compare their sets of unique values V_C and V_S , and allows the chooser to retrieve some extra information $ext(v)$ about records V_S , that match records V_C on a given parameter. Thus, sensitive data marked as V_C and V_S , such as date of birth, address or credit-card number, describing the data subjects in two datasets can be compared in their encrypted forms using the PEqT primitive, in order to find the equijoin between the two datasets. The equijoin shows where the list of the items requested match the lists of the items in the

dataset and nothing else. Then the PE uses the 3Pass primitive to reveal the information that the *sender* wants to make available to the *chooser*, the $ext(v)$, for the items in equijoin only. However, the sender is “blind” at this stage, as s/he does not know the records that are in the equijoin. Consequently, the investigators could encrypt their list of the suspects V_C and receive data, $ext(v)$ on the individuals matching the criteria in the encrypted set V_S . Please note that v stands for a single record/data-subject in dataset V_C or V_S . Thus the uppercase letters refer to sets. The PE protocol involves the following steps:

1. Both parties apply hash function h to the elements in their sets, so that $X_C = h(V_C)$ and $X_S = h(V_S)$. *Chooser* picks a secret PH key E_C at random, and *sender* picks two PH keys E_S and E'_S , all from the same group Z_p^* .
2. *Chooser* encrypts entries in the set: $Y_C = E_C(X_C) = E_C(h(V_C))$.
3. *Chooser* sends to *sender* set Y_C , reordered lexicographically.
4. *Sender* encrypts each entry $y \in Y_C$, received from the *chooser*, with both E_S and E'_S and for each returns 3-tuple $\langle y, E_S(y), E'_S(y) \rangle$.
5. For each $h(v) \in X_S$, *sender* does the following:
 - (a) Encrypts $h(v)$ with E_S for use in equality test.
 - (b) Encrypts $h(v)$ with E'_S for use as a key to lock the extra information about v , $\kappa(v) = E'_S(h(v))$.
 - (c) Encrypts the extra information $ext(v)$:
$$c(v) = K(\kappa(v), ext(v))$$

Where K is a symmetric encryption function and $\kappa(v)$ is the key crafted in Stage 5b.
 - (d) Forms a pair $\langle E_S(h(v)), c(v) \rangle$. These pairs, containing a private match element and the encrypted extra information about record v , are then transferred to *chooser*.
6. *Chooser* removes her encryption E_C from all entries in the 3-tuples received in Step 4 obtaining tuples α , β , and γ such that $\langle \alpha, \beta, \gamma \rangle = \langle h(v), E_S(h(v)), E'_S(h(v)) \rangle$. Thus, α is the hashed value $v \in V_C$, β is the

hashed value v encrypted using E_s , and γ is the hashed value v encrypted using E'_s .

7. *Chooser* sets aside all pairs received in Step 5, whose first entry is equal to one of the β tuples obtained in Step 6. Then using the γ tuples as symmetric keys it decrypts the extra information contained in the second entry in the pair $\langle E_s(h(v)), c(v) \rangle$.

The above protocol can perform the basic functions required for the purpose of investigative data acquisition. Its use in investigative scenarios is described in the following section.

3 . IDAP VS. PRIVATE EQUIJOIN

This section evaluates our proposed use of the PE protocol as basis for IDAP. The operations required during investigative data acquisition from a third party in general consist of:

1. Identification of the type of the information that is required. These could be h parameters that contain answers to investigator's questions, referred to as return parameters rp_{1-k} , e.g. Date of Birth (DOB), address, location of a card payment, or numbers called by a given subscriber. In a formal, legally prescribed environment, it ought to be able to demonstrate later that these criteria matched those on the warrant application, adding an additional level of legal scrutiny and accountability.
2. Specification of any circumstantial request constrains, or l different input parameters, ip_{1-l} , with values ip_val_{1-l} , e.g. time frame of the transactions being requested.
3. Specification of the relevant data subject e.g. by identifying the individual whose data is to be retrieved, or by providing the mobile phone number of the suspect, etc. This parameter is referred to as the record of the interest, ri with value ri_val .
4. Retrieval of the relevant records.

Then, if we refer to the dataset as the *source*, the request for investigative data could be mapped into the following SQL query:

```
SELECT  $rp_1, rp_2, \dots, rp_h$   
FROM source  
WHERE  $ri=ri\_val$  AND  $ip_1=ip\_val_1$  AND  $ip_2=ip\_val_2$  AND ... AND  $ip_l=ip\_val_l$ 
```

In most cases the names of the return parameters, as well as the names of the input parameters, and values of these input parameters can be openly communicated. But the value of the interesting record, *ri_val* is used to uniquely identify the suspect and must be hidden. This can be achieved by running a database query for the return parameters of all the records that satisfy the conditions defined by the input parameters and then collecting the interesting record from the sender using a PE protocol. Consequently, the query that is actually run on the sender's database can be rewritten to:

```
SELECT  $ri, rp_1, rp_2, \dots, rp_h$   
FROM source  
WHERE  $ip_1=ip\_val_1$  AND  $ip_2=ip\_val_2$  AND ... AND  $ip_l=ip\_val_l$ 
```

The results of such query (3) would be an input to a PE that would enable the chooser to privately select only the record of interest that match given *ri_val*.

