Enforcing Privacy on the Internet

by

Frans Adriaan Lategan

Thesis

submitted in the fulfillment
of the requirements for the degree

Philosophiae Doctor

in

Computer Science

in the

Faculty of Natural Sciences

at the

Rand Afrikaans University

Promoter: Prof MS Olivier

May 2002

UNIVERSITY
OF
JOHANNESBURG

Contents

Acknowledgements

I would like to thank the following people for their support during the creation of this work:

- My promoter, Prof. Olivier for his guidance and support.
- My friends and family, especially my wife, Annalize for their support and encouragement.
- My Creator for the abilities, patience and perseverance.

Abstract

Privacy of information is becoming more and more important as we start trusting unknown computers, servers and organisations with more and more of our personal information. We distribute our private information on an ever-increasing number of computers daily, and we effectively give target organisations carte blanche to do what they want with our private information once they have collected it. We have only their privacy policy as a possible safeguard against misuse of our private information. Thus far, no reliable and practical method to enforce privacy has been discovered.

In this thesis we look at ways to enforce the privacy of information. In order to do this, we first present a classification of private information based on the purpose it is acquired for. This will then enable us to tailor protection methods in such a way that the purpose the information is acquired for can still be fulfilled. We propose three distinct methods to protect such information.

The first method, that of nondisclosure, is where private information is required not for the contents, but as input to verify calculations. We shall present an encryption method to protect private information where the private information consists of a set of numeric values S on which some function G has to be applied and the result $\alpha = G(S)$ has to be supplied to a target organisation. The calculation of the result α must be verifiable by the target organisation, without disclosing S.

The second method, that of retaining control is a method by which we can grant limited access to our private information, and thus enforce the terms of privacy policies.

The final method we present is a conceptual method to extend P3P in order to add more flexibility to the decision on whether or not a given item of private information will be supplied to a target organisation by using the Chinese Wall security policy. This will enable a user to not only define rules as to which items of private information he would disclose, but also to define what collection of private information any given organisation would be able to build about him.

Afrikaanse Oorsig

Die privaatheid van inligting raak al hoe belangriker soos wat ons onbekende rekenaarstelsels, bedieners en organisasies begin vertrou met ons private inligting. Ons versprei ons private inligting op al hoe meer rekenaarstelsels daagliks, en in werklikheid kan die organisasies doen net wat hulle wil met ons inligting sodra hulle dit bekom het. Ons het slegs hulle privaatheidsbeleid as moontlike beskerming teen die misbruik van ons private inligting. Tot dusver is daar geen betroubare metode om die privaatheid van inligting af te dwing nie.

In hierdie proefskrif kyk ons dus na die ontwikkeling van metodes om privaatheid meer afdwingbaar te maak. Om dit te bewerkstellig stel ons eers 'n klassifikasie van private inligting voor wat gegrond is op die doel waarvoor sulke inligting bekom word. Dit sal ons in staat stel om beskermingsmeganismes so te ontwikkel dat die doel waarvoor die private inligting bekom is steeds haalbaar is. Ons stel drie verskillende metodes voor om private inligting mee te beskerm.

Die eerste metode is waar private inligting nie onthul word nie, omdat dit slegs benodig word om die resultaat van 'n berekening na te gaan. Ons lê 'n enkripsie metode voor om private inligting te beskerm waar sulke inligting bestaan uit 'n versameling numeriese waardes S waarop 'n funksie G toegepas word en die resultaat $\alpha = G(S)$ moet dan aan 'n teiken organisasie verskaf word. Die berekening van die resultaat α moet deur die teiken organisasie nagegaan kan word sonder dat hulle toegang tot S verkry.

In die tweede metode probeer ons om beheer van die private inligting te behou, om sodoende die bepalings van 'n organisasie se privaatheidsbeleid af te dwing.

In ons laaste metode beskryf ons 'n manier om P3P sodanig aan te vul dat die besluit om private inligting aan 'n organisasie te verskaf gebaseer kan word op die Chinese Muur sekerheidsbeleid. Dit sal 'n gebruiker in staat stel om nie net reëls te bepaal waarvolgens individuele private items vrygestel word nie, maar om ook die versameling van private items wat 'n organisasie van hom kan versamel te beperk.

Chapter 1

Introduction

This first chapter serves as a road map for the thesis. It contains the problem statement, defines the scope of the research and describes the layout of the rest of the chapters.

A journey of a thousand miles must begin with a single step
Lao-tzu
The Way of Lao-tzu

1.1 Background

In our daily lives we constantly need to divulge some of our private information to do business. Past earnings and loan repayment history might be necessary to apply for a home loan. A criminal record is kept on a police database somewhere, and medical records must be accessible for accurate treatment. In order to pay our taxes, we must supply the Internal Revenue Service (or similar government body) with detailed information on our financial dealings. They use that information to calculate the amount of tax we must pay.

This is the way it has always been, and no-one has complained about privacy in the past, but with the advent of networked computerised files things have changed somewhat. When our private information is stored in a hard copy format where very few people have physical access to it, it is reasonably secure. However, the same private information can be several orders of a magnitude more vulnerable to unauthorised access when stored electronically, since millions of people can potentially access it should it not be protected.

These days private information is stored in electronic files, and those files can be sold to advertisers, placed on the Internet for public viewing, hacked into by malicious hackers or even browsed by curious employees. The latter is particularly hard to prevent, since traditional access control is of little use where people are supposed to have access to private information about others to do their work, but should not be allowed access to the same information when motivated by curiosity.

1.2 Problem statement

Current privacy protection mechanisms rely largely on ethical and legal safeguards; very little of this can be enforced before a privacy violation to prevent such a violation. In this thesis we aim to remedy this by allowing the owner of private information to actually enforce the privacy of his private information. A compromise between total privacy (and an impossibility to transact on the Internet) on the one hand, and total disclosure of private information (and a total loss of privacy) has to be reached. For a more effective compromise between these extremes, we shall develop a classification of private information based on the purpose it has been acquired for, so that protection

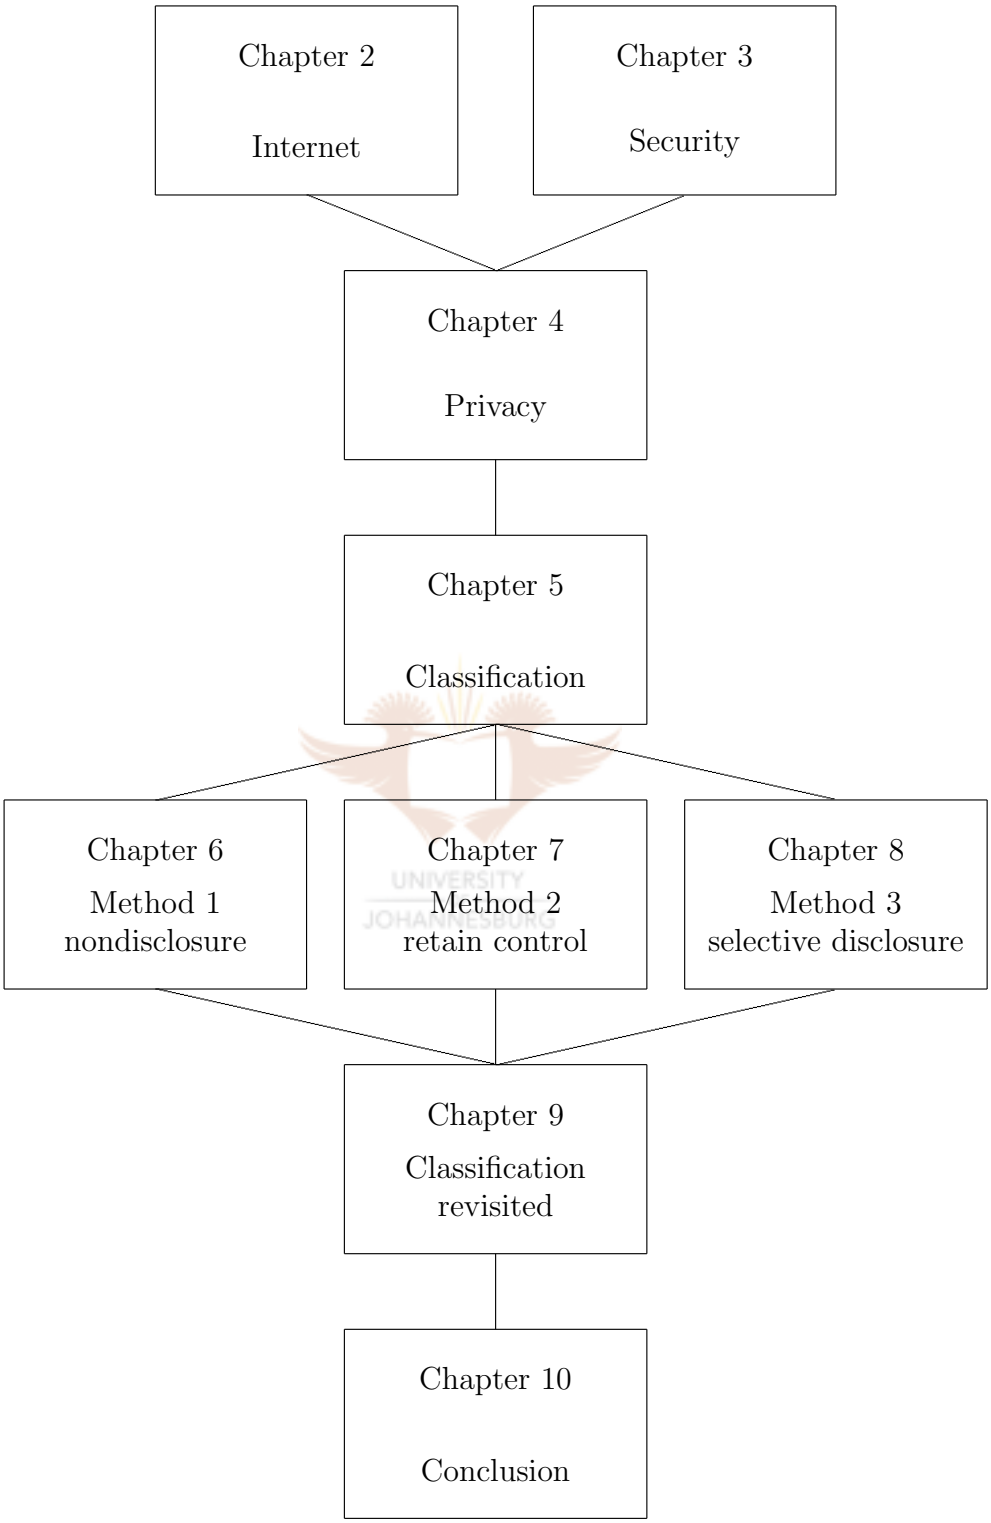

Figure 1.1: Chapter layout

mechanisms can be effectively tailored to protect the privacy of data without invalidating the purpose it has been acquired for in the first place.

