PrivGuard: A Model To Protect Private Information Based On Its Usage

FA Lategan[a] MS Olivier[b]

[a]*Department of Computer Science, Rand Afrikaans University, PO Box 524, Auckland Park, 2006, South Africa*
[a]*fransl@discoveryhealth.co.za*
Department of Computer Science, University of Pretoria, Pretoria, South Africa

Abstract

No reliable method currently exists to safeguard the privacy of private information. Privacy policies are insufficient as compliance can not be enforced automatically. In this paper we propose a model to improve the control the owner of private information has over its protection. This is achieved by classifying private information based on the purpose it is acquired for, and then designing methods to protect each class of private information. Private information is then encrypted using homomorphic functions where such information is only required for validation. The validation can then be performed without divulging the actual private information. In cases where private information is required for other usages, a system based on Kerberos and trusted third parties is used in order to maintain as much control over private information as possible.

Keywords *Privacy, access control, encryption*
Computing Review Categories *K.4.1, E.3, K.6.5*

1 Introduction

In most cases where a subject is granted access to data, such access is limited. Typical limits are usually an expiry date on such access, restrictions on what can be done with the data or restrictions on where the data can be accessed from. No organisation would give a subject carte blanche access to their data on a vague promise to take good care of it. Unfortunately, this is exactly what happens when individuals supply their private information to some target organisation on the Internet — that target organisation might have a privacy statement, but effectively has total control of the private information, and might resell, redistribute or modify the data without the owner's knowledge or consent. Furthermore the owner can not "take back" his or her private information, since there is no secure way to force the target organisation to delete the data.

Privacy is an abstract concept. For our purposes, we shall define privacy as a state that exists when access to private information about a particular individual can be effectively controlled and managed by that individual even after a third party has collected such private information. The aim of privacy is not to prevent the use or collection of private information, but rather the misuse (intentional or not) thereof.

In order to safeguard private information, while still allowing it to be used for the intended purpose, we categorise it according to the purpose the information is required for. This allows us to tailor a protection mechanism for each category without invalidating the usefulness of the information.

In this paper we present a model we call PrivGuard to allow an individual to supply his private information to a target organisation in a way that limits the access such an organisation has to the information. To do this, we require several well known tools such as tickets, public key encryption and a trusted third party.

PrivGuard uses a two-pronged approach based on whether or not private information actually has to be disclosed in order to fulfill the purpose it was requested for. An early version of the classification and protection for the second class of private information was presented at the WWW10 conference [7]. The method used to verify calculations in the first class of private information was presented at the IFIP TC11 sixteenth annual Working Conference on Information Security [6].

This paper is structured as follows: in Section 2 we shall supply some background information. Our classification of private information is outlined in Section 3. We then present the implementation in Sections 4–6 followed by examples in Section 7. Related work is discussed in Section 8 and Section 9 contains the summary.

2 Background

For effective electronic commerce every individual *i* is forced to reveal some private information *m* about herself at some stage to a target organisation *o*. Even if *i* trusts *o* to perform the electronic transaction, it does not imply that *i* would like *o* to keep a permanent record of *m* in a database. Unfortunately, *i* has no control over *m* as soon as *m* is disclosed to *o*, and stored in *o*'s database. These databases are increasingly compromised, misused, sold, or

even made freely available to the public over the Internet. See [17, 15, 3, 2] for a sample of some of the concerns over the safeguarding of private information.

One of the ways that private information can be safeguarded, is by using privacy policies (see [3]). A lot of research has gone into automating such policies, as described in the P3P protocol [11, 16] but these policies are still not enforceable. However, P3P allows an individual *i* to define a set of acceptable usage rules of private data *m*. Any organisation *o* with a published privacy policy *P* that does not violate these rules can request *i*'s browser to automatically supply those parts of *m* that are required. This saves *i* from having to read and interpret every site's privacy policy before sending private information to such a site, and from having to manually enter the private information every time. Although P3P can define arbitration authorities for disputes, adherence to published policies is not enforced; safe transmission of data between parties is also not described.

There is therefore a clear need for privacy to be enforceable to prevent misuse before it occurs, and not only to have a legal recourse after the fact. In certain cases no legal proceeding can undo the damage done by such misuse; it is far better for such misuse not to happen at all.

Kerberos [10, 5, 12], Secure Electronic Transaction (SET) [14], Digital signatures, public key encryption and homomorphic encryption algorithms [13, 8, 1] are also used in PrivGuard.

3 Classification

3.1 Introduction

Access to private information is more complex than simple authentication and authorisation. In most cases where private information is required for an online transaction, authentication has occurred by the time that the information is requested (usually by logging on to the organisation's website securely with SSL and verifying the site's certificate) and authorisation is implicitly given by the user at the time by supplying the private information to the organisation. Our overall purpose is to restrict what can be done with such information once access to it has been granted, and to prevent misuse, intentional or not. In order to adequately protect this information, no single method will suffice, and we first require a way to group the information into several classes for which protection can then be designed.