3.1 PE's Performance

The previous section discussed different types of protocols available that could enable the *chooser* to download a record from the *sender's* database, maintaining the secrecy of the record selected. We also mentioned that most available protocols could not achieve IDAP on their own, and a combination of two or more protocols is required. Such combination typically results in high computational and communicational complexity, because each protocol usually requires its own preparation phase. The PE protocol described in

Section IV is capable of both private matching and performing SPIR, and has a low overhead. Table I defines the computational complexity of the protocol.³

TABLE I
Computational Complexity of the PE protocol

	Symmetric Crypto.	Asymmetric Crypto.	
	crypto. operation	key generation	crypto. operation
Step 1	-	$O(3)$	-
Step 2	-	-	$O(m)$
Step 4	-	-	$O(2m)$
Step 5	$O(n)$	-	$O(2n)$
Step 6	-	-	$O(2m)$
Step 7	$O(m)$	-	-
Total Complexity	$O(n + m)$	$O(3)$	$O(5m + 2n)$
Cost (ms/operation)	0.33	7	30

The complexity of each of the steps in the proposed initial solution.

Where n is the number of the data rows in the source, and m is the number of interesting records. Cost is the measured average time in ms to perform given cryptographic operation from managed C# .NET code.

In practice this particular solution based on the PH cipher and implemented in C# .NET can process thousand records a minute, on average. The following

³ For the research purposes the PE protocol has been implemented on a desktop computer running Microsoft Windows XP Professional operating system with an AMD Turion 64 X2 Mobile 1.58GHz CPU, and 3GB of RAM. The implementation was based on the Bouncy Castle cryptographic API. MS SQL GUIDs acted as input to hashing protocols, while the produced hashes were used as an input to the asymmetric algorithms (as in the OT and PE protocols). The AES128 protocol was tested using a 1kB input (that is approx. 150 words of ASCII text) this is expected to be larger than necessary to simulate records returned by the dataholder (similar amounts of data are used in Iliev and S. W. Smith (2005) and *Cristofaro et al 2009*). Using the implementation the research team has confirmed that some of the experiments can be simulated based on the computational complexity and cost measured in millisecond for operation. The values for the cost presented in Table I are based on an average time for the execution of 1 million cryptographic operations of the given type.

section discusses the performance in context of investigation, and discusses issues that could possibly limit the usability of the solution presented.

3.2 Advantages of PE in data acquisition process

Following our general philosophy outlined in the first part, the PE protocol allows for acquiring more than one interesting record at the time, and adding more records to the enquiry increases the processing time by a negligible value (~151ms) per each extra interesting record in an enquiry. Use of the PE would also satisfy the condition that the dataholder remains in full control of data, and decides which data can be disclosed. This addresses several current legal concerns regarding whether or not the police should be given direct access to traffic data in particular, or, as in the present system, the data controller should remain in control of the data and can if necessary refuse the request and challenge its legitimacy in court. In the PE protocol each record is processed separately and there are no chances of the records being mixed up by the privacy-preserving process. Thanks to this fact unnecessary data of non-suspects could be discarded on reception by the authorities and still the encrypted interesting records received would form valid evidence for use in a court of law. The costs involved in building and deploying PE based IDAP are anticipated to be low since it is a software system and the architecture is based on a protocol that is in the public domain.

3.3 Limitations of PE in the data acquisition process

The processing time required for the protocol to run is the main drawback of the PE protocol. If there are a thousand records in the database it only takes approximately one minute for the complete run of the protocol, however, the processing time is linear to the number of records in a dataset and data acquisition from a database with five million records would take three and a half days to run on an ordinary PC. During an urgent enquiry, especially where there is a clear danger to life, the police can currently get access to relevant location data from a mobile network operator in less than half an hour. Such a result could not be expected of PE if the database has more

than thirty thousand records. Additionally, even if the data requested is relatively small in size, e.g. 100kB per record, the results from a database of five million records would be more than 500MB of data that would need to be transferred over the Internet. Clearly, there is a requirement for the PE to run on a subset of the sender's database rather than the whole database or another solution would need to be chosen. The first approach is described further below.

The more serious problem is that PE alone does not solve the issues that we discussed in the introduction. A PE based system would work reasonably well in those situations where the issue is merely the secure matching of a single value per record, e.g. an IP address, name or a credit card number. In some scenarios it may however be necessary to request records based on a number of secret input parameters. Consider a scenario where the police has a profile of a suspect based on a complex investigative hypothesis about a planned terrorist attack, for instance the simultaneous kidnapping of key politicians and other high profile targets. This investigative hypothesis provides the investigators with a number of search parameters even if they do not yet have a specific theory about any individual in the terrorist group. This could be for instance people who showed a particular interest in the diaries of possible targets, provided that they don't have a profile that makes such an interest plausible. When matching now this profile against people working in certain organizations, information about this investigative hypothesis could be deduced by those organisations, which may harm both the investigation and also entire groups of people that match the profile.. For example if the case being investigated has the potential to create public tension, disclosing through the search parameters that the suspect's profile matched individuals in a local minority could have serious consequences to the members of this minority if it is leaked to the press that the police interests are going in that direction. Alternatively, the data holder may learn more about current thinking of the police on how to make effective bombs, or who the police thinks possible targets could be. IDAP should be able to assist the police in such a scenario by hiding the search parameters from the data holder, thus some

modifications needs to be introduced to the protocol, which are proposed in the next section.

4. PROPOSED MODIFICATIONS

The previous section has listed the drawbacks of using PE in the pursuit of IDAP. Here these drawbacks are addressed by two correcting measures that modify the PE protocol for the specific purpose of investigative data acquisition.