With this classification, we shall then use three possible approaches to improve existing privacy protection mechanisms. Our first method, namely nondisclosure of private information will use cryptographic techniques. Our second method, retaining control over private information that is disclosed as far as possible will be developed using a model based on trusted third parties. In our final method, selectively disclosing private information in order to prevent any one organisation from learning too much about one individual, we shall apply the Chinese Wall security model to the P3P protocol.

In this thesis we shall not be looking at extending ethical and legal means of protecting privacy.

1.3 Roadmap

The rest of the thesis is organised as depicted in Figure 1.1.

Chapter 2 gives background information about the Internet, its history, growth and technology. Chapter 3 contains some background information on computer security and the security tools we shall use in the rest of this work. Chapter 4 then links the Internet and security together to discuss the state of privacy on the Internet.

In Chapter 5 we develop a classification of private information which we then use in the following chapters to tailor privacy protection models: in Chapter 6 we discuss nondisclosure of private information, followed by ways to retain control over private information in Chapter 7 and selective disclosure of information in Chapter 8.

We then take a second look at our classification in Chapter 9 and finish the thesis with our conclusion in Chapter 10.

Chapter 2

Internet

The Internet is the domain in which we shall investigate the privacy of an individual's information. As such, it is important to be familiar with the relevant aspects of the Internet, and its phenomenal growth. In this chapter we shall present a general overview of the Internet and where it came from, its history and some of the technologies in current use to provide content and functionality.

No man is an Island, entire of itself;
every man is a piece of the Continent, a part of the main
John Donne
Devotions XVII

2.1 Introduction

The Internet has existed before the World Wide Web and Internet Explorer, and before nice graphical views, Domain Name Servers (DNS) and search engines. People who only recently started using the Internet (and based on the exponential growth of the Internet, at least half of what is available on any given date did not exist six months before) would scarcely recognise it as it was five years ago, let alone thirty.

The exponential growth of the Internet has a very important effect on the privacy of information. When our private information is stored in a hard copy format where very few people have physical access to it, it is reasonably secure. However, the same private information can be several orders of a magnitude more vulnerable to unauthorised access when stored electronically, since millions of people can potentially access it should it not be protected.

2.2 History

When we talk about the Internet, we are usually referring to the World Wide Web (WWW). The other portions such as Gopher, Telnet and FTP have taken the back seat to the ubiquitous WWW (WWW overtook Gopher and Telnet in 1994, and FTP in 1995 [68]). The WWW started when Tim Berners-Lee created the first browser on the NeXTStep platform, and developed it from there. An early article [3] highlights the state of WWW in 1994. It is important to note that WWW was not designed with privacy in mind at all, but to promote collaboration between researchers working in the same field. The full history is outlined in RFC 2235 [68].

2.2.1 Growth

The World Wide Web has grown at an exponential rate — see Figure 2.1 (a logarithmic scale was used). The numbers were taken from the ISC's domain survey [27] and RFC 2235 [68]. Note that the method used to measure the number of hosts was changed slightly in January 1998.

Through the years, as the number of people accessing the Internet has continued to grow exponentially, and their demographics have changed from English speaking academics to technologically savvy young people to the general population, the primary purpose that they use the Internet for has

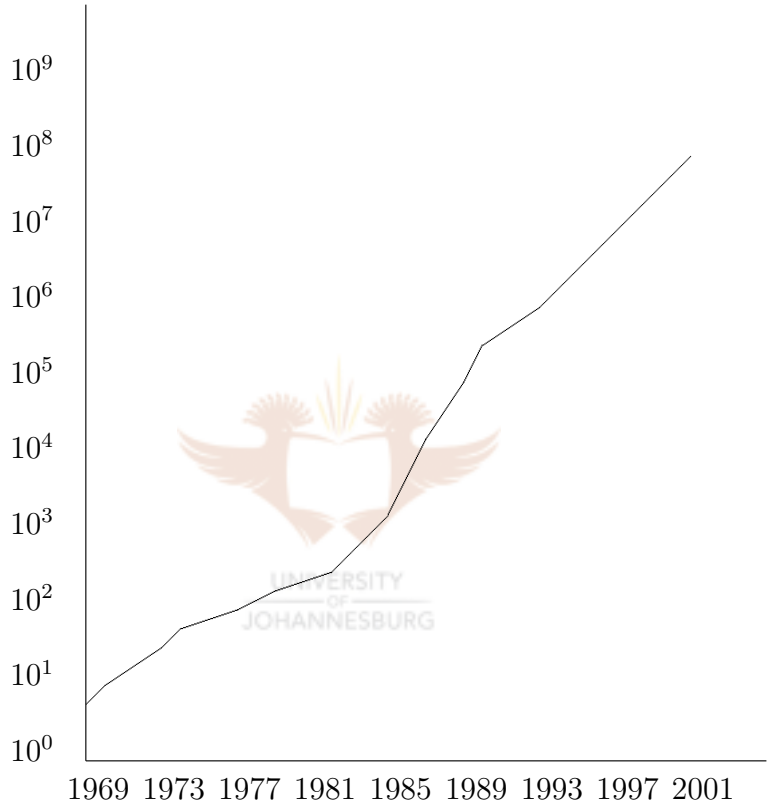

Figure 2.1: Growth of Internet Hosts (The numbers were taken from the ISC's domain survey [27] and RFC 2235 [68].)

also shifted. At the start the Internet was mainly used for the exchange of academic information, but today the focus is more on general information, news, commerce and entertainment. With this shift, the requirements have changed from an open system accessible to all to a quasi open system, where some access control is applied. In 1990 the Internet was a one-way flow of information, from the servers to the users (the PUT part of HTTP was very limited [3, p. 78]). Currently people are submitting so much personal information on the Internet that privacy concerns are multiplying [10].

Finally we have now reached a point where we have to start rethinking the design of the Internet [6], as more and more issues come up that were never a factor when it was originally designed. The privacy of data on the Internet will be discussed in more detail in Chapter 4. The rest of this chapter will give a quick overview of the functional components of the Internet, as well as future trends.

2.3 Functional Components of the Internet

All of these functional components are described on the W3C organisation's website, http://www.w3.org which will contain the latest incarnation of each as these components as they continue to evolve.

2.3.1 URI, URL and URN

A Uniform Resource Identifier (URI) is a compact string of characters for identifying an abstract or physical resource [5]. It can specify a locator (URL), a name (URN) or both.

A Uniform Resource Locator (URL) is an extensible way in which every resource on the Internet can be located. The first part of a URL specifies the protocol or scheme used to access a specific resource, and the second part of the URL specifies the scheme-specific information. URL schemes that use the IP protocol directly contain the location of the resource in the scheme-specific part of the URL. The location section of a URL consists of a server address (either in name or dotted IP form) followed by an access path on the server to reach the resource.

Optionally, the port used for access can be specified if different from the standard, and a user name and password can also be supplied along with the access. Note that specifying the user name and password on the URL is not secure, which is why it is rarely used.

An example of a URL is

http://myuser:mypassword@www.w3.org:80/Consortium/

where user "myuser" with password "mypassword" will use the http protocol on port 80 to access the Consortium page on the www.w3.org server. More information on URLs can be found in RFC 1738 [4].

A Uniform Resource Name is the subset of URIs that have to remain unique and exist even if the resource pointed to no longer exists. A URN is a persistent labelling of a resource with an identifier, such as

urn:foo:a123,456

and has to start with "urn:". The full URN syntax is defined in RFC 2141 [38].

2.3.2 HTTP

The Hypertext Transfer Protocol (HTTP) protocol is used to access Internet content, and defines the exact requests and responses used. The server specified by the URL is contacted, with the header specifying the type of request. A response is then sent back, containing a message header and optionally a message body. A status code details the nature of any error that might have occurred. The HTTP protocol is fully specified in RFC 2616 [20].

The HTTP protocol also specifies the way in which proxies, caches and pipelines can be used to filter content, improve performance or reduce lag time. For web browsing, the response often consists of an HTML file.

2.3.3 HTML

The Hypertext Markup Language (HTML) is the standard for publishing documents on the Internet. HTML was originally conceived for the exchange of scientific and technical documents, but hypertext and multimedia support were added later. Unfortunately HTML outgrew its original purpose, and there is a drive to replace HTML with XHTML, a more rigourous markup language, based on XML. The first version of XHTML, XHTML 1.0 [43] is just a restatement of HTML 4.01 [45] in XML terms.

This will finally allow the document publishing standard to expand as rapidly as the content being deployed using it. Eventually any XHTML compliant content will be usable on any XHTML compliant platform. XML's extensibility will enable new tags to be added without creating incompatible versions.

2.3.4 Cookies

Internet sites can use cookies to maintain state about a user, such as a shopping cart or a wish list. Cookies have originally been used for such in-

nocuous uses, and as a cookie can only be read by the domain that set it, no privacy impacts were expected. Unfortunately embedded banner advertisements could write cookies accessible from their home server, but could also often access data on the page they were embedded on, resulting in a leak of private information. Sites such as Double Click [16] now have special opt-out cookies that use a generic user id in order to invalidate such information collected, if the user chooses to opt out. The cookie mechanism is fully specified in RFC 2965 [30].