As the purpose of our classification is to simplify the protection of the private information, we propose a classification of private information based on the purpose the information is required for. This will allow us to tailor the protection of the information for each class in such a way that the maximum protection can be given whilst still allowing the purpose the information was requested for to be fulfilled. For this reason our classification consists of three different levels of access required in order to fulfill the purpose it was requested for, starting with no access required as the first level up to full access required on the third. Protection of the private information which will be discussed later will then obviously vary from very good in class one to none at all for class three.

3.2 Classes

3.2.1 Class 1

The first class shall consist of private information not actually required for the transaction taking place, but used for some lateral purpose, such as to validate a calculation or generating statistics. An example of this is the submission of a tax return, where a lot of private information is given such as the exact amount of interest earned, royalties received, medical expenses incurred, etcetera. This information is not actually required directly by the tax authorities, as they intend using it only to verify the calculation of the taxable income.

This class of private information can be very well protected against misuse as the information might not actually be divulged to the organisation requiring it if methods can be devised to satisfy the purpose the information was requested for without revealing the information itself.

3.2.2 Class 2

The second class we define will contain the cases where direct access to private information might be required at some stage, but not necessarily at the time it is requested, or by the organisation requesting it. This class can also be subdivided into smaller classes.

3.2.3 Class 2.1

This class we define to contain all the instances where private information is requested during the processing of a transaction on the Internet that is not required by the organisation requesting it, but will be passed on to a third party for a purpose related to the transaction being performed. There are many examples of this happening, such as online vendors requiring banking details and shipping addresses of customers. The online retailer can not directly use the customers' bank account numbers; they have to pass the banking details on to their bank, and the bank will use the actual banking details to perform the fund transfer. In the same manner the online retailer does not use the actual shipping address of the client; they just print it on a label and the shipping company uses the address to deliver the package to the client. In this example it is useful to note that the zip code is often used to calculated shipping costs payable by the customer and the organisation would require the actual zip code, but the rest of the shipping address is normally not directly used.

3.2.4 Class 2.2

In the second subclass we consider the cases where an organisation requests private information from a customer for possible future usage related to a transaction being performed, but it is not guaranteed that the information would actually be used at some future date. There is also the added risk to the organisation that such information could become incorrect or expire before its eventual use, making it worthless. Online hotel reservations are a case in point — the hotel usually requires a credit card number to confirm the reservation, but this is not necessarily the way that the final bill will be settled. The credit card will only be used in case of a no-show to charge a late cancelation penalty. In cases where the reservation is several months in advance, it is quite possible that the credit card, while being valid at reservation time, might have been closed or suspended before the actual date of the accommodation. Even though credit might have been reserved on the account, the delay between the reservation of credit and actual payment request could present problems.

3.2.5 Class 2.3

In our third subclass we shall place private information collected for the sole purpose of uniquely identifying customers to an organisation or to allow customers to log in using an account with the organisation. Examples of such private information include social security number, e-mail address, mother's maiden name, etc. Quite often more than just identity is associated with such information, or the same information is used at more than one site, increasing the risk to the customer should such private information be compromised.

3.2.6 Class 2.4

This subclass shall consist of all cases where the private information is directly required for the current transaction, and any amount of obfuscation will make it impossible to fulfill the purpose the information was acquired for in the first place. This could be information such as the list of books ordered from an online bookshop (when packing the box), or an individual's taxable income on an electronic tax return (when determining the amount of tax payable).

3.2.7 Class 3

This class contains all private information, and is a container for all the other classes, in that all the other classes are subsets of this class.

3.3 Observations

It is clear that any item of private information from classes 2.1 to 2.3 can be treated as though in class 2.4 without invalidating the purpose it was acquired for. The principal difference between classes 2.4 and 3 is that class 2.4 will still be included in our protection mechanisms, though nothing is preventing an organisation accessing class 2.4 information from actually storing it after the first access and bypassing the protection for subsequent accesses, while class 3 information will not be included in our protection mechanisms. Private information in class 3 will typically include any other items of private information about an individual that exist, but are not needed or requested for the transaction taking place.

4 Details of PrivGuard

The private information is encrypted if it is part of class 1. Tickets are used to access the actual private information for the other classes. A ticket granting ticket (TGT) describes the types of access allowed, and is used to request tickets from a trusted third party S that can be used to access the actual data. We describe the use of tickets in more detail, by applying our categorisation defined in Section 3.

4.1 Class 1: Validation and calculation

In this class private information pertaining to a particular individual is not used directly; it is collated or joined in some way, giving an end result which is then used. More formally, information attributes $m = \{x_1, x_2, x_3, \ldots, x_n\}$ about an individual i is collected from a number of sources to which some function G is applied to give a result α where $\alpha = G(x_1, x_2, x_3, \ldots, x_n)$. The result α is then passed on to a target organisation o.