4.1 Lowering Processing Time

Above we recommended minimising the processing time required for each run of the protocol in large databases, such as those belonging to ISPs and mobile telephone providers. Theoretically, in order to maintain privacy of the suspect, the *chooser* needs to request from the *sender* to process all the records in the database. Only this way no information about the interesting records is revealed. The correctness of this scheme can be proven under the requirements of the multiparty computation (Asonov & Freytag, 2003). In its current form the system would not be capable of processing any urgent requests due to the processing time required, and this would be a major drawback. This could be mitigated by limiting the numbers of records that need to be processed and then sent by the *sender* per enquiry. Privacy of the alleged suspect should be protected, but if the probability of the *sender* guessing the ID of the interesting record is for example 1:1000 and not 1: n , and the dataholder has no other information that could help infer any knowledge as to the identity of the suspect, then this research argues that the privacy of the suspect and the investigation is maintained. As we discussed above, also during traditional face-to-face investigations, *diffusion* is used - hiding the suspect's identity by asking open-ended questions about a larger group of individuals rather than about a single person. As we noted, this is a widely accepted technique which would however in a digitalised environment fall foul of the prohibition of fishing expeditions. We are therefore from a legal perspective required to balance various conflicting – and sometimes

converging - interests. The interests of the police in a speedy investigation converge with those of other data subjects that the police should only receive the minimal amount of data necessary – this points to a solution that limits the number of “camouflage records” that they receive. From the perspective of the suspect, it matters just how detrimental an inference would be drawn by the mere fact of being the suspect of a criminal investigation. Thirdly, the nature of the data is also relevant. In an investigation against a suspected paedophile for instance, even otherwise innocent behaviour like browsing catalogues for children wear can be indirect evidence for the police case. In this situation, were it to become public knowledge that someone is suspected by the police of paedophilia would be particularly severe on an innocent suspect. At the same time, the mere fact that someone was looking at clothing catalogues is not particularly sensitive data outside the context of such an investigation; after all, even Amazon’s recommender system will make use of it. Therefore, the customers of the online retailer who are asked to provide “camouflage” for our suspect do not risk anything personally, even if the data were compromised, as the fact that they too looked at clothing catalogues is in itself uninteresting. In this scenario, it seems reasonable to increase the number of foils, as the risk for each is negligible, but the privacy gain for the suspect considerable. However, if the data is sensitive or possibly embarrassing regardless of whether or not it is analysed in the context of an investigation, for instance information about buying Viagra, then the number of foils should be reduced to minimise the risk for them as third parties. In both situations, another parameter would be the speed with which the police needs the information. Our approach allows to “scale” the protection of both the suspect and that of the other customers, taking this type of legally required balancing as a starting point.

The problem is to decide on the technique of narrowing down the scope in a way that ensures the record of interest are among the results returned. If the list of the record identifiers is public, such as the list of the Internet Protocol (IP) addresses or telephone numbers served by a given network operator, then the *chooser* could simply select records to be processed at random

from such directory. However, in case such list is not publicly available it would be possible to split the PE protocol back into separate parts: PEqT; and OT, and an additional off-line preparation phase. This way the initial off-line phase could be run against the whole database, but the information retrieval would be performed against a smaller set of records. If as a number of records requested per each interesting record is defined as the diluting factor - α the protocol IDAP would be defined as follows:

Phase A - Preparation

1. *Sender* applies hash function h to the elements in the input set V_S , so that $X_S = h(V_S)$.
2. *Sender* picks a encryption PH key E_S at random from a group Z_p^* , where p is a strong prime.
3. *Sender* encrypts each $h(v) \in X_S$ with the key E_S , the result is a list of encrypted identities $Y_S = E_S(X_S) = E_S(h(V_S))$

If more records need to be added to the set these can be processed using steps 1 and 3, and then added to the list.

Phase B - PEqT

1. Following a request for data, *sender* provides *chooser* with a complete list of encrypted identities prepared during Phase A, reordered lexicographically.
2. *Chooser* applies hash function h to the elements in set containing the identities of the interesting records, so that $X_C = h(V_C)$.
3. *Chooser* picks a commutative cryptography key pair, encryption key E_C and decryption key D_C , at random from the same group Z_p^* that was used by *sender* in the Phase A.
4. *Chooser* encrypts entries in the set X_C , so that: $Y_C = E_C(X_C) = E_C(h(V_C))$.
5. *Chooser* sends to *sender* set Y_C , reordered lexicographically.
6. *Sender* encrypts with key E_S each entry $y \in Y_C$ received from *chooser*.
7. *Sender* returns set of pairs $\langle y, E_S(y) \rangle$ to *chooser*.
8. *Chooser* decrypts each entry in $E_S(Y_C)$, obtaining $E_S(X_C) = D_C E_S(E_C(X_C)) = D_C E_S(Y_C)$.
9. *Chooser* compares each entry in $E_S(X_C)$ to the entries of Y_S constructed in Step A3 (Step 3 of Phase A) and received by the *chooser* in Step B1. This way the interesting records can be identified.

Phase C - OT

1. After identifying the interesting records in Y_S the *chooser* selects at random $\alpha - 1$ other unique records from Y_S for each interesting record in V_C . These are the diluting records, that together with the records of interest form a shortlist for the enquiry. If the number of interesting

records multiplied by o is greater than n , the size of the dataset V_S , then the complete Y_S is shortlisted.

2. Send the shortlist to *sender*.
3. *Sender* picks an encryption PH key E'_S at random from the group Z_p^* .
4. *Sender* identifies entries $h(v)$ from X_S that have been shortlisted and processes each shortlisted record in the following way:
 - (a) Encrypts $h(v)$ with E'_S to form the key used to lock the extra information about v , i.e. $\kappa(v) = E'_S(h(v))$.
 - (b) Encrypts the extra information using a symmetric encryption function K and the key $\kappa(v)$ crafted in the previous step:

$$c(v) = K(\kappa(v), ext(v))$$
 - (c) Forms a pair $\langle E'_S(h(v)), c(v) \rangle$.
5. The pairs formed in C4(c), containing a private match element and the encrypted extra information about record v , are then transferred to *chooser*.
6. *Sender* encrypts each entry $y \in Y_C$, received from *chooser* in Step B5, with key E'_S to form set of pairs $\langle y, E'_S(y) \rangle$.
7. Pairs $\langle y, E'_S(y) \rangle$ are then transferred to *chooser*.
8. *Chooser* removes the encryption E_C from all entries in the 2-tuples received in Step C7 obtaining tuples α, β such that $\langle \alpha, \beta \rangle = \langle h(v), E'_S(h(v)) \rangle$. Thus, α is the hashed value $v \in V_C$, and β is the hashed value v encrypted using E'_S .
9. *Chooser* sets aside all pairs received in Step C5, whose first entry is equal to one of the first entry of any two-tuples obtained in Step B9. Then uses the appropriate β tuple associated with a given interesting record as a symmetric key to decrypt the extra information contained in the second entry in the pair received in C5. This is performed for all the matching entries.