2.3.5 XML

The Extensible Markup Language (XML) is a markup language that can be extended by using self-describing tags. It is more rigourous than HTML; for example tags have to be closed in the reverse order that they have been opened, to form a nested structure. It allows validation to be performed by using Document Type Definition (DTD) files, or by using schemas.

The biggest advantage that XML has over HTML is that it is possible to extend the markup language by defining custom tags and DTDs. This allows different knowledge domains to standardise on the way in which they wish to exchange information, and removes the ambiguities present in HTML that complicates machine readability. XML documents are searchable, and modern databases can store them directly, whilst still allowing querying on them.

XML is still evolving; XML Schema was approved as recently as May 2001. The up to date standards can be found at [61].

2.4 E-Commerce

With the wider acceptance and penetration of the Internet, it is being used more and more for e-commerce. In the late 1990's a lot of hype was attached to e-commerce, and the dot com revolution caused investors to sink millions of dollars into e-commerce. The dot com bubble eventually burst, as most e-commerce ventures were not built on sound business principles, but expected that the exponential growth of the Internet would cause exponential growth in profits too. Unfortunately their costs usually grew even faster than their revenue. Many startups went from broke to millionaires back to broke within a very short time.

Despite these setbacks, e-commerce is generating more revenue every year, as more and more people are doing business over the Internet. Since the world wide web was not originally designed for e-commerce, the protocols had to

evolve quickly to cater for this need. Although security did receive some attention, it is only recently that privacy concerns are starting to surface.

2.5 Global availability

2.5.1 Accessibility

As the Internet started moving away from the academic community and the demographics of Internet users started mirroring the demographics of the outside world, accessibility to people with disabilities became more important. Currently the World Wide Web Consortium (W3C) has accessibility guidelines that make sites more accessible [64]. Things like text-to-speech devices and voice recognition can make web sites available to the visually impaired. Sites that are useable without fancy graphics or multimedia are also more accessible to users browsing with phones or personal digital assistants (PDAs).

2.5.2 Internationalisation

Availability also has to do with language. At the start of the Internet, most of the content was in English, but this is rapidly changing. More and more content is being developed in languages other than English. Even the ASCII character set can no longer be taken for granted, as languages such as Arabic and Chinese require their own character sets. Computer translators are getting better, and hopefully in the future they will improve to a level where the language that a site was originally published in won't matter anymore. As an example, the following paragraph has been translated by Voila [59] from English to French, and then back to English again.

Original English Computer translations are getting more sophisticated by the day and can currently make the contents of a web page in an unknown language understandable. Unfortunately no computer translation algorithm is perfect.

French translation Les traductions d'ordinateur deviennent plus sophistiquées et peuvent actuellement faire le contenu d'une page Web en langue inconnue compréhensible. Malheureusement aucun algorithme de traduction d'ordinateur n'est parfait.

English translation The translations of computer become more sophisticated and can at present make(do) the contents of a web page in under-

standable unknown language. Regrettably no algorithm of translation of computer is perfect.

The W3C also has guidelines for internationalisation [65].

Most of the future growth of the Internet will be in languages other than English. This is obvious from the fact that a very large percentage of the English-speaking world is online, whereas only a fraction of the Chinese-speaking world is online (and significantly more people speak Mandarin or Cantonese than English). With China's phenomenal economic growth and its huge population, a significant portion of future growth has to come from China, with Chinese content.

2.5.3 Device Independence and Java

Although most early adopters accessed the Internet from PCs or workstations, the Internet was designed from the start to function across platforms, and this has continued. It is now possible to access the Internet from cellular phones and personal digital assistants (PDAs) and applications written in Java and distributed over the Internet can run independent of the operating system, as long as a virtual machine is available.

Household appliances and cars, airport lounges and in airplanes, set-top boxes and consoles — the Internet is available everywhere. Even phone lines are not required as wireless access continues to mature.

All of this means that information published on the Internet becomes accessible to more and more people every day, and that privacy needs are therefore expanding as the statistical likelihood of private information being disclosed increases.

2.6 Summary

The Internet, a relative unknown concept as little as ten years ago, has penetrated the lives of millions of people, and is continuing to grow. Initiatives like internationalisation, accessibility and extensibility make this growth possible. As more sites are available, and they offer more functionality, more and more people are going to be doing more of their business online. Due to this rapid growth, the standards had to evolve so quickly that little if any thought has been given to privacy concerns.

With people putting their private information on more sites, and more people online, the threat to privacy is still growing. Comparatively little thought is being given to keep private information safe, as efforts are directed to solving the current growing pains, such as limited bandwidth, limited IP

addresses, too few high level domains and denial of service problems (both deliberate and accidental). Although these are valid concerns, security and privacy also require a closer look.

Chapter 3

Security

This chapter gives some background information on the most important security concepts and tools used later on to safeguard the privacy of a user's private information.

They shall not pass.
R.–G. Nivelle
Order, 1916, used as slogan throughout the defence of Verdun

3.1 Introduction

Security in computer systems has a relevance on the privacy of an individual's information in that it supplies us with the tools and models to effectively protect such private information. In this chapter we shall present a general overview of the security tools that we shall use in the rest of the work. For a more detailed introduction to computer security the reader is referred to Pfleeger's book [44].

3.2 Properties of secure data

For any system or data item to be considered secure, it must have three important properties, namely secrecy, integrity and availability. These properties might seem conflicting but without all three of them the data to be protected loses its value and purpose.

3.2.1 Secrecy

The secrecy property of the data requires that no access is allowed to the information unless expressly granted. This implies that the information is being kept secret from everyone not specifically allowed access to it. The biggest stumbling block to this property when viewed from the Internet user's perspective is that granting legitimate access to data usually implies sending the data to another party, where such a party then stores that data on computer systems neither known to or controlled by such a user. This means that although access to information might be revoked in a conventional system, it is almost impossible to revoke access to data used for transacting over the Internet.

An example of this might be a user's credit card details, which are generally stored in online vendors' systems after requesting it the first time. Until such a user has transacted on the Internet, her credit card details can only be obtained from her bank, or herself. As soon as she transacts on the Internet, every vendor she has used could also possibly be a source (that she has no control over) of her credit card details. Even if she were to contact a vendor and request that they delete her details, she has no way of ensuring that they have in fact done so.

It seems that the secrecy property can only be guaranteed up to the point where at least one other party requires access to such information.

Although the other party might maintain the secrecy of the data, there are no guarantees to that effect. With each additional party that has access to private information, the risk of accidental or intentional disclosure increases, as the data is only as safe as the weakest of all systems it is stored on.

3.2.2 Integrity

The second property of secure data is integrity. This means that as long as legitimate access is allowed to the information, such information must be accurate. Once again the Internet presents a major stumbling block to this property, namely that once a user has sent her information to another party, and that party has stored such information, the integrity of the information has been lost. Any subsequent updates to this information will not automatically apply to the copies stored at various locations over the Internet, unless the user kept a record of each organisation that stored a copy of it, and notifies them individually. Even then the user has no method of ensuring that they comply with her request to update the information.

The delivery address of an online book purchase is a good example of this problem. Should the customer change his address after ordering the books, but before delivery, the copy of his address at the vendor will not be correct unless such a vendor is expressly notified of such a change. Even if the package has not yet been shipped, the shipping label (another copy of the information) might already have been printed, resulting in an incorrect delivery.

The integrity of information used on the Internet seems to be comprised just like secrecy as soon as access to such information is granted to (the information is supplied to) another party on the Internet.

3.2.3 Availability

The last property of secure information that we consider is that of availability. This requires information to be available for access if such access has been granted. In the context of commerce on the Internet this is usually not a problem as any party being granted access to information generally has a locally stored copy, and as long as their systems are available, the information sent to them should also be available. However this means that the availability of the information is only as good as the system of the organisation it is stored on — there is no way for the owner of the information to improve availability through conventional means such as parallel servers or by storing it on high availability servers.

The availability of information used on the Internet is generally very good, but as it is not under control of the user granting access to it, there is no way for the user to verify at any given point that some organisation given access to information can in fact access that information.

3.3 Encryption

In the Internet world, secrecy is most often obtained by using encryption. Several types of encryption are combined to attempt to ensure the secrecy of information in transit. (Note that the point concern expressed above about secrecy once the information has been received by the second party is still valid, despite the fact that such information might have been secret whilst in transit).

Encryption is the process in which information is transformed from an intelligible format (clear text) to an unintelligible format (cypher text) by the application of a mathematical function E (encryption algorithm). Decryption is the process by which the original clear text is recovered from the cypher text by the application of another mathematical function D. Note that the encryption function E could be the same as the decryption function D but this does not have to be the case.

The mathematical functions used require as input the information as well as a seed value (the key). In order to recover the clear text from the cypher text, both the function used as well as the key is required. More formally, we can represent encryption of a message m with key k yielding cypher text M as $E(m, k) = M$. Similarly, decryption with key k' is $D(M, k') = m$.

Once again the encryption key k could be the same as the decryption key k'. In such cases the encryption is called symmetric key encryption. If $k \neq k'$ the encryption is called public key encryption. Most current encryption algorithms use the same function for encryption and decryption so that $E(E(m, k), k') = m$.

Historically, secrecy was assured by keeping both the encryption algorithm and the key secret, but since strong encryption algorithms are notoriously hard to design, extensive peer review the only way to provide a measure of strength. Therefore the current practice is to use known algorithms that have been vetted by experts, and only keep the keys secret. Once again it is useful to keep in mind that just because several experts have failed to find a vulnerability in an encryption algorithm does not imply that no vulnerability exists. A very good general introduction to encryption and history of both encryption and cryptanalysis can be found in Singh's book [55].

We now briefly discuss some commonly used types of encryption, namely

Group size	Keys required
2	1
3	3
5	10
10	45
20	190
50	1225
100	4950

Table 3.1: Number of symmetric keys required for different group sizes

symmetric key encryption, public key encryption, one way functions and cryptographic hash functions, as well as random number generators and SSL, the most commonly used protocol for secure transmission of information in the Internet commerce environment.