What makes problems of this type interesting, and indeed what allows us to protect the private information, is the fact that although the target o needs to verify the calculation of α, o does not strictly require access to $x_1, x_2, x_3, \ldots, x_n$.

We describe a method to prove to o that $\alpha = G(x_1, x_2, x_3, \ldots, x_n)$ without disclosing $x_1, x_2, x_3, \ldots, x_n$. This would enable us to retain exclusive control over our private information in this class.

4.2 Class 2.1: Third party requirement

When a target organisation o_1 requires private information m from an individual i to pass on to a third party o_2, o_1 can send a ticket to such a party allowing it to get the data directly from S.

4.3 Class 2.2: Future use

In this subclass of private information where o would like to store m for some future use, o can just keep the TGT, and request a ticket from S when such access becomes needed. A further benefit is that m remains current, since all updates at S will filter through when o needs to access m. The availability of m is linked to the expiry date on the TGT.

4.4 Class 2.3: Identification

As i's public key is part of the TGT, o can just encrypt a random message with it, and request i to decrypt it using i's private key. No actual knowledge of or access to private information is required by o in such a case.

4.5 Class 2.4: Actual contents

In the last subclass where actual access to m is required to actually deliver a package, or actually transfer money (used by a bank) the real data is required, and is retrieved from S with a valid, unused ticket. The information is protected in the sense that no intermediary will have access to it.

The first four uses can be achieved without actually revealing m to o. The ticket granting ticket is all that is required. Only in the last case is the actual private information required, but privacy is still protected in a way, since it is only revealed at the last possible stage of any transaction. (And then it is minimal information such as: ship package number 1342 to this address, or transfer $23.54 from account number 352 to account number 2435).

5 Details of Class 1 implementation

We discuss the implementation of PrivGuard to Class 1 problems by referring to the income tax example.

We propose to encrypt $x_1, x_2, x_3, \ldots, x_n$ using a function f such that $y_i = f(\varepsilon, x_i)$ where ε is the encryption key. The result of the encryption will give us $y_1, y_2, y_3, \ldots, y_n$. Note that most of these values are usually supplied by third parties such as employers or banks, and that they can use these encrypted values on the tax certificates issued by them instead of the original values. This prevents the tax payer from modifying these values.

If we submit our tax returns using $y_1, y_2, y_3, \ldots, y_n$ instead of $x_1, x_2, x_3, \ldots, x_n$ we can protect a lot of sensitive information, but the IRS loses the ability to verify α.

What we now require is a function G' such that $G'(y_1, y_2, y_3, \ldots, y_n) = \alpha'$ where $\alpha' = f(\varepsilon, \alpha)$. This would allow the IRS to verify the correctness of α using the encrypted values $y_1, y_2, y_3, \ldots, y_n$ without ever knowing $x_1, x_2, x_3, \ldots, x_n$ by applying the function G' to $y_1, y_2, y_3, \ldots, y_n$ to obtain α' and then apply function f to α' and comparing the results. If $f(\alpha) = \alpha', \alpha$ has been correctly calculated from $x_1, x_2, x_3, \ldots, x_n$.

5.1 Restrictions

It is clear that we have to supply o with f and ε, which rules out symmetric encryption as a possibility for f. Furthermore for a given mapping $f(x_1) = y$ it should be hard or impossible to find x_2 such that $f(x_2) = y$. The function f would still protect the private information without this restriction, but then a possibility of fraud exists: assume $(\exists x_1, x_2)(x_1 \neq x_2)(x_1, x_2 \neq 0)$ such that $f(x_1) = f(x_2) = y$ with x_2 easy to find; then the sender could replace x_1 in all calculations with x_2 without affecting the result. To prevent this, we therefore assume that f has to map one and only one x to each y. We can summarise these restrictions more formally.

Suitable functions for f, G and G' would be functions that have the following properties:

Property 5.1 *The function $f(\varepsilon, x) = y$ should obscure x so that there is no easy way to determine x given f, ε and y.*

Property 5.2 *If $f(\varepsilon, x_1) = f(\varepsilon, x_2) = y$, then $x_1 = x_2$*

Property 5.3 $G'(f(\varepsilon, x_1), f(\varepsilon, x_2)) = f(\varepsilon, G(x_1, x_2))$

It would also be convenient, but not a requirement, if G and G' can be the same functions.

5.2 Function definitions

We shall define suitable functions for f based on the type of operation that we want to perform on x and then prove that each function satisfies the properties stated above.

5.2.1 Multiplication

If the function G to be performed is normal multiplication, we define f to be f_* where f_* is RSA encryption as defined in [13]. We use p and q to denote the primes, and N to be their product. We shall define ε to be the public key, and k to be the private key. Thus $f_*(x) = x^\varepsilon \bmod N$ and α' is obtained by $\alpha' = y_1 * y_2 * y_3 * \ldots * y_n \bmod N$.