In this improved protocol the initial processing depends on the size of the dataset - n , but it needs to be performed only once in a given period of time, e.g. once per month, or per year. There is no need that the camouflage data is up to date, since the police is ex hypothesis investigating a past event, so might well be interested in a former client, or a client whose circumstances have changed. The remaining operations are less processing savvy as illustrated in Table II. The IDAP protocol has been implemented in the same fashion as the PE protocol described in Section 3.1. The results from the empirical evaluation matched the results that simulated using the computational complexity and cost presented in Table II.

TABLE II

Computational Complexity of Improvement 1

		Symmetric Crypto.	Asymmetric Crypto	
		crypto. operation	key generation	crypto. operation
Phase A (run periodically)	Step 1	-	-	-
	Step 2	-	$O(1)$	-
	Step 3	-	-	$O(n)$
Phase B (run per enquiry)	Step 3	-	$O(1)$	-
	Step 4	-	-	$O(m)$
	Step 6	-	-	$O(m)$
	Step 8	-	-	$O(m)$
Phase C (run per enquiry)	Step 3	-	$O(1)$	-
	Step 4(a)	-	-	$O(m \times o)$
	Step 4(b)	$O(m \times o)$	-	-
	Step 6	-	-	$O(m)$
	Step 8	-	-	$O(m)$
	Step 9	$O(m)$	-	-
Total Complexity for k enquiries, where $n < m \times o$		$O(km(o + 1))$	$O(2k + 1)$	$O(km(o + 5) + n)$
Cost (ms/operation)		0.33	7	30

The complexity of each of the steps in the proposed improved solution. Where n is the number of the data rows in the source, m is the number of interesting records. Also the diluting factor o, as well as the number of the protocol runs k affect the processing time required by the protocol. Cost is the measured average time in ms to perform given cryptographic operation from managed C#.NET code.

Fig. 3 illustrates the processes involved in this improved version of acquisition protocol. It is worth noting that there are only five communication rounds required in this protocol. This is two rounds more than in the original PE protocol, still, most of efficient SPIR protocols require considerably more rounds. This method provides significant improvements to the processing time required for enquiries if total number of records in the *sender's* database is higher than $o \times m$, i.e. higher than the number of interesting records m multiplied by the diluting factor o . This is illustrated in Fig. 4. Furthermore, the true strength of this version of the protocol is seen when multiple enquiries are run of the same database using a single encrypted catalogue of the records, compiled by the *sender* in Phase 1 (shown in Fig. 5).

Fig. 3 IDAP Process Flow

Graphical representation of the improved IDAP

Fig. 4 Processing time per enquiry depending on the number of interesting records

This proposed modification of the protocol improves significantly the processing time required for the protocol to run for the cases where the product of the number of the interesting records m and diluting factor o is smaller than the number of the records in the database n .

Fig. 5 Processing time depending on the number of enquires

This proposed modification improves significantly the processing time required for the protocol to run for the cases where more than one enquiry is run against the same database.

4.2 Allow multiple selection criteria

The PE protocol can be used to privately retrieve data if the data is identified by a single parameter, such as ID number, credit card number, IP address, etc. However, this is not always the case. Consequently, if IDAP is used to

find a suspect based on circumstantial knowledge, or a suspect's profile, the PE protocol needs to be modified. Query (4) shows the way the request (3) would be modified for such enquiry, here sip_{1-j} stand for j secret input parameters:

```
SELECT sip1, sip2, ..., sipj, rp1, rp2, ..., rph
FROM source (4)
WHERE ip1=ip_val1 AND ip2=ip_val2 AND ... AND ipl=ip_vall
```

A computationally expensive solution to this problem has been published by Kwecka, Buchanan, and Spiers (2010). The authors suggest that symmetric encryption should be used to lock the return parameters and the symmetric keys should be secured with relevant commutative encryption keys that are unique to each value of the secret input parameter returned for the given row. Despite being computationally expensive, this solution has a unique benefit of allowing semi-fuzzy matching of the results if the underlying commutative protocol is ElGamal-based. In this paper a simplified approach is proposed. Since the query (4) replaces the ri parameter with j different sip parameters, the list of these j parameters could be used as a complex ri in the improved IDAP protocol. Thus, in Steps B2 and A1 a list of all values of given sip parameters would be hashed together to form records in sets V_C and V_S . This way the security of the protocol nor its complexity is affected by this improvement.

4.3 Correctness and Security

IDAP is a modification of the PE protocol that has its correctness and security proofs provided in Agrawal Evfimievski, and. Srikant (2003). The goals and logic of IDAP and PE are similar; however, IDAP is streamlined to provide better performance than PE in the specific use scenario of investigative data acquisition. There is an assumption that there is a method of authenticating other parties and securing the channel for communication. In order to evaluate the correctness and security of IDAP the inputs and outputs need to be clearly stated (Cristofaro et al 2009):

Chooser's input: set V_C containing IDs of interesting records.

Sender's input: set V_S containing IDs of the records in the dataset, together with extra information about these records – $ext(v)$.