3.3.1 Symmetric Key Encryption

Encryption can be split into two types, symmetric key encryption and public key encryption. In the first case, the same key is used both for encryption and decryption, and in the second, a different key is used for the encryption and the decryption. The encryption key is called the public key and the decryption key is called the private key.

In symmetric key encryption, the same key is used for both the encryption and decryption processes. More formally, $D(E(m,k),k) = m$. The fact that the same key is used for both encryption and decryption causes the most problems in an Internet commerce environment. The two main problems are key distribution and non-repudiation.

For two parties to communicate securely, they must share a secret key that they use to encrypt communication between them. They can not use this secret key for any other communications with other parties, as such other parties would then also be able to decrypt the secret communications between them. In order for any group of size n people to communicate securely, they will require a total of ${}_nC_2$ keys. For a group size as small as 10 people, this requires 45 keys, and if we double the group to 20 people, 190 keys are required. This is clearly an exponential function, as shown in Table 3.1.

Another problem is that if the one of the parties does not have a secure channel to the other, they have no way of exchanging a first set of keys for communication. This is a scenario that happens quite often in the Internet world — in most electronic commerce situations the user has never met the organisation he is doing business with.

The second major problem with symmetric key encryption is that of non-repudiation. In case of a dispute between two parties as to the true content or origin of a message, there is no way for the recipient to prove to an independent third party that the sender did indeed send the message if the sender claims that the recipient falsified or modified it, since they both know the secret encryption/decryption key.

Let's say for example, that Alice is a client of Bob's stockbroking firm, and sends Bob an encrypted message to double her stockholding in FlyByNight.com; Bob purchases additional shares for Alice, right before FlyByNight.com's share price sets a new record for the biggest loss in a single day. Alice then claims that her message stated that she wanted to sell her holdings in FlyByNight.com, not purchase additional shares, and that Bob falsified the message. Since Bob also knows the key required for encryption, he can not prove that Alice did send the original message.

In general, though, conventional encryption algorithms are reasonably fast to use, and if a secure channel between two parties can be established, a session key can be generated to be used for encrypting the session's communications only, thereby solving the key distribution problem.

3.3.2 Public Key Encryption

Public key encryption was independently developed by Diffie, Hellman and Merkle in the USA [13] and the Government Communication Headquaters (GCHQ) in the United Kingdom. GCHQ actually developed the concept first, but due to secrecy requirements, published their findings more than 20 years later [55, p. 279–292]. Diffie, Hellman and Merkle's idea was made practical by Rivest, Shamir and Adleman when they developed the RSA algorithm [48]. This is a highly ingenious encryption system in which the encryption and decryption keys are different, and given one key, the other one can not easily be deduced from it. Formally, this can be represented as $D(E(m,k),k') = m$ with $k \neq k'$.

Any party wishing to communicate securely chooses a private and public key pair. The private key is kept secret and is known only to the original party. The public key is disseminated as widely as possible, to ensure that any party who wishes to communicate securely with the key owner can get hold of it. It is important to note that once encrypted with the public key, the original plain text can not be recovered from the cypher text by using the public key. It can only be recovered by using the private key. (Note that current implementations of public key cryptography depend on mathematically "hard" problems such as factorisation in the case of RSA or discrete logarithms in the case of Diffie-Hellman, and that the security of

these algorithms depend on the "hardness" of the underlying mathematical problems. See [21, 31]).

Public key encryption is ideally suited to the Internet commerce situation, as it voids the key distribution problem. If Alice wants to establish a secure channel to Bob, she sends Bob k_a, her public key and requests k_b, Bob's public key. Messages are encrypted using the other party's public key, and decrypted with each party's private key, k'_a in Alice's case, and k'_b in Bob's case. Another useful side effect from one of the properties of public key encryption $D(E(m,k),k') = E(D(m,k'),k)$ is that it makes digital signatures possible: if Alice decrypts a message with her private key, Bob can then encrypt it with her public key to retrieve the original message. Since Alice is the only person in the world to know her private key, the message must have originated from Alice, exactly as recovered by Bob.

Due to its importance both in Internet commerce and in later chapters, the RSA public key encryption is discussed in more detail.

RSA algorithm

The central concept of the RSA (Rivest Shamir Adleman) algorithm [48] is based on the hardness of factoring very large composite numbers, as opposed to the relative ease of generating and multiplying large primes. Should it ever become feasible to factor large numbers quickly (as would be the case with quantum computers [55, p. 325–331]) this would no longer be a secure encryption method.

Considering the case where Alice wants to communicate securely with Bob, she first has to generate the keys. Alice chooses two large (150 digits or more) prime numbers p and q, and calculates their product n. She then chooses a large integer d (the decryption key) that is relatively prime to $(p-1)*(q-1)$. This means that $(p-1)*(q-1)$ and d have no common factors greater than 1. The encryption key e is calculated as the multiplicative inverse of d modulo $(p-1)*(q-1)$. Alice's private key consists of n and d, and her public key consists of n and e, which she sends to Bob. Bob encrypts his message m by computing $M = m^e \bmod n$, and sends M to Alice, where she retrieves m by $m = M^d \bmod n$.

To clarify this process, we use a short numerical example. Alice chooses primes $p = 47$, $q = 59$. Their product $n = 47 * 59 = 2773$. Choose $d = 17$, an integer that is relatively prime to $(p-1)*(q-1) = 46*58 = 2668$. As d happens to be prime, and is not a factor of $(p-1)*(q-1)$ it is easy to see that d is relatively prime to $(p-1)*(q-1)$. We can now calculate e to be 157. Since $e*d = 17*157 = 2669$ and $2669 \bmod 2668 = 1$, it is clear that e is the multiplicative inverse of $d \bmod (p-1)*(q-1)$. To encrypt a message,

Bob maps the alphabet to the numbers 1–26, and uses $n = 2773$, $e = 157$. Lets says that Bob's message m is "hello". Thus $m = (0805\ 1212\ 1500)$ and $M = (805^{157} \bmod 2773\ 1212^{157} \bmod 2773\ 1500^{157} \bmod 2773)$, $M = (2639\ 2179\ 2153)$. Alice then uses $n = 2773$, $d = 17$ to decrypt the message $m = (2639^{17} \bmod 2773\ 2179^{17} \bmod 2773\ 2153^{17} \bmod 2773)$, $m = (0805\ 1212\ 1500)$ which translates to "hello".

Homomorphic encryption algorithms

"A homomorphism is a mapping from one algebraic structure to another under which the structural properties of its domain are preserved in its range in the sense that if $*$ is the operation on the domain, and $\circ$ is the operation on the range, then $\theta(x * y) = \theta(x) \circ \theta(y)$." [7].

Some public key encryption algorithms are homomorphic, specifically the RSA algorithm is homomorphic for the multiplication operation mod n, with both $\circ$ and $*$ being multiplication. This means that $E(a * b) = E(a) * E(b) \bmod n$ if E is the RSA public key encryption system. We present a proof of this statement:

Theorem 3.3.1 *$E(a * b) = E(a) * E(b)$ mod n where E is RSA encryption*

Proof: Let E be an RSA encryption, $E = m^e \bmod n$.
Then $E(a * b) = (a * b)^e \bmod n$.
Thus $E(a * b) = a^e * b^e \bmod n$
and $E(a * b) = (a^e \bmod n) * (b^e \bmod n) \bmod n$
Finally $E(a * b) = E(a) * E(b) \bmod n$

◇

This property will be used in a later chapter to design privacy protection mechanisms.

Digital Certificates

Public key encryption can only be trusted when the public key of the party you are communicating with does actually belong to that party — consider the following scenario.

Alice wants to communicate securely with Bob, but Charlie wants to intercept the communication and exploit the fact that Alice and Bob have not exchanged public keys before. When Alice requests Bob's public key, Charlie keeps a copy of Bob's public key, and substitutes his own in the message to Alice. He then also keeps Alice's public key, and sends his own public key to Charlie. Messages from Alice to Bob are then actually encrypted with Charlie's public key. Charlie then decrypts the message with his own private

key, modifies it if required, and encrypts it with Bob's public key before sending it on. Similarly the communications from Bob to Alice are tampered with.

The only way to prevent such a scenario, is for all parties to distribute their public key as widely as possible, or to make use of a certifying authority. The public key of the certifying authority is widely known (typically distributed with every web browser). Any party wishing to be contacted securely requests a certificate from the certifying authority. This basically means that the certifying authority digitally signs the public key and identity of the requester. Any third party wanting to communicate with the requester is given the certificate instead of just the public key, so that the authenticity of the public key and identity of the certificate owner can be verified. Currently the largest provider of digital certificates is Verisign [58].

3.3.3 One Way Functions

One way functions are functions that are easy to compute in one direction, but very hard to compute in the other direction. This might be likened to breaking an egg, which is easy, but putting a broken egg back together is impossible. A good example of a one way function is the modulo arithmetic used in the RSA algorithm; it is easy to calculate $M = m^e \bmod n$, but given M, e and n, it is very hard to recover m. The encryption process requires $O(\log_2(e))$ steps [48], whereas recovering m would require $O(n)$ steps, making it unfeasible to do so practically for larger n.

Hash functions are a specific type of one way functions that reduce any given input to a fixed number of bits, called the hash value. A simple example of a hash function is found in a normal address book — the first letter of any given input is used as the hash value, allowing the owner to quickly find an entry if a person's name is known. It must be easy ($O(l(s))$ is desired) to compute the hash of any given input s, where $l(s)$ is the length of s in bits.

For cryptographic purposes, they are often used to generate keys from random input strings, but not all hash functions are suitable for cryptography. A cryptographic hash function h needs to satisfy several other properties [14]:

1. Given a hash value m, it must be computationally infeasible to find an input s such that $h(s) = m$.

2. Given $h(s) = m$, it must be computationally infeasible to find a second input s' such that $h(s') = m$.

3. For a collision resistant hash function h it must be computationally infeasible to find any two inputs $s \neq s'$ so that $h(s) = h(s')$.