Theorem 1 *f_* satisfies the properties defined in Section 5.1 where G is multiplication and G' is multiplication mod N.*

Proof:

1. This is trivial, from the usage of RSA for f.
2. This will hold if the modulus N used for the RSA encryption is larger than the largest possible x to be encrypted.
3. This follows from the homomorphic property of RSA:
   $((x_1{}^\varepsilon \bmod N) * (x_2{}^\varepsilon \bmod N)) \bmod N$
   $= (x_1{}^\varepsilon * x_2{}^\varepsilon) \bmod N$
   $= (x_1 * x_2)^\varepsilon \bmod N$

◇

Note that the final product α of all the unencrypted values has to be less than N. In practice N would be at least a 768 bit number, the current minimum value for reasonably secure RSA encryption.

5.2.2 Addition

If the function G to be performed is normal addition, we use a new public key cryptosystem as described in [8] for f. We define f to be f_+ where $f_+(\varepsilon, x) = \varepsilon^x \bmod N$ and ε is a generator for GF(N). N is the product of two large primes p and q as for RSA. For a full description the reader is referred to [8].

Theorem 2 *f_+ satisfies the properties defined in Section 5.1 where G is addition and G' is multiplication mod N.*

Proof:

1. If we examine $f = \varepsilon^x \bmod N$ it is clear that the encryption can easily be reversed if N can be factored to create the decryption key, or if x can be recovered by solving the discrete logarithm. Since factoring large numbers and solving large discrete logarithms are unpractical using current technology, we argue that f_+ satisfies the first condition.

2. This will hold if N is larger than the largest possible x to be encrypted.

3. This follows from the homomorphic property of the cryptosystem:
$((\varepsilon^{x_1} \bmod N) * (\varepsilon^{x_2} \bmod N)) \bmod N$
$= (\varepsilon^{x_1} * \varepsilon^{x_2}) \bmod N$
$= \varepsilon^{x_1+x_2} \bmod N$.

◇

Note that the final sum α of all the unencrypted values has to be less than N. In practice N would be at least a 768 bit number, the current minimum value for reasonably secure RSA encryption.

5.2.3 Comparison

Comparisons between encrypted values seem impossible, since ordering is lost through encryption, and if this were not the case, the encryption would be severely weakened. Assume an encryption function g preserves ordering, and that $y = g(x)$ is an encrypted value. Since we have to supply o with g, o can encrypt as many constants with g as needed, and a binary search will quickly determine the value of x. For a tax calculation, this can fortunately be circumvented by disclosing the intermediate values that require comparison. This will reduce the privacy of the tax return slightly, but this will still be an improvement over the unprotected return. See Table 2, where the total interest calculation is disclosed, but the various sources of interest are hidden for a concrete example.

The above then allows us to use this approach in a simplistic tax return, where various amounts are added and subtracted or multiplied and divided to calculate a taxable income.

5.3 Exhaustive searches

A possible problem with this method is that the exhaustive searches required to decrypt $y_1, y_2, y_3, \ldots, y_n$ would not take too long using computers, since the values of $x_1, x_2, x_3, \ldots, x_n$ are reasonably small compared to the computing power available today (usually less than 10^6, and almost always less than 10^{10}). The magnitude of this problem can be reduced by using a very big number of significant digits after the decimal point, and randomly adding a very small number (less than 1 cent) to each amount before the encryption function f is applied. This will not modify the taxable amount in any significant way, but should protect the amounts.

5.4 Decimal amounts

Although it seems that decimals can not be used due to the use of the modulus function, this can easily be overcome by only using cents for amounts (or millicents, picocents, etc. as applicable to prevent exhaustive searches).

6 Details of Class 2 implementation

When an individual i wants to send some private data m to a target organisation o, o's privacy policy P is checked using P3P. This policy is then used by i to give o a ticket granting ticket $T(o,i,P)$ granting access to m as described in the intersection between i's own privacy policy and P for a limited time. The actual private information is stored at a trusted third party S, for round the clock availability. When o needs part of m, o presents T as well as d, a description of the required subset of m and optionally o_2, (the other organisation who actually requires access to m in class 2.1) to S. S then verifies that P has not changed, and sends o a ticket t_{d,o_2}, where $o_2 = o$ if this was not supplied, and if allowed by P. A ticket can not be reused, and can only be used by the party it has been issued to. If o has to send private information to another party o_2, o has to request another ticket for o_2 from S if allowed by P. If o has to reuse information, a new ticket can be requested with T, unless T has expired. The information flow is depicted in Figure 1.

6.1 Properties of Tickets

In order to function as described above, tickets need certain properties, similar to the Kerberos implementation.

6.1.1 Security

The TGT and tickets must be tamper proof. The contents should not be modifiable by any party other than S. The necessity for this requirement should be obvious. Such security could be achieved by having S sign them with S's private key. Any party could verify the signature using S's public key. Information in the ticket and TGT that are only meant for certain o could be encrypted by S with their public keys, keeping it private.