Output: *chooser* learns $|V_S|$ (the size of the set V_S), $V_C \cap V_S$, and $ext(v)$ for $v \in V_C \cap V_S$, while *sender* learns $|V_C|$. *Proxy* learns only the sizes of the sets.

Normally both parties learn the sizes $|V_S|$ and $|V_C|$, as by default all the encrypted identities in V_S are sent to the *chooser*, while the *chooser* in order to find the interesting records among these encrypted identities and in order to decrypt the $ext(v)$ for these records provides the *sender* with encrypted elements of the set V_C . There is no requirement by the public authorities to know the size of the dataset, but since there is now a way to run IDAP and avoid providing the authorities with the dataset size, this needs to be accepted as an outcome of the protocol. The fact that the *sender* learns the number of interesting records is beneficial in the data acquisition scenario, as the *sender* can then verify that the *chooser* follows the data acquisition notice that would previously outline the IDs of the interesting records, and under IDAP would specify the number of the interesting records.

IDAP is based on *Shamir's commutative protocols*, a variant of PH protocol where the prime p is public and common between the communicating parties. An adversary with the knowledge of the ciphertext C and the prime p would need to solve the following hard problem to break the commutative PH protocol (Schneier 1995):

$$e = \log_p C \bmod p$$

Just like RSA, the ciphertext created using the PH algorithm may leak some information about the input plaintext message. Therefore, this algorithm is

suitable for uses where the input is formed from random data. This is the case in the PE and IDAP, as the commutative PH is used to encrypt hashed IDs of the records. While it is normally recommended to use padding schemes in any implementation of RSA (Kaliski 2003), and thus PH implementation as well, the PE and IDAP mitigate this requirement by using fixed size hashes as the input.

The proofs of the correctness and security of PE can be found in . Agrawal, A. Evfimievski, and R. Srikant (2003), while IDAP has modified this protocol by introducing the following improvements:

- Lowering processing time, by narrowing the scope of the enquiry.
- Allowing for multiple selection criteria.
- Restoring the balance between the privacy of the innocent and the suspects.

In order to narrow down the scope of the enquiry IDAP splits the PE protocol into three parts. However, the only way the operations of the protocol are affected is the fact that under IDAP the *chooser* request extra information for only $(m \times o)$ records, rather than for the whole dataset n . The main consequence of this approach in respect to the security of the protocol is that the *sender* knows that there are m interesting suspects in the set of identities the size of $(m \times o)$. This could become an issue if the same request is run against a number of parties and the parties collude, but this thesis has shown that the investigative data acquisition process can be treated as a single database scenario, if requests are made against CSPs. Therefore colluding is not possible. On the other hand, for small organisations with less than 100,000 IDs, there is no need to narrow down the results. Consequently, in IDAP, the privacy of the suspect is affected by the diluting factor o , and the *sender's* probability of guessing the interesting records IDs is $1:o$ and not $1:n$. As long as o is reasonably large, and the *sender* has no other sources of information about the suspects, the privacy of the suspects should be safe.

IDAP allows for the multiple selection criteria by hashing together different selection criteria and using it within the PE protocol as an ID of a record. This

does not affect the security of the PE protocol. On the other hand adding a semi-trusted third party – the *proxy* – in order to restore the balance between the privacy of the innocents and the suspects that we will discuss in the next section would somewhat modify the security of the protocol. The *proxy* filters out the records not classified as interesting from the *sender's* response. Assuming that the semi-trusted party behaves as expected, the security of the $ext(v)$, the data records contained in the *sender's* database is information theoretic from the *chooser's* perspective. On the other hand, if the *proxy* and the *sender* cooperate, they can easily work out the identities of the interesting records. The main aim of IDAP is to hide those identities from the *sender*. However, under current practice, the identities of the suspects are provided in every data acquisition notice. Consequently, if the semi-trusted party were to cooperate with the *sender*, this would only reveal information that is currently openly communicated to the dataholders anyway, making the worst case scenario not worse than current best practice.

4.4 Communicational Complexity

The cost of communications should also be considered when discussing IDAP. This cost depends on the diluting factor, just as the cost of processing does. Thus, for low values of o , such as 1,000, the cost of communications should be reasonable. However, where higher degree of privacy and secrecy is required, the costs of on-line communications could prove to be prohibitive. In such cases, it would be possible to exchange encrypted data via the post or couriers, as there are a small number of communication rounds between the parties.

5. Assessing privacy risks and Data Protection Directive compliance

In this final section we return to our discussion from the beginning and evaluate the wider legal and societal issues that this proposal raises and assess the privacy risks that are involved. We recommend in response two

institutional aspects to complement the technological solutions described above.

Let us recap quickly the main features of the system that we have described so far. The police is interested in our target, John Doe. They make a request for data about Doe to the online provider X. Since X does not need to know the identity of the suspect, and may draw adverse inferences about him if he knew that the customer was target of an investigation, the police requests data from a larger set of people (the foils), chosen randomly. Since the retailer knows that only one of the people whose information he hands over is the suspect, he can't any longer draw an adverse inference against any individual; the community hides the identity of the suspect from the retailer behind a wall build by them all, just as in the Spartacus example. At the same time, the data of all the customers is encrypted in such a way that the police can only make sense of the data that belongs to the suspect – a key has been created prior to making the data request that opens only that data for the specific subject under investigation. The encryption renders the records unusable to the authorities in the sense that they are secure against attacks in polynomial time. This prevents “fishing expeditions”, and ensures that the data of the innocent customers can't be used by the police for other purposes.