For an ideal cryptographic hash function h giving hash values of length c bits, the first two properties would take $O(2^c)$ steps to generate. Contrary to intuition, the third property would only take $O(2^{c/2})$ steps to generate, using the well-known birthday attack [25][p. 58, 111–112].

3.3.4 Random Number Generators

"A random number sequence is a sequence of numbers with the property that no member can be predicted from the preceding elements; in particular these cannot form a progression or follow any regular or repetitive pattern" [7].

Since computers are deterministic, it is obvious that no software algorithm can generate a truly random number sequence. This is a big problem in cryptography, as there are many cases where a random number is required, for example the prime number seeds for RSA encryption or a session key for conventional encryption. Pseudo random number generators are deterministic functions used to generate number sequences on computers that appear to be random to a lesser or greater extent. The "randomness" of these functions depend on both the function itself, and on the predictability of the seed value used to generate the first number. There are several tests that can be used to determine the quality of the random number generator, and some examples are explained by Marsaglia [36]. An overview of the types of random number generators can be found in another article by Marsaglia [35], and some of the more popular pseudo random number generators' strength are analysed by Kelsey et al. [28].

The importance of random number generators and the seed value can not be overemphasised in the Internet security context. Netscape's first implementation of SSL was insecure despite using well known and secure algorithms, because their choice of seed for MD5 was predictable [22]. Once again the old adage holds that a chain is only as strong as its weakest link.

3.3.5 Prime numbers

Prime numbers form a very important part of the RSA algorithm. In order to have practical implementations of RSA, it must be possible to generate very large primes. At first glance this seems as hard as to factor a very large number, but in fact there are ways to verify primality probabilistically in polynomial time.

Lets say that there is a polynomial time function $t(c, w)$ which, when given a candidate prime c and a possible witness w outputs "yes" in all cases if c is prime, and a better than 50% probability of outputting "no" if c is composite. Running t twice with two different witnesses w_1 and w_2 results in a better than 75% certainty that c is prime if both results were "yes". This can then be done as many times as required in order to be certain enough that the number is prime, as the probability of error halves with every iteration. A suitable certainty for RSA would be such that the chance of being wrong about the primality of p and q is less than the chance of an attacker randomly guessing the private key correctly.

A more complete description of such an algorithm can be found in Harel's book [24][p. 308–313]. Most of the technologies described above are used in the Secure Socket Layer (SSL) protocol, built into most Internet browsers.

3.3.6 Secure Socket Layer (SSL)

Secure Socket Layer or SSL is a seamless way of ensuring a secure channel between a user and an Internet site. This protects the integrity and secrecy of any information sent on the channel.

SSL uses certificates to prove that an Internet vendor's key belongs to that organisation. Public key encryption is used to authenticate the origin of the traffic as being from the site certified in the certificate, and to exchange session keys. Pseudo random number generators and cryptographic hash functions are used to generate the session keys, and symmetric key encryption is used to encrypt most of the communication. Message authentication codes are used to detect tampering with any message.

SSL is mentioned here as it is the first protocol cited by any Internet site when asked about security, and because it implements so many different security tools. The important point to remember is that SSL only secures data whilst in transit, and though the protocol is secure, implementations thereof might not be [22]. Secrecy of a user's information at the Internet site is not covered by SSL, and currently no mechanism exists to guarantee such secrecy. More information about SSL and an analysis can be found in [66].

3.4 Trusted Third Party

A trusted third party is a concept frequently used in security. Basically, it is a third party not necessarily involved in a transaction that can be trusted by both the other parties. Finding such a party in conventional (closed) security systems is relatively easy — usually one of the two parties involved

can play that role, for example during normal authentication on a network or database, the network server or database server act as the trusted party. On the Internet, this becomes more problematic, as there are more nodes, owned by different corporations and countries, and communications can easily be eavesdropped, rerouted, erased, spoofed or interrupted.

One solution to finding a trusted third party is to use de facto standard organisations with a reputation to protect. An example of this is Verisign, for certificates. Another solution when trusting a third party with information is to use a third party which already has access to that information, such as a bank for credit card details. We shall revisit trusted third parties later in the work.

3.5 Kerberos

Kerberos is a network authentication service that allows a party to be authenticated once, and then access several resources which might be distributed over the network. Kerberos in effect depends on a trusted third party, the Kerberos server, which will do the authentication. The Kerberos server is trusted by both the user being authenticated, and the resources to which access could be granted.

The central concept in Kerberos is that of a ticket, which allows an authenticated user to access resources by using it. The ticket is effectively a certificate issued by the Kerberos server, and contains a session key for communication between the user and the resource being accessed. The ticket also contains the user's id, and an expiry date. The user then presents the ticket to the resource it has been issued for, and they can then communicate securely using the session key. Access to the resource is granted once per ticket, and tickets are not allowed to be reused. In order to eliminate some of the effort of obtaining keys, a ticket granting ticket can be used, which has a longer expiry date, and is used to automatically request tickets from a ticket server.

More information on Kerberos can be found in [41, 29, 51].

3.6 Summary

We are now acquainted with the security tools we shall require later on in order to improve the privacy of an individuals information when it is supplied to a target organisation over the Internet. In the following chapters we shall focus more on the privacy issues and how privacy should be protected.

UNIVERSITY
OF
JOHANNESBURG

Chapter 4

Privacy

The privacy of information that an individual supplies over the Internet is the central theme of this thesis. As such it is very important that we are familiar with the issues surrounding privacy, and the current state of privacy protection; what has been done and where the shortfalls are.

Big Brother is watching you.
George Orwell
1984

4.1 Introduction

Privacy is an abstract concept. To work with privacy in computers, we first need a way to define what we mean by privacy in more concrete terms. 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.

It is also important to note that there is a distinction between private and personal information. Personal information is information containing personal attributes of a person, such as a phone number or an address, that might not always be private information, although this would depend upon the person and the application. The essence is that not all personal information is necessarily private information, but by treating it as such, involuntary privacy violations could be avoided. This is especially important as all personal information could be considered private under some circumstances.

We start by looking at the need for privacy, then we shall investigate current privacy protection mechanisms available. Finally we shall look at current needs, and where more development is required.

4.2 The need for privacy

There are two sides to Privacy. It is something a lot of people are willing to sell very cheaply in certain instances. (We gladly pay for our purchases with credit cards for the convenience of not having to carry cash with us. The down side is that a complete record of our every purchase could be kept in a database somewhere and might be available for all to browse). In other cases, it is something that we are not willing to compromise on lightly (such as our exact earnings) or even not at all (such as the medical records of an AIDS sufferer or a past criminal record).

When our private information is stored in a hard copy format where very few people have physical access to it, it is reasonably secure. However, the same private information can be several orders of a magnitude more vulnerable to unauthorised access when stored electronically, since millions of people can potentially access it should it not be protected.

These days private information is stored in electronic files, and those files

can be sold to advertisers, placed on the Internet for public viewing, hacked into by malicious hackers or even browsed by curious employees. The latter is particularly hard to prevent, since traditional access control is of little use where people are supposed to have access to private information about others to do their work, but should not be allowed access to the same information when motivated by curiosity.

There are countless examples of privacy being invaded, and private information is being increasingly compromised, misused, sold, or even made freely available to the public over the Internet — even when the government controls them [34]. In a panel discussion Samarati gives a good summary of the increasing concerns of privacy on the Internet [56]; others agree — see also [10, 11, 37, 67].

4.2.1 "Big Brother"

One of the biggest potential threats to privacy is posed by government organisations, as they have both the resources and manpower to make massive invasions of privacy possible. George Orwell predicted such a scenario in his book "1984" [42]. Newer technologies like ECHELON [9, 50] are starting to make such constant surveillance possible, as computers are used to filter and eliminate the bulk of the intercepted traffic, allowing human agents to focus on the pertinent messages.

The FAQ at Echelon Watch [17] estimates that up to 90% of Internet traffic is intercepted by ECHELON. ECHELON intercepts e-mail, faxes, telephone calls, Internet downloads and more. It then uses artificial intelligence to sift through the information and target specific communications. Far from being responsible with such powers, government organisations have often misused it [54, 34]. There have even been allegations that the United States used ECHELON intercepts to give US corporations a competitive advantage against their European counterparts [54].

What make this situation even more critical is that the American Constitution does not guarantee an individual's right to information privacy, despite beliefs to the contrary [60]. The only way to ensure privacy is to use encryption in normal communications, and enforceable privacy protection mechanisms when transacting on the Internet.

Possibly the best place for more information on ECHELON is on the ECHELON Watch home page [17]. Government interception is not the only threat to privacy — the criminal element has already moved to the Internet.

4.2.2 Crime

Criminals are making the most of the current lack of privacy protection on the Internet. Several avenues of weak security are open to them to get access to private information about individuals, which is then used for illegal or unethical purposes. Most often it is credit card details that are so misused — the online e-commerce accounts for a small percentage of total credit card transactions, yet it accounts for more than 50% of the chargebacks in some countries [18]. In France, where most credit cards are smart cards (improving real world security) and by law chargebacks are not allowed for unsatisfactory service or non-delivery, the Internet chargebacks accounted for 83% of disputes.

This is mainly due to poor privacy protection. E-commerce sites are storing credit card information, where they actually only require it for a single transaction. Software security holes, incorrect security configurations [49] and general laxness about security make private information easy to come by. Some sites get the security setup right, but then send an (unencrypted) e-mail confirmation message containing the private information — thereby undoing all the good done by using SSL.

This means that individuals need to protect their private information not only for the sake of privacy, but also to avoid becoming the victims of cybercrime. Although criminal access to private information is usually unauthorised, users also have to deal with cases where authorised access to their private information is being misused.

4.2.3 Commercialisation of private information

In many cases a user voluntarily discloses private information to an organisation for a specific purpose. Often such private information is then commercialised and sold as a way to generate additional revenue for the organisation concerned, without the consent of the user.