6.1.2 Time limit on the ticket granting ticket

Some of the access granted on the TGT might be for a limited time only. This time limit should be visible to a target organisation o, but should not be modifiable by o. This is

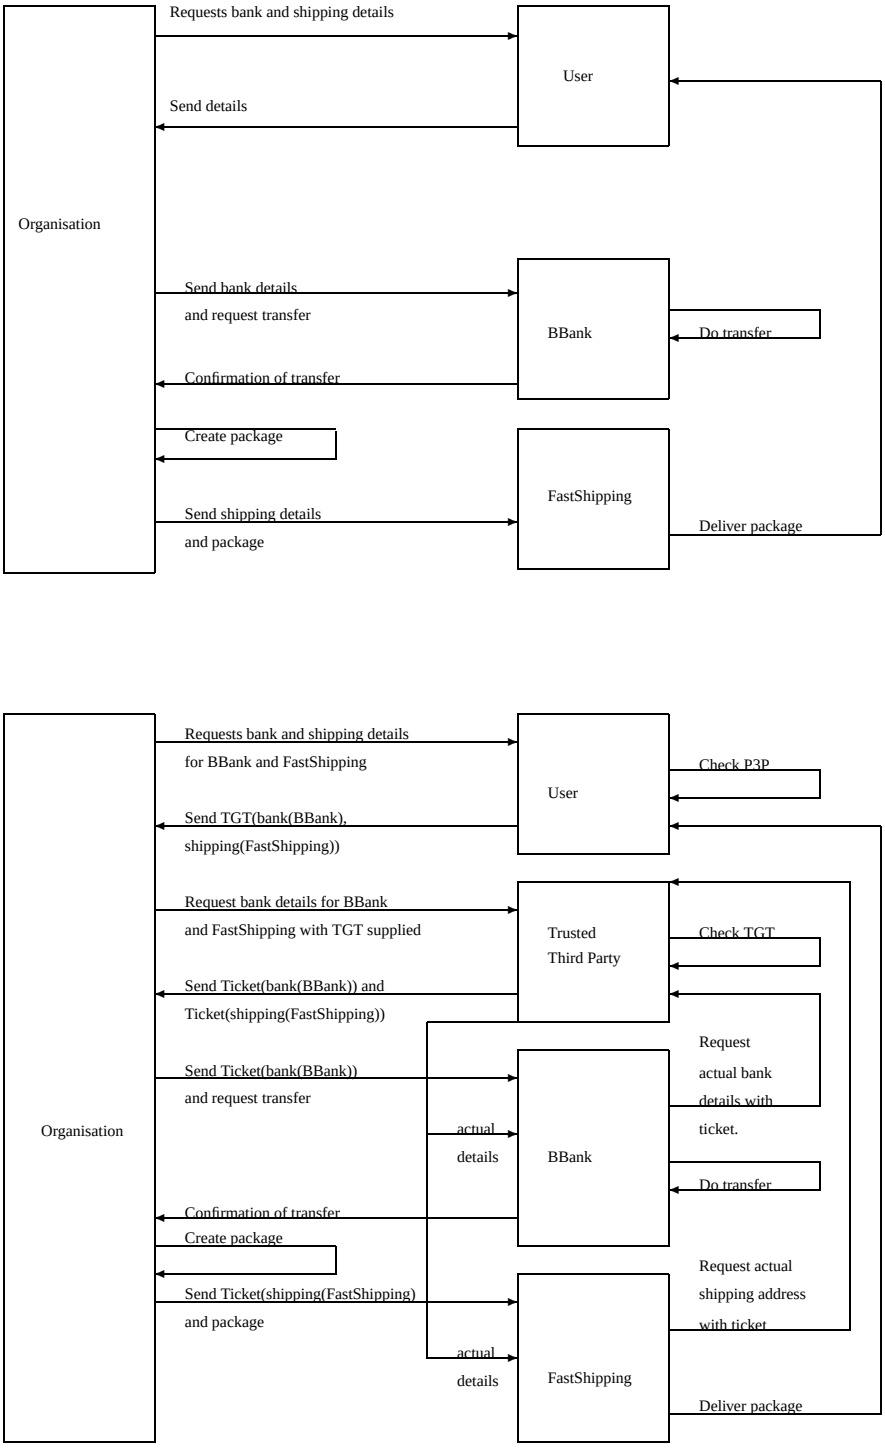

Figure 1: Graphical example of the protocol compared to normal situation

required so that *o* can ensure adequate time is allowed for shipping, billing, etc. and also in the case where *i* subscribes to a service requiring private information, *o* can ensure that the TGT does not expire before the subscription does. This requirement can be achieved by signing the time limit with *i*'s private key in cases where *i* set the limit, or *S*'s private key otherwise.

6.2 The Trusted Third Party

The trusted third party *S* forms an important part of this protocol, and if no such party exists other alternatives to a single trusted third party have to be found. We present a couple of examples.