However, this still involves providing government agencies with records of individuals that are “innocent bystanders”, which raises legal issues as well as issues of public acceptance. There are some additional actions that may reassure the public that the data is safe. First, if the technique for minimising the processing time (Improvement 1) is employed, the chances that investigators will retrieve encrypted records of a particular individual that is not a suspect are small in large datasets. Thus, for a dataset with n records, during investigation with m interesting records and the diluting factor o the probability of this event A can be defined as (5)

$$(5) P(A) = \frac{(o-1) \times m}{n-m}$$

Consequently for investigations with five interesting records, with diluting factor of a thousand and a dataset consisting a million records, the probability of this event occurring during a single run of the protocol would be less than 0.5%. Since the runs of the protocol are independent this probability would stay the same. This also means that the investigators would need to first break the encryption key used by the *sender* to hide identities (Phase A), before they could attempt to obtain the data about a specific individual that is not a suspect, otherwise the probability of the encrypted data being provided to them would be small. Additionally, if the identity of a data subject is never encrypted under the same key as the data records, then investigators would need to successfully brute force two separate keys in order to make use of the retrieved encrypted records. Otherwise the information would be unintelligible.

The chosen encryption method makes it all but impossible for the police to get access to the data of the “foils”. Even if they could access this data, it would in all likelihood be of no interest to anyone, and have no potential of privacy harm, as was generated randomly. In addition to the relevant information that Doe bought large quantities of fertiliser – relevant giving the investigative hypothesis that he is a bomb maker – the police would learn nothing more significant than e.g. that a Mr. Smith bought a shovel and Mrs Jones a wheelbarrow from the same farm equipment company. At the same time, the police would become exposed to a significant risk themselves for violating their legal obligation to destroy this data unseen. This random character of the camouflage information therefore prevents the police from using this data strategically. However, some of the data could expose the data subject to risks other than privacy risks. If for instance the data from the “foils” happens to be credit card details, and the police were to lose this data before destruction, people may fear that they have been exposed to an unacceptable risk that the data can fall into the hand of criminals. That the data is highly encrypted may be insufficient to alleviate this fear. Acceptability therefore depends also on the public trust into the data handling and security procedures used by the police institutionally, not just the technology provided by our approach. Most security professionals trust into a security process more than they trust in encryption. The solution proposed here is that in order

to reassure the public, a semi-trusted third party needs to be involved and integrated into the protocol. The following modifications to the IDAP are proposed:

1. All communication between *chooser* and *sender* goes through *proxy*.
2. Chooser provides *proxy* with the identifiers of the interesting records encrypted by *sender*, $E_s(h(v))$. This is done over a secure channel or with use of a 3Pass protocol once the parties are authenticated.
3. At the stage where data is transferred from *sender* in Step C4, *proxy* filters the response and discards the records that were not specified by *chooser's* request, i.e. the records other than the ones identified in Step 2.

The semi-trusted party should have no interest in finding out the object of the investigation or the content of the data records returned by the dataholder. The party that is chosen must not cooperate with the *sender* or the protocol will be broken, since simple matching exercise would reveal the identities of the suspects. A key concept is that the *proxy* has no incentives to find out the detail of the investigation, thus it is not going to invest in expensive cutting - edge decryption technology to decode the data, nor it is going to cooperate with the *sender* in order to establish the identity of the suspect. On the other hand, if the need arises to verify the *chooser's* requests in front of a court of law, the *proxy* and the *sender* could work together to establish the identities of the records requested by the *chooser*, or verify that the data request by the police was in conformance with the warrant that was granted. This introduces an additional “price” for the police – in return for more secrecy vis a vis the data controller (the online retailer or bank) and a more efficient search, they are also subject to more scrutiny and transparency, as data requests are now necessarily lodged with a third party, that can also check if the formulation of the search query was law compliant. Since under current law, any camouflage data would have to be destroyed immediately after the data of the suspect is transferred to the police (as storage after this point would be unnecessary, in violation of the data minimisation principle), the problem that the police might

be tempted to “store” the excess data until such a time that decryption technology has improved is addressed.

Nonetheless, from a legal perspective even encrypted data is still personal data under the Data Protection Directive, even though the UK implementation of this European Directive is less clear on this point. This means one of the six legally valid grounds for processing the data must hold. The most obvious one is the consent of the data subjects, and we will come back to this option below. Another basis can be a legal duty created through statute.⁴ Currently, no such duty to shield each other exists in any EU member state. However, the arguments that we developed in the first part of this paper would at least permit legislators to create such a duty. Even though it would impose a (minimal) privacy risk for the “foils”, since this is required to reduce the much greater privacy risk of the suspect, we argue that such a *prima facie* infringement would be proportionate, efficient and necessary. Finally, using again the notion from the beginning of this paper that privacy is as a common good that is fundamental for a free, democratic order, it may even be possible to permit such an approach even in the absence of new legal duties. Art 7(e) of the directive creates a blanket exception if the processing of the data “necessary for the performance of a task carried out in the public interest” This in turn *might* make it unnecessary to require consent from those customers whose data is used merely to hide the identity of the suspect. Just as our privacy can be violated as part of a criminal investigation to further the public good of efficient law enforcement, so one could argue that we are also required to shoulder a purely abstract privacy risk to maintain the foundations of a democratic. Similar arguments have been made in the past regarding medical research data and “benefit sharing”: as long as I benefit in the long run from medical research, solidarity requires that I take a marginal privacy risk in making some of my data, in an anonymised, encrypted format, available for research (Wicks et al 2010, Laurie and Sethi 2013)). We have a similar benefit sharing here – everybody can become subject of a police investigation, so in the long run, I share the benefits from a system that pools

⁴ Data Protection Directive Art 7(c)

all our records and selects randomly a few of them each time a the modern equivalent of a Roman General asks is: Which one of you is Spartacus?