This is also a case of misuse of private information, and as such requires protection.

4.3 Ethics

Ethical considerations play a big role in the protection of private information, as many people have legitimate access to a user's private information, but should not misuse such access. An example could be an employee at the IRS having access to people's tax returns (with electronic submission over the Internet, any employee of the IRS could theoretically access any person's

online tax return. In the paper based scenario, this was much harder to do, as only one physical copy of the tax return existed, unlike the electronic case). Accessing a person's tax return in order to verify taxes paid is ethically fine, but browsing the tax returns of famous people out of curiosity is unethical and should be prohibited.

Unauthorised commercial exploitation of private information is also part of the ethical dilemma, as is the allegations that the ECHELON intercepts were used to give US corporations a competitive boost over their European counterparts.

This is one of the hardest cases of privacy invasion to safeguard as it is very hard to distinguish when legitimate access is used for unethical purposes. We now present an overview of the current state of privacy protection mechanisms that do exist.

4.4 Current privacy protection mechanisms

Currently privacy can be safeguarded in one of three logical ways, according to Cranor in [11]:

1. Private information is not disclosed at all.
2. The source of the private information is hidden, that is, anonymity is preserved.
3. Privacy policies are in effect that promise the responsible usage of private information.

Several mechanisms exist that use one or more of these logical methods to protect online users' private information. Although the use of privacy policies is becoming one of the most used methods to protect private information on Internet sites, other methods such as anonymising services, TRUSTe and Crowds also exist.

Although private information could also be protected using laws and regulations, we do not specifically consider legal aspects per se in this work, although we briefly touch upon it when discussing privacy policies.

4.4.1 Privacy policies

A privacy policy is basically a promise by an online organisation to treat private information submitted by users in a certain way. As this could constitute a legally binding contract between the user and the service provider,

it is currently one of the stronger methods to protect private information [60].

Some of the more common clauses in privacy policies are that the user has access to her information in order to rectify mistakes, that information collected will be kept to that required to complete the current transaction, and that such information will not be shared with third parties unless explicitly stated at collection time.

Despite these advantages, privacy policies have several downsides that limit their effectiveness:

- A user has to read the privacy policy first before submitting private information, interpret it and then decide if it is acceptable. If not, negotiating with the organisation is very hard, as this process is not defined and would essentially consist of manual communications through e-mail or some other medium.

- An organisation's privacy policy can be changed at any time. Any change could apply to information already collected, possibly under a pervious version of the policy which offered broader protection. There is also no automated protocol which would automatically inform users of a change in an organisation's privacy policy.

- No physical or computerised system prevents a breach of the privacy policy. In case of breach, the onus is on the user to find out about it, and the user then has to take legal steps in order to rectify such a breach. Massive damage could have been dealt to the user and taking legal action across international borders (as many e-commerce transactions cross borders) is not very easy.

Fortunately the W3C has created P3P [63], a way to automate at least the first concern mentioned above.

4.4.2 P3P

The platform for privacy preferences, or P3P is a framework designed by the W3C to allow a the specification of a privacy policy by an organisation. A potential visitor to the site can configure her browser or user agent with her own privacy preferences. When she then visits a site that supports P3P, the site's privacy policy can be automatically compared with her preferences. If there is a match, the browser or user agent could use an icon or pop-up message to let her know that it is safe to supply the information required. In case of a mismatch, the page could either not display, or a warning message

with a user confirmation could be displayed. In this manner the user could always be advised of the match between her own privacy settings and what the target organisation requires.

The original specification of P3P allowed for the negotiation of a privacy policy between the user and a target organisation, but this was too complex to be adopted in practice, as this required not only extensions to the current HTTP protocol, but also required a new protocol for the negotiation. P3P V1 removed the negotiation phase, and a user could either accept or reject an organisation's policy, and nothing more. The problem with this approach is that it is too simple, and provides little more than automatic verification of an organisation's privacy policy. See Grimm's article [23] for more concerns about P3P.

There are already some web sites that implement P3P. Although P3P eliminates the need for a user to read and verify a site's privacy policy before submitting private information, privacy is still not enforceable. After information has been submitted to a web site, and the organisation's privacy policy changes, there is no method to "take back" such private information. Breach of the privacy policy would still have to be manually detected and remedies have to be sought outside the P3P protocol. Basically then P3P is only useful as long as the organisation requesting private information can be trusted. Even if this is the case, P3P defines no physical controls to ensure that no misuse is made of private information.

In short then, P3P is a significant enhancement to normal privacy policies, but does not allow enforceability. The current P3P specification is at the W3C [63]; a higher level overview is given by Reagle and Cranor [46].

P3P policies can be evaluated by the user agent using preferences specified in APPEL.

4.4.3 APPEL

A P3P Preference Exchange Language (APPEL) [62] is a language used to describe a collection of P3P preferences in a ruleset for implementation by a user agent. The syntax is highly dependent on P3P itself, and as such ties in very narrowly with the P3P syntax. The rules in a ruleset are evaluated in order, so higher priority rules should be placed at the top of the ruleset.

A user defines his privacy preferences in APPEL, and the user agent then evaluates the rules and compares that to a site's privacy policy. The outcome of the rule evaluation can be one of three possibilities, namely request, limited or block, with an optional user prompt. Request means that the site should be accessed normally (HTTP request); limited means that all but the most necessary request headers should be suppressed when requesting the page,

and block means that that URI should not be accessed as it violates the user's privacy preferences. The optional prompt argument could allow this behaviour to be overridden by prompting the user if set to "yes".

An APPEL ruleset should always have at least one rule that evaluates to true, so it is advisable to have a default rule at the bottom of the ruleset that will always fire to implement default behaviour. The individual rules contain expressions that are evaluated. If all of a rule's expressions are satisfied, the rule evaluates to true.

Within an expression, special connective attributes describe the way contained expressions are evaluated. The six connective attributes are:

or Returns true if at least one contained expression matches an element of the privacy policy being checked. If the policy contains additional elements not listed, they are ignored.

and Returns true if all of the contained expressions match elements of the privacy policy being checked. If the policy contains additional elements not listed, they are ignored.

non-or The opposite of or, returns false where or would return true, and vice-versa.

non-and The opposite of and, returns false where or would return true, and vice-versa.

or-exact Returns true if at least one contained expression matches an element of the privacy policy being checked. If the policy contains additional elements not listed, returns false.

and-exact Returns true if all of the contained expressions match elements of the privacy policy being checked. If the policy contains additional elements not listed, returns false.

Some sites do not automate their privacy policies, but use an independent third party such as TRUSTe to audit their policies.

4.4.4 TRUSTe

TRUSTe is a mechanism to reassure users that a target organisation displaying the TRUSTe logo is serious about protecting their privacy, and undertakes to (among others) the following requirements:

Notice The organisation's web site will have a link to its privacy policy from its home page, so that a potential user can easily access and review it.

Choice At a minimum the target organisation will allow users the choice to opt out of having their private information shared with third parties for other purposes such as advertising.

Security The target organisation must implement reasonable security measures to protect a user's private information where it is stored from unauthorised access.

Data quality and access Users must be able to view and correct inaccuracies in their private information as stored by the target organisation.

When a user visits a site displaying the TRUSTe logo, she can also click through to the TRUSTe website to verify the integrity of the TRUSTe seal. The TRUSTe organisation also randomly visits TRUSTe sites under assumed identities to verify compliance with their terms. The biggest failing of TRUSTe is once again that these requirements are not programmatically enforceable. More information about TRUSTe is available in Benassi's overview [2] and on the TRUSTe home page [57].

4.4.5 Anonymisers

Anonymisers target the source of private information. Their aim is to hide the exact source of private information, thereby improving privacy. For example, if Bob knows that Alice earns $1000 a month, her privacy is compromised, but if he just knows that a visitor to his site earns $1000 without being able to specifically identify Alice, her privacy is not compromised although her personal information has been given out.

The biggest problem with anonymisers is that they are of very limited use in an e-commerce situation, as in most cases the purchaser needs to be identified for funds transfers and deliveries. Anonymisers are very useful in cases where information only is required — if Alice is HIV-positive and would like to find out more about the condition, she can use anonymisers to avoid any information sites from being able to trace the requests back to her, thereby protecting her private information.

There are several anonymising services on the Internet, and the services they offer range from anonymising to anti-logging. A comparison of some of these services is contained in Table 4.1.

The comparison criteria are defined as follows:

Anonymise The user's identifying information is hidden from the site being visited (at least the IP address, previous site visited and domain).

	Safeweb	Anonymizer	Rewebber
Anonymise	Yes	Yes	Yes
Cookie block	Yes	Yes	Yes
Anti-filtering	Limited	Yes	Yes
Anti-log	No	No	Yes
Anti-censorship	No	No	Yes
Anti-eavesdrop	Yes	Yes	Yes

Table 4.1: Comparison of anonymising services

Cookie block The site intercepts cookies, and prevents surfing patterns from being tracked with cookies.

Anti-filtering Web addresses are encrypted or obfuscated so that proxy filters can not deduce which ones to filter.

Anti-log Web addresses are encrypted as for anti-filtering, but the encrypted link is only valid for a short time, so that subsequent access attempts will fail, invalidating any activity logs.

Anti-censorship The anonymising service registers a separate unrelated domain name and uses that as an entry point for each specific customer, so that access to the anonymising service can not be blocked.

Anti-eavesdrop SSL is used for visited sites, so that ISPs and Internet access providers can not intercept communications.

Another shortfall of anonymisers is that there is a single point of failure — if an anonymising service is compromised (or even a dummy service set up by a malicious organisation) all traffic through that service is compromised. In such a case the use of an anonymising service is even worse that not using one, as all web traffic is routed through the service. A possible solution is to string several anonymisers together (with a performance penalty) or to use Crowds.

4.4.6 Crowds

Crowds is a mechanism to protect the anonymity of requests to the Internet, like anonymisers. Unlike anonymisers, though, Crowds is a peer to peer system, which could improve anonymity in that there is no single point of failure where the source of private information could be compromised.