The first solution is to integrate the trusted third party with the browser, so that an individual's browser stores the private information, and grants the TGTs and tickets. This is the option currently implemented in the P3P protocol, as the private information is stored by the browser. The biggest disadvantage to this method is that the information would not necessarily be available around the clock, and the organisation involved in a transaction might therefore not be adequately reassured of availability, for instance someone orders books from an online vendor, and disconnects from the Internet, making their shipping address inaccessible.

The second solution is to use more than one "semi-trusted" third party, each one only storing information that they can reasonably be trusted with, perhaps because they already have access to it. A good example might be to store billing addresses and financial information at a bank, since they already have that information. E-mail addresses could be stored at ISPs, and phone numbers at the phone company. The possible drawbacks here are that some element of private information can then be inferred by looking at the party storing it, such as who an individual's ISP is.

The third solution also uses more than one "semi-trusted" third party, but stores only part of each private information item at each party. A good example might be to store only every second number of a phone number at two semi-trusted parties, rendering the information useless without both parts. Access to the information is then required from all third parties involved before the information is actually usable. This increases both the complexity and the trust of the solution.

6.3 Collusion Between Participants

Another important question to consider is that of collusion between some of the parties, say between a big online retailer and its preferred shipping company. The same situation would occur if the retailer acquired the shipping company, and did not inform its customers of such an acquisition. Such collusion might then possibly be curbed by using several shipping companies, or by policing by the trusted third parties (they could refuse to give TGT's to such guilty parties for any individuals registered with them, which might make the possible repercussions of such collusion too severe for any *o* to risk). Further discussion of this issue might be the basis for future work.

6.4 Assurances

When making use of PrivGuard, the online retailer would like to be sure that the information requested is actually stored and available at the trusted third party. This can not be guaranteed, just like the retailer would not know whether a credit card number or address supplied actually exist in the real world when not using PrivGuard. Using this protocol could give the same level of assurance, if the trusted third party carried out the same validation checks that an online retailer would have been able to do. If the validation rests with the trusted third party, it can actually perform another validation that an online retailer would not be able to do — the trusted third party can check whether such information has been successfully used by another party before, and perhaps attach a confidence factor to any tickets issued.

6.5 Misuse of Protection

Finally, the better any individual's private information is protected and anonymised, the bigger the risk that such a system could be used for nefarious purposes (imagine someone ordering marijuana from the Netherlands, where it is legal, and having it anonymously delivered to the United States, where it is not. Interception of the package by customs might prevent delivery, but the recipient might not be identified).

In cases where a single *S* is used, a court order might possibly be used to reveal private information about *i* in cases where criminal intent can be proven. In such a case the expired TGT might be presented to *S* with the court order or warrant requiring disclosure. *S* can then supply the original information. This could be even more effective than just storing the individual's private information in the traditional way, as the information stored at *S* might be more up to date. However, if *S* is physically located in an area outside the jurisdiction of authorities requiring such disclosure, *S* might not divulge the private information. This can then be countered by *o* requiring non-expiring third-party access to the private information *m* on the TGT for such authorities. The individual's private information could then only be accessed by the relevant authorities, and not by *o*.

7 Examples

For more clarity we attach an example using class 1, and another using the second class.

7.1 Class 1

A fictional example using the South African tax law. For this example we shall set $N = 19697446673$ and $\varepsilon = 131$. All the **bold figures** can be kept private; the others, as well as all the encrypted figures are to be supplied to the IRS, along with the algorithms, f_* and f_* as well as N and ε.

Table 1 demonstrates how income and deductions are encrypted. Note that $18182403837 * 2608534625 * 7122548528 * 9973525707 * 847550577 \bmod 19697446673 = 6922828699$ and that 122229 encrypted for addition also yields 6922828699.

The interest sub-calculation is performed in Table 2. Note that the product of the encrypted values mod $N = 19697446673$ is 4224346997, the same as the encrypted value for 9365. We now disclose the total amount, since tax is not payable on the first (say) 2000 interest received.

Note that the totals can be verified as above.