Basing the approach on Art 7(c) or 7(e) respectively would result in slightly different legal regimes, and therefore also slightly different implementations of the approach. In neither case, consent of the foils” is necessary. However, if governments were to decide to impose a new duty under Art 7(c), the approach proposed in this paper, or a functionally equivalent solution, would become legally mandatory and therefore used by all online organisations that store customer data. Art 7(e) by contrast simply creates a permission for online retailers to implement this solution if they consider it beneficial for them, and we would expect a much less widespread uptake, with market forces ultimately deciding on its acceptance.

However, the absence of case law makes it difficult to assess if this argument, which rests exclusively on the strength of the jurisprudential analysis of privacy outlined in the first part of this paper, would withstand scrutiny by the courts. A legally safer option is therefore to ask for a generic consent from customers – “are you willing to put your data in a pool if and when there are police inquiries in the future”? This anticipatory consent prevents time delays during investigations. Whether or not a sufficient number of customers would be willing to subscribe to such a scheme requires further, empirical research that should also address the question how adequate incentives could be designed. We noticed above the possible conflict between a conception of privacy as a public good and the notion of consent as the ultimate “trump” that can lead to individuals opting out of their legally guaranteed protection. *Prima facie*, the situation is different here. While enlightened self – interest is one reason people may have for allowing their data to be used as camouflage in an investigation, the aim ultimately is to protect a common good. However, if solidarity alone is not sufficient to incentivise customer’s to protect in a mutual privacy protection scheme, other incentives could be found. One possibility would be to require reciprocity in order to be protected under the scheme –

only those who “donate” their data will benefit if they themselves should come under the spotlight.

However, it is at least questionable if this approach would be legally sound – after all, if the police makes an inquiry regarding someone who is not participant in the scheme, his data would be treated with less concern for privacy than possible in principle, which would arguably mean that the data controller, that is the company, is in violation of data protection law. This again reflects that ultimately privacy would be treated in such an approach as an alienable property, to be assigned away provide consent is given

6. Conclusion and further work

Our investigation started with a common privacy problem in online investigations: In order to obtain data about a suspect, the police must disclose to the data controller (a bank, and ISP etc) the identity of the “person of interest”. This poses a privacy and reputation risk to the suspect: people often assume that “where there is smoke, there is fire”, and even being subject of a police investigation carries substantial reputation risks – holders of public office e.g. will frequently resign even at such an early stage of a criminal investigation. It also poses a risk for the police investigation and its integrity, as it can warn off suspects and increase their flight risk. A combination of technical and legal factors prevents the use of strategies to minimise these risks that are used in the offline environment. Traditional approaches to privacy protection online also struggle with this scenario, as they typically pitch state interests (here, the police) against those of the citizen. In our setting though, these interests converge. By looking at new and emerging conceptions of privacy that understand it less as an individual right only, but as a communal good that enables important social institutions in a democratic society, we were able to overcome this gridlock and suggest a combination of technical, attitudinal and legal measures. Novel about this approach is in particular the notion of privacy as communal responsibility, which allows accepting small privacy risks for a larger number of people to prevent more serious privacy risks for other individuals.

Because this conception of privacy differs from the traditional jurisprudential conceptualisation, it raises several questions about the legal evaluation of our proposal. We discussed possible legal foundations that allow the necessary data transfer, concluding that for both the most promising ones, a significant degree of public acceptance is crucial. The success of our proposal will therefore ultimately depend on empirical, social factors regarding risk assessment, solidarity and community loyalty. Further research should in particular look at social attitudes to “privacy risk sharing”, and how, if at all, it differs between different online communities. We noted above a slow but noticeable shift towards a greater concerns for privacy, and a greater willingness to prioritise it over investigative interests by the police. We should therefore expect uptake to be highest in those environments where mutual solidarity and a feeling of belonging is strongest, for instance voluntary internet based associations such as the community of Wikipedia editors, and the lowest where the community” is one of mere convenience, such as the community of Amazon customers”.

Developing appropriate incentives is therefore one of the key tasks for future research. As the initial problem was caused by a combination of traditional legal concepts and their lack of “fit” with modern online environments, our solution too employed a combination of legal and technological approaches. Further research is therefore needed on legal, technological and organisational aspects alike. From a technological perspective, improving further our idea that for specific queries, different ratios between “camouflage” and “real” data are better than a “one size fits it all approach” will be further explored. This involves studying further the balance between number of foils, sensitivity of data and resulting risks. The challenge here is also to balance protection from risk against communication complexity in both legally and technologically sound ways. Exploring different ways to balance communication complexity, different key sizes and the ratio between interesting/extra data that is sent to the investigators should result in a

number of typical risk profiles, which can lead to partly automated choice of protocols.

A different task will be to extend our approach beyond the simple model of a one off query of the type typically encountered in police investigations. Were the police to make several queries about the same suspect to the same data controller in a short period of time, the controller might be able to triangulate the identity of the suspect after all. This would still require much more effort than they have to invest at present, but would at least be theoretically possible. Multiple queries of this type are rare, due to police operational reasons (and also legal constraints), much more common however are of course request for the long term surveillance of an account in situations where the goal is prevention of future crimes rather than investigation of a past crime as in our scenario. A natural extension of our idea would therefore to be the study of long term, real time surveillance operations which inevitably would demand much more from the “foils”. Our approach to think of PETs as communal tasks should either way make a valuable contribution to the range of PET tools that are available. In the past, they reflected the libertarian, individualistic concept of privacy law, equipping individuals with protective tools that “build walls around them” within which they can keep their data safe. By contrast, our approach is a tool for the emerging understanding of privacy as a public good, where the protection of anonymity becomes a communal task, where we are strong only when united.