Every user runs a process called a jondo on their computers, and sets up their browser to use the jondo as a proxy server. All the jondos in a group

then form a Crowd. A routing path is randomly formed through some of the jondos in the Crowd for each individual jondo, and the path is cached until a new jondo joins the Crowd, in which case all path are rerouted. Each jondo is aware of its predecessor and successor on the path, but nothing more. All requests are then sent along the path where the end node submits it to the actual website. The reply is similarly sent back along the path. Requests originating from a particular jondo can not be distinguished from requests merely being passed on; the only difference is that the jondo will not pass the reply for a request it originated back to its predecessor.

Some of the apparent problems with using Crowds is that server logs might point to objectionable sites being visited by a person in the Crowd who did not initiate the request, but as Crowds is designed not to distinguish between requests being passed on and requests originating locally, this might be a good defence. Another apparent problem is that it is possible that due to routing on the path the initiator of a request could be the one sending it directly to the target website, but once again since no differentiation exists between originating and forwarded requests, this is not really an issue.

The real risks is that Crowds increases network load and decreases response time as experienced by the end user. Crowds should be formed among users with similar bandwidth availability, so that a dial-up user does not end up rerouting a broadband user's data requests.

In summary, Crowds is very interesting privacy protection mechanism, but it does not protect the security of private information, only the source. Crowds can be dowloaded from the Crowds home page [12], and is explained in more detail by Reiter and Rubin [47].

4.4.7 Namesafe

Namesafe.com [40] 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.

Protection mechanism	Method (Cranor) [11]
Privacy policy	3 (Policy)
P3P	3 (Policy)
TRUSTe	3 (Policy)
Anonymisers	2 (Anonymity)
Crowds	2 (Anonymity)
Namesafe	1 (non-disclosure) and 2 (Anonymity)
iPrivacy	1 (non-disclosure) and 2 (Anonymity)

Table 4.2: Summary of privacy protection mechanisms

4.4.8 iPrivacy

iPrivacy.com [26] 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 just the right to access 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 card company, which, of course, makes it inaccessible to people without credit cards. As of April 2002, no credit card companies have yet signed up with iPrivacy, making it still a future service.

A summary of these existing protection mechanisms is presented in Table 4.2.

4.5 Need for development

Based on the discussion of existing privacy mechanisms it is clear that there is a real need for development of better mechanisms. If measured against the criteria for secure data of secrecy, integrity and availability, the secrecy is still lagging the furthest, in that none of the existing mechanisms allow access to be revoked once given. Apart from TRUSTe requiring access for a user to update information, the integrity of private information is not addressed either. Even in the case of TRUSTe, the user would still have to notify all sites individually, making integrity difficult to achieve.

4.6 Summary

We have given an introduction to the importance of privacy, and discussed the current mechanisms available to protect private information. We have also shown that they all have shortcomings, and that there is a clear need for the development of more advanced privacy protection mechanisms. In the following chapters we shall proceed to develop such a mechanism.

Chapter 5

Classification

We propose a classification of private information based on the purpose it is being requested for. Such a classification would enable us to tailor protection methods for each class of information so that the information can be protected without making the data useless to the target organisation.

Give us the tools, and we will finish the job.
Winston Churchill
To President Roosevelt, 1941

5.1 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 thereof, 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 six 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 sixth. 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 six.

5.2 Classes

5.2.1 Class 1 — Validation

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. Various mathematical operations are then performed on these amounts to derive a taxable income. All of these amounts are not actually required directly by the tax authorities, as they intend using them 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.

5.2.2 Class 2 — Third party use

The second class we define will 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 funds 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 calculate 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.

5.2.3 Class 3 — Future use

In the third class 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 cancellation penalty. In cases where the reservation is several months in advance, it is quite possible that the credit card, whilst being valid at reservation time, might have been closed or suspended before the actual date of the accommodation.

5.2.4 Class 4 — Identification

In our fourth class 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.

5.2.5 Class 5 — Direct usage

This class shall consist of all cases where the private information is directly required for the current transaction, and any degree 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, or an individual's taxable income on an electronic tax return.

5.2.6 Class 6 — Other

Our final class contains all private information, even such items that have neither been requested by nor supplied to any online retailer and is a container for all the other classes, in that all the other classes are subsets of this class. No protection mechanism will be developed for this class of information.

5.3 Observations

It is clear that any item of private information from classes 1 to 4 can be treated as though in class 5 without invalidating the purpose it was acquired for. The principal difference between classes 5 and 6 is that class 5 will still be included in our protection mechanisms, though nothing is preventing an organisation accessing class 5 information from actually storing it after the first access and bypassing the protection for subsequent accesses, while class 6 information will not be included in our protection mechanisms. Private information in class 6 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.

Information item	Purpose	Class
Name	Identification	4
Name	Shipping Label	2
Shipping Address	Shipping Label	2
Zip code	Shipping Label	2
Zip code	Shipping costs	5
Credit Card details	Payment	2
List of books	Actual shipping	5
E-mail Address	Communication	2
E-mail Address	Special offers	3
Phone Number	Shipping Label	2

Table 5.1: Classification of private information for online book order

5.4 Example

We present an example classification of all the items of private information required when Alice buys books from Bob's online bookshop.

Bob requires Alice's name both for identification, and to use on the shipping label. He also requires a shipping address, phone number and zip code to ship the books to, as well as Alice's credit card details for payment. The list of books that Alice is ordering will be used to fill the order. Bob would also like Alice's e-mail address to inform her of new releases by the publishers of the books that she is ordering. A summary is listed in Table 5.1.

It is interesting to observe that the same item of private information can be required at different times for different purposes, so that it can belong to more than one class for the purpose of a single transaction, such as the zip code in the example, which is used both for the shipping address and to calculate shipping costs. In such cases the effective protection would be less than that of any of the two classes it falls into — typically class 2 information would not be visible to the organisation the user is transacting with, and class 5 information would not be visible to any third party such as the shipping company. In this example, however the information is available to both the organisation and the shipping company.

P3P also attempts to categorise the usage that will be made of private information through its purpose elements, but there are some core differences.

5.5 Comparison with P3P Purpose elements

P3P, first discussed in Chapter 4, will now be used to compare the P3P Purpose element as listed in the P3P specification [63] with our classification. In P3P the purpose element on a site's privacy policy is used to specify the reason a particular data element is collected. These reasons can be one or more of the following:

<**current**/> Completion and support of current activity

<**admin**/> Web site and system administration

<**develop**/> Improve and further develop the site (not customisation)

<**tailoring**/> Once-off customisation of site to individuals (not to be used for future customisation)

<**pseudo-analysis**/> Analyses user behaviour without using personal identification of individuals

<**pseudo-decision**/> Customises pages without identification of individuals

<**individual-analysis**/> Analyses user behaviour with identification of individuals

<**individual-decision**/> Customises pages whilst using personal identification of individuals

<**contact**/> Contacting of visitors for marketing of services and products other than through the telephone

<**historical**/> Historical preservation of social history — must include disputes information and researchers allowed access

<**telemarketing**/> Contacting visitors for marketing of services and products through voice telephone

<**other-purpose**/> Other uses

Additionally, each of these reasons can have the level of requirement specified to be either always (assumed), opt-in to explicitly opt-in before data is collected, and opt-out, where data is collected unless the user explicitly opts-out. These elements can be further grouped into six high level groups:

Site operation and administration; R & D Contains the elements <admin/>, <current/> and <develop/>, and does not include personal information for customisation or future contacting of the user.

Site customization This contains only the <tailoring/> element and is information used only for a single visit to a site.

Anonymous user tracking Consists of <pseudo-analysis/> and <pseudo-decision/> to gather information about user habits and preferences without specifically identifying the user.

Individual user tracking Uses personal information about a user to either customise the current page directly for the user <individual-decision/> or to build a profile of an individual user <individual-analysis/>.

Contacting the user This groups together the elements used to contact a user for purposes not related to the current activity, such as marketing. It contains the <contact/> and <telemarketing/> elements.

Other purposes Information collected for other purposes such as <other-purpose/> and <historical/>.

There are some similarities between this set of purposes and our classification. Firstly, class 5 (direct usage) maps to purpose <current/>. Class 4 (identification) has some similarities with <individual-analysis/> and <individual-decision/>, and <historical/> would obviously be part of class 3 (future use), but not vice-versa. Finally, class 6 (all other private information) fulfills a similar role to purpose <other-purpose/>.

Note that this is a very loose mapping, as the primary aim of the P3P purpose element is for the user to decide whether or not some item of private information will be supplied to the requestor, whereas our classification is used after the user has already decided to supply some item of private information to the requestor, but wishes to retain some form of control over such information. In this case the class that the information falls into will be used to select an appropriate protection mechanism in order to retain some control over the information supplied.

In a way we could also argue that all the P3P purpose elements map to the same class, namely class 5, as no protection mechanism is applied to the private information in P3P after the decision has been made to divulge such information, whereas the purpose of classes 1 through 4 is to tailor some form of protection around the data. Based on these different mappings to the P3P purpose elements, it is clear that our classification is required, and can not be replaced with the P3P purpose elements.

Mechanism	Class
Privacy Policies	1, 2, 3, 4 and 5
P3P	1, 2, 3, 4 and 5
APPEL	1, 2, 3, 4 and 5
TRUSTe	1, 2, 3, 4 and 5
Anonymisers	1, 2, 3, 4 and 5
Crowds	1, 2, 3, 4 and 5
Namesafe	1, 2, 3, 4 and 5
iPrivacy	1, 2, 3, 4 and 5

Table 5.2: Classification of existing protection mechanisms

5.6 Classification of existing protection mechanisms

In Section 4.4 we described several current privacy protection mechanisms. In this section we shall investigate the classes they attempt to protect. As can be seen in Table 5.2, each of the existing protection mechanisms attempt to protect all of our classes of private information, which is both their greatest strength and their greatest weakness — because they protect all classes of private information they are simpler to use, but their protection can not be enforced. This is because class 5 requires direct access to the private information, and any attempt to enforce protection will invalidate the purpose the information was required for. The only way to effectively tailor protection for the different classes of private information is to treat each class differently, as discussed in the next section.