7.2 Class 2

Let us use the case where Alice buys books from Bob, to be shipped using FastShipping. Alice's private details m consist of her name m_{name}, her payment details $m_{payment}$, her e-mail address m_{e-mail} and her shipping address $m_{address}$. Bob's privacy policy states that he requires a customer name for identification, payment details for a once-off payment for the order, as well as a shipping address for a once-off shipment of the order. He would also like Alice's e-mail address, to notify her of specials, and would like to distribute it to book clubs and other customers for reference purposes. Alice's privacy policy allows all of the above, except that she does not want her e-mail address distributed to others, but would like to be notified of specials. Alice now gives Bob a TGT T_1 allowing him access to $m_{name}, m_{payment}$ and $m_{address}$ for 7 days, limited to one transaction, and allowing him access to m_{e-mail} for 3 months. T_1 also limits the use of $m_{payment}$ to Bob as beneficiary, and $m_{address}$ to FastShipping. Bob now presents T_1 to our trusted third party S, and requests a ticket t_p for payment using Bob's bank, BBank, as well as a ticket t_s for shipping using FastShipping. S verifies the validity and policies of T_1, and then issues t_p and t_s. Bob now sends t_p to BBank, requesting payment. BBank presents t_p to S, who supplies BBank with Alice's credit card information for the transaction. BBank notifies Bob of the successful transfer. Notice that Bob never knew Alice's credit card details. Bob now sends the package and t_s to FastShipping, who presents t_s to S to get Alice's shipping address, and deliver her books. Note that Bob also did not know Alice's shipping address. For the next three months, Bob can send mail to Alice by requesting a ticket from S with T_1. This ticket is then sent with the message to a trusted messenger service (which might also be S), who will then forward it to Alice. When T_1 expires, S will no longer issue tickets to Alice's e-mail address. The same will happen if T_1 is presented to S by anyone other than Bob, or if Bob requests a ticket for another recipient. Any attempt reuse the payment information contained in T_1 during its 7 day validity for another transaction will also be rejected by S. So, Bob can not reuse or redistribute the payment information, unless BBank conspires with him (unlikely — banks are all about trust. If they can not be trusted, public scorn will soon force them to close). Bob also can not reuse the shipping address for similar reasons, unless FastShipping conspires with him (a possibility, especially if Bob does a lot of business with FastShipping, and in reality there are not that many shipping companies with the ability to ship world wide at reasonable rates. This could be prevented by sending the package to S, who can forward it to Alice for a nominal fee, or to use a series of shipping companies, each knowing only the next step in finally delivering the package).

8 Related Work

Previous efforts have been made to improve the privacy of information used in online transactions, but they all have flaws making them either under utilised or of limited worth. We briefly discuss some of these initiatives.

Namesafe.com [9] provides a service with certain partners where a unique username is generated for every user registering with them. The namesafe username and details is then used when making a purchase. This provides good protection of identities, shipping addresses and e-mail addresses, but not anything else. Packages are shipped to a Mail Boxes Etc. location, from where it can be picked up. This limits the usefulness to locations within a reasonable distance of a Mail Box Etc. collection point. Adding protection for say phone numbers would require the addition of another partner, or that an existing partner expands the service they offer. Although the shipping address is hidden from the target organisation, any subsequent packages sent will also reach the user's Mail Box, so that junk mail can still reach the user. This is similar to using a post office box for packages — your home address remains safe, but you can still receive junk.

iPrivacy.com [4] offers a service in partnership with credit card companies. Their service encrypts part of the private information, but leaves other parts, such as the city, state and zip code open. The information that is encrypted, such as the address, can the be decrypted by the delivery company. This means that the information could still be out of date if the individual moves, as the encrypted version of an address is stored by the target organisation, instead of a TGT granting right of access to the information. The user can not limit the time that the organisation has access to address — as long as the delivery company retains the decryption key, subsequent submissions for deliveries is still possible, without the user being able to revoke such access. iPrivacy.com's software is obtained from an individual's credit

	Description	*Clear*	*Encrypted*	*For (A)ddition or (M)ultiplication*
	Salary	**110000**	18182403837	A
less	Pension Fund Contribution	**-9000**	2608534625	A
plus	Car Allowance	**40000**	7122548528	A
plus	Taxable Interest	7365	9973525707	A *From Table 2*
less	Travel expenses	-26136	847550577	A *From Table 3*
	Total taxable income	122229	6922828699	

Table 1: Income and deductions for example

	Description	*Clear*	*Encrypted*	*For (A)ddition or (M)ultiplication*
	Unit Trusts	**3556**	5403644383	A
plus	Bank account	**345**	15350080410	A
plus	Fixed Deposit	**5432**	11953675544	A
plus	Share trading account	**32**	11140409637	A
	Total	9365	4224346997	

Table 2: Interest sub-calculation

card company, which, of course, makes it inaccessible to people without credit cards. As of March 2002, no credit card companies have yet signed up with iPrivacy, making it still a future service.

Secure Electronic Transactions (SET) [14] is an initiative to prevent the misuse of credit card numbers for online transactions by encrypting the credit card and purchase details so that the organisation requesting payment can not access the credit card number, and that the payment instruction can not be reused. This protection is only applicable for credit card details, and does not protect any other aspect of an individual's privacy. SET, though heavily backed by credit card companies, has failed to gain acceptance as users have no real incentive to campaign for its use. The risk generally lies with either the merchant or the credit card company when a credit card number is misused — in typical "card not present" transaction disputes it is hard to prove that the card owner authorised the transaction, so the card owner does not suffer the loss. Embedding banking sites within an organisation's website also makes SET redundant, as the organisation would not have access to the user's banking details.

A quick comparison can be seen in Figure 2.