Literature

- Agrawal, R., Evfimievski, A., & Srikant, R. (2003). *Information sharing across private databases*. Paper presented at the Proceedings of the 2003 ACM SIGMOD international conference on Management of data, San Diego, California.
- Aiello, B., Ishai, Y., & Reingold, O. (2001). Priced Oblivious Transfer: How to Sell Digital Goods. In B. Pfitzmann (Ed.), *Advances in Cryptology — EUROCRYPT 2001* (Vol. 2045, pp. 119-135): Springer-Verlag.
- Asonov, D., & Freytag, J.-C. (2003). Almost Optimal Private Information Retrieval. In *Privacy Enhancing Technologies* (pp. 239-243).

- Balz, D., & Deane, C. (2006, 11/1/2006). Differing Views on Terrorism. *The Washington Post*, p. A04.
- Bao, F., & Deng, R. (2001). Privacy Protection for Transactions of Digital Goods. In *Information and Communications Security* (pp. 202-213).
- Bignami, F. (2007) Privacy and Law Enforcement in the European Union: The Data Retention Directive. *Chicago Journal of International Law*
- Bloustein, E J. (1964), Privacy as an aspect of human dignity: An answer to Dean Prosser, *NYUL Rev.* 39: 962
- Cachin, C. (1999). Efficient private bidding and auctions with an oblivious third party. In *Proceedings of the 6th ACM conference on Computer and communications security* ACM. 120-127
- Cristofaro, E. et al (2009). *Privacy-Preserving Policy-Based Information Transfer. Proceedings of the PETS '09: 9th International Symposium on Privacy Enhancing Technologies*, Seattle, WA, 164-184.
- Du, W., & Atallah, M. J. (2001). Privacy-Preserving Cooperative Scientific Computations. *Privacy-Preserving Cooperative Scientific Computations. csfw Vol 1* 273.
- Frikken, K. B., & Atallah, M. J. (2003). *Privacy preserving electronic surveillance*. Paper presented at the Proceedings of the 2003 ACM workshop on Privacy in the electronic society, Washington, DC.
- Goldwasser, S., & Lindell, Y. (2002). *Secure Computation without Agreement*. Paper presented at the Proceedings of the 16th International Conference on Distributed Computing.
- Home Office. (2007). *Acquisition and Disclosure of Communications Data - Code of Practice*. Retrieved. from <http://security.homeoffice.gov.uk/ripa/publication-search/ripa-cop/acquisition-disclosure-cop.pdf?view=Binary>.
- Home Office. (2009). Protecting the Public in a Changing Communications Environment [Electronic Version]. Retrieved 20 April 2009 from <http://www.homeoffice.gov.uk/documents/cons-2009-communications-data?view=Binary>.
- Iliev, A. and. Smith S. W, (2005) Protecting Client Privacy with Trusted Computing at the Server," *IEEE Security and Privacy*, vol. 3, 20-28
- Kaliski, (2003). RSA Problem, in *ACM SIGKDD Explorations: MIT Laboratory for Computer Science*, 2003, 10
- Kwecka, Z., Buchanan, W., Spiers, D., & Saliou, L. (2008, 30 June – 1 July). *Validation of 1-N OT Algorithms in Privacy-Preserving Investigations*. Paper presented at the 7th European Conference on Information Warfare and Security, University of Plymouth.
- Kwecka, Z., Buchanan, W. J., & Spiers, D. (2010). *Privacy-Preserving Data Acquisition Protocol*. Paper presented at the Sibircon, Irkutsk.
- Lawless, M.D. (2007). The third party doctrine redux: Internet search records and the case for a " Crazy Quilt" of Fourth Amendment protection. *UCLA JL & Tech.* 2-6.
- Laurie, G Sethi. N. (2013) Towards Principles-Based Approaches to Governance of Health-Related Research Using Personal Data. *European Journal of Risk Regulation* 1/2013: pp. 43-57
- Ostrovsky, R., & William E. Skeith III. (2007). A Survey of Single-Database PIR: Techniques and Applications. In O. Tatsuaki & W. Xiaoyun (Eds.), *Public Key Cryptography* (Vol. 4450, pp. 393-411). Berlin: Springer

- Rasmussen Reports. (2008). 51% Say Security More Important than Privacy. Retrieved 01/09/2009, from http://www.rasmussenreports.com/public_content/politics/current_events/general_current_events/51_say_security_more_important_than_privacy
- Regan, P. M., (1994), *Legislating Privacy: Technology, Social Values and Public Policy*, Chapel Hill: The University of North Carolina Press
- Rivest, R. L., Shamir, A., & Adleman, L. (1978). A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21(2), 120-126.
- Schneier, B. (1995). *Applied Cryptography: Protocols, Algorithms, and Source Code in C*: John Wiley & Sons, Inc.
- Shamir, A. (1980). *On the Power of Commutativity in Cryptography*. Paper presented at the Proceedings of the 7th Colloquium on Automata, Languages and Programming.
- Shannon, C. (1949). Communication theory of secrecy systems. *Bell System Technical Journal*, 28
- Simitis, S., (1987), Reviewing Privacy in an Information Society, *University of Pennsylvania Law Review*, Vol. 135, No. 3 pp. 707-746
- Tene, O. (2008)., What Google knows: Privacy and internet search engines. *Utah Law Review*, 1433-1492
- Swire, P., & Steinfeld, L. (2002). *Security and privacy after September 11: the health care example*. Paper presented at the Proceedings of the 12th annual conference on Computers, freedom and privacy, San Francisco, California.
- Weis, S. A. (2006). *New Foundations for Efficient Authentication, Commutative Cryptography, and Private Disjointness Testing*. Unpublished PhD Thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Wicks, P, et al. (2010) Sharing health data for better outcomes on PatientsLikeMe." *Journal of medical Internet research* 12.2
- Young, B. C., Kathleen, E. C., Joshua, S. K., & Meredith, M. S. (2006). Challenges Associated with Privacy in Health Care Industry: Implementation of HIPAA and the Security Rules. *J. Med. Syst.*, 30(1), 57-64.