PrivGuard adds significant enhancements to the solutions mentioned here. First of all, by using a trusted third party to store the actual information, only one copy of each information item is kept, enhancing the integrity of the data. This protocol is applicable to any information item the user might have, even though some measure of privacy protection is only added to classes 1 and 2.1 – 2.3. This means that all a user's private information can be routed through this protocol, making a seamless integration with the browser possible, as with P3P. Finally the user retains a degree of control over the information, by being able to set a time limit on the access to the information. The biggest weakness of PrivGuard is that it requires a trusted third party, but this can be address as discussed in Section 6.2. The only remaining stumbling block is that PrivGuard increases the complexity of online transactions, as the protocol requires more steps as already depicted in Figure 1.

9 Summary

We have presented a classification of private information based on the purpose for which it is acquired, and created a protocol to protect private information in several of these classifications.

With this protocol we have effectively prevented unauthorised reuse and redistribution of private information in all cases where the target organisation *o* did not require direct, unprotected access to an individual *i*'s private data. In such cases the confinement problem has also been sidestepped, as the actual information has not been divulged. We have also managed to protect private information by disclosing it at the last possible stage in any transaction, virtually preventing all intermediaries from accessing our private information. However, note that this protocol does not prevent the last stage organisation to store and misuse the actual private information. The impact of this is lessened by the fact that this last stage usually has very little information, which might not be useful per se and the fact that in e-commerce applications these last stage organisations will typically be very large, such as banks and shipping companies, with a lot to lose should they misuse private information.

Further study could possibly integrate this approach

	Description	*Clear*	*Encrypted*	*For (A)ddition or (M)ultiplication*
	Fixed allowance per 1000 km	**1022**	4201531621	A
plus	Fuel allowance per 1000 km	**226**	7357125748	A
plus	Maintenance allowance per 1000 km	**204**	15157866258	A
	Total per 1000 km	1452	18998024939	
	Total per 1000 km	1452	7632560095	M
times	Business km travelled (in 1000)	**18**	10205764354	M
	Total deduction	26136	12475071519	M

Table 3: Travel expenses sub-calculation

Product	Namesafe	iPrivacy	SET	PrivGuard
Prerequisites	US resident	Credit Card	Credit Card	Trusted Third Party
Scope of protection	Medium	Medium	Narrow	Wide
Implementation	Practical	Future	Practical	Theoretical
P3P Integration	No	No	No	Possible
Users retains control	No	No	N/A	Yes
Up to date data	No	No	Yes	Yes

Figure 2: Comparison of privacy solutions

with P3P, in order to automate this process and make it totally transparent to the end user.

References

[1] Niv Ahituv, Yeheskel Lapid, and Seev Neumann. Processing encrypted data. *Communications of the ACM*, 30(9):777–780, 1987.

[2] Roger Clarke. Internet privacy concerns confirm the case for intervention. *Communications of the ACM*, 42(2):60–67, 1999.

[3] Lorrie Faith Cranor. Internet privacy. *Communications of the ACM*, 42(2):29–31, February 1999.

[4] iPrivacy home page. http://www.iprivacy.com.

[5] John Kohl and B. Clifford Neuman. The Kerberos network authentication service. Technical report, 1991.

[6] Frans A. Lategan and Martin S. Olivier. Enforcing privacy by withholding private information. In S. Qing and J. H.P̃. Eloff, editors, *Information Security for Global Information Infrastructures*, pages 421–430. Kluwer, August 2000.

[7] Frans A. Lategan and Martin S. Olivier. On granting limited access to private information. In *The tenth international World Wide Web conference on World Wide Web*, pages 21–25. ACM Press, 2001.

[8] David Naccache and Jacques Stern. A new public-key cryptosystem. In *Theory and Application of Cryptographic Techniques*, pages 27–36, 1997.

[9] Namesafe home page. http://www.namesafe.com.

[10] B. Clifford Neuman and Theodore Ts'o. Kerberos: An authentication service for computer networks. *IEEE Communications Magazine*, 32(9):33–38, September 1994.

[11] Joseph Reagle and Lorrie Faith Cranor. The platform for privacy preferences. *Communications of the ACM*, 42(2):48–55, 1999.

[12] The Kerberos network authentication service. http://www.isi.edu/gost/gost-group/products/kerberos/.

[13] Josef Seberry and Jennifer Pieprzyk. *Cryptography: An Introduction to Computer Security*. Prentice Hall, 1989.

[14] SET home page. http://www.setco.org.

[15] Bhavani Thuraisingham, Sushil Jajodia, Pierangela Samarati, and Martin S Olivier. Security and privacy issues for the World Wide Web: Panel discussion. In Sushil Jajodia, editor, *Database Security XII: Status and Prospects*, pages 269–284. Kluwer, 1999.

[16] Platform for privacy preferences (P3P) project. http://www.w3.org/P3P/.

[17] Huaiqing Wang, Matthew K. O. Lee, and Chen Wang. Consumer privacy concerns about Internet marketing. *Communications of the ACM*, 41(3):63–70, 1998.

F. A. Lategan and M. S. Olivier, "PrivGuard: A model to protect private information based on its usage," *South African Computer Journal*, 29, 58–68, 2002.

Preprint
Source: http://mo.co.za