5.7 Protection methods for each class

We now discuss the requirements for protection mechanisms for each class, based on the properties for secure data mentioned in Chapter 3, namely secrecy, integrity and accessibility. These idealistic requirements will form the basis of what we shall aim for when designing protection mechanisms for these classes later on. A quick overview of the protection an ideal method would provide for each class is presented in Table 5.3.

5.7.1 Class 1 — Validation

In the case where private information is only required to validate a calculation or to generate statistics, we would ideally like to keep the information

Class	Organisation	Other
1	No	No
2	No	Yes
3	Later	No
4	No	No
5	Yes	No
6	N/A	N/A

Table 5.3: Private information access table of ideal protection mechanisms for each class

secret (secrecy) since the individual terms of the calculation are not directly required. Only the result of the calculation would have to be divulged (accessibility). Despite the individual terms being kept secret, the party to which the result is sent must be able to verify the correctness of the calculation (integrity).

5.7.2 Class 2 — Third party use

For our second class, where information is requested in order to pass it on to a third party which requires it, we would also like to hide the information from the organisation (secrecy) since it does not need the information itself. The third party requiring the information should be able to access it (accessibility). Updates to the information should also filter through to the third party directly, without the organisation being involved (integrity).

5.7.3 Class 3 — Future use

In cases where private information is requested by an organisation for future use, we would like to reveal the information only when it is actually required, and not ahead of time (secrecy). When the information is then actually required, the organisation should have guaranteed access to it (accessibility). Such information should be up to date, in case it changed since access was granted but before access is required (integrity).

5.7.4 Class 4 — Identification

Where private information is required for identification purposes only, it would be best not reveal the information at all (secrecy) but to establish some other mechanism for identification that is reliable (accessibility) but that does not require knowledge of private information about the individual.

Identification should be possible without false matches or invalid rejections (integrity).

5.7.5 Class 5 — Direct usage

In the fifth class where private information is directly required for some purpose, the secrecy requirement can not be fulfilled, implying that the information would have no enforceable security. Secrecy in these cases would be achieved through privacy policies, legislation and trust. Accessibility is good if the organisation stores its own copy of the information, but this compromises secrecy even further and destroys integrity. If the organisation does not store a copy but requests the information every time it is required, accessibility is compromised, although integrity is good.

5.7.6 Class 6 — Other

Class 6 is mentioned for completeness only, as no protection mechanisms will be developed for this class.

5.8 Summary

In classifying private information according to the purpose it is required for allows us to measure the effectiveness of any protection mechanism we develop. More specifically, it allows us to verify that one of the core properties of secure data, namely accessibility is still satisfied after a particular protection mechanism has been applied to the data. If the other two properties of secure data (integrity and secrecy) can then be guaranteed by the protection mechanism, we can be satisfied that we have increased the security and therefore the privacy of the data concerned.

Without such a classification it will not be possible to effectively measure the availability of the data, or to develop a protection mechanism for it, as any protection mechanism applied to the general case does limit its accessibility in some way or another. If no method existed to prove accessibility to authorised users, such a mechanism could not be applied. This classification of private data therefore is quite a valuable tool when searching for ways to secure private information.

UNIVERSITY
OF
JOHANNESBURG

Chapter 6

Method 1 — Nondisclosure

In this chapter we shall present an encryption method to protect private information where the private information consists of a set of numeric values S on which some mathematical function G has to be applied and the result $\alpha = G(S)$ has to be supplied to a target organisation. The result α must be verifiable by the target organisation, without disclosing S. We apply this method to the specific case of protecting the privacy of electronic income tax returns, and discuss other possible applications.

The most visible joy can only reveal itself to us when we've transformed it, within.
Rainer Maria Rilke
Duino Elegies, 7

6.1 Introduction

The purpose of this chapter is to propose a method to safeguard our private information in cases where this information is required not for the contents, but as input to verify calculations. We shall present an encryption method which allows certain mathematical operations to be performed on a set of encrypted values and to cross check this result with the result obtained by performing the same operations on the unencrypted values. This would allow us to use this protection method on problems falling into class 1 of our privacy classification in the previous chapter.

More formally, this method is applicable to cases where the private information consists of a set of numeric values S on which some function G has to be applied and the result $\alpha = G(S)$ has to be supplied to a target organisation o. The result α must be verifiable by the target organisation o, but we would like to keep S private. A tax form submitted to the IRS is a very well known practical application hereof.

This method protects the private information against unauthorised access from outside the organisation it was sent to. It does not depend on the organisation to protect the data adequately, as the private information is encrypted. It also prevents people in the organisation with valid access to browse this information out of curiosity, something that usually was out of the hands of the originator of the private information.

We propose a way to apply the first of Cranor's options (see page 36) to a class of applications where previously the third option was the best alternative, by using homomorphic encryption algorithms (see page 25). Although many researchers have mentioned or devised homomorphic encryption algorithms [1, 15, 39, 52], we believe that there is still scope to combine such algorithms and apply them to protect privacy.

This is specifically needed for sensitive issues such as income tax records, which are still vulnerable despite controls being implemented, as mentioned in [34]. In the rest of the chapter we shall continually refer to the income tax example problem.

6.2 Problem description

Let's start by restating our definition of privacy.

Definition 6.2.1 *Privacy is defined as the 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.

6.2.1 Problem statement

In certain cases 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 $x_1, x_2, x_3, \ldots, x_t$ about an individual i is collected from a number of sources to which some function G is applied to give a result α where $\alpha = G(x_1, x_2, x_3, \ldots, x_t)$. 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 organisation o needs to verify the calculation of α, o does not strictly require access to $x_1, x_2, x_3, \ldots, x_t$.

If we could find a way of proving to o that $\alpha = G(x_1, x_2, x_3, \ldots, x_t)$ without disclosing $x_1, x_2, x_3, \ldots, x_t$, we would be able to solve our problem. This would enable us to retain exclusive control over our private information.

6.2.2 The IRS example

Every year we supply the IRS with a lot of very private information about ourselves. The sole purpose of supplying them with this information is to allow them to verify that the amount of tax payable by us as stated on our tax returns has been correctly calculated.

In the terminology of our problem statement, the IRS is o and the taxable income is represented by α. The taxable income α is calculated by adding various amounts representing income, and subtracting other amounts representing allowable deductions. We shall use $x_1, x_2, x_3, \ldots, x_t$ to represent the various amounts that are used in the tax calculation, which we shall denote with the function $G(x_1, x_2, x_3, \ldots, x_t)$.

In the traditional or unsecured tax return, the IRS is now supplied with α as well $x_1, x_2, x_3, \ldots, x_t$ to enable them to verify α using G which is known to all parties.

6.3 Proposed solution

In order to keep our private information secure, we propose to encrypt $x_1, x_2, x_3, \ldots, x_t$ using an encryption function f such that $y_i = f(\epsilon, x_i)$ where

ϵ is the encryption key. The result of the encryption will give us y_1, y_2, y_3, $\ldots, y_t$. Note that most of these values are usually supplied by third parties such as employers or banks in the tax example, 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 the encrypted values $y_1, y_2, y_3, \ldots, y_t$ instead of $x_1, x_2, x_3, \ldots, x_t$ 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 $G'(y_1, y_2, y_3, \ldots, y_t) = \alpha'$ where $\alpha' = f(\epsilon, \alpha)$. This would allow the IRS to verify the correctness of α using the encrypted values $y_1, y_2, y_3, \ldots, y_t$ without ever knowing $x_1, x_2, x_3, \ldots, x_t$ by applying the function G' to $y_1, y_2, y_3, \ldots, y_t$ to obtain α' and then apply function f to α' and comparing the results. If $f(\alpha) = \alpha', \alpha$ has been correctly calculated from $x_1, x_2, x_3, \ldots, x_t$.

6.3.1 Restrictions

It is clear that we have to supply the target organisation o as well as the other organisations such as banks and employers issuing tax certificates with both the encryption function f and encryption key 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 6.3.1 *The function $f(\epsilon, x) = y$ should obscure x so that there is no easy way to determine x given f, ϵ and y.*

Property 6.3.2 *If $f(\epsilon, x_1) = f(\epsilon, x_2) = y$, then $x_1 = x_2$*

Property 6.3.3 $G'(f(\epsilon, x_1), f(\epsilon, x_2)) = f(\epsilon, G(x_1, x_2))$

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

6.3.2 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.

Multiplication

If the function G to be performed is normal multiplication, we define f to be f_* where f_* is RSA encryption as described on page 24. 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 $f_*(x) = x^\epsilon \bmod n = \alpha$ and α' is obtained by $\alpha' = y_1 * y_2 * y_3 * \ldots * y_t \bmod n$.

Theorem 6.3.1 *f_* satisfies the properties defined in Section 6.3.1 where G is multiplication and G' is multiplication mod n.*

Proof:

1. This is trivial, from the usage of RSA for f: there is no easy way to determine x_i given y_i, ϵ and n for any value of i.
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:
   $((x_1{}^\epsilon \bmod n) * (x_2{}^\epsilon \bmod n)) \bmod n$
   $= (x_1{}^\epsilon * x_2{}^\epsilon) \bmod n$
   $= (x_1 * x_2)^\epsilon \bmod n$

◇

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, as the original RSA article in 1978 [48] already suggested using a 200 digit number for n, which is a 670 bit binary number.

Addition

If the function G to be performed is normal addition, we use a new public key cryptosystem as described by Naccache and Stern [39] for f. We define f to be f_+ where $f_+(\epsilon, x) = \epsilon^x \bmod n$ and ϵ is a generator for GF(n). As for RSA, n is the product of two large primes p and q. For a full description the reader is referred to [39].

Theorem 6.3.2 *f_+ satisfies the properties defined in Section 6.3.1 where G is addition and G' is multiplication mod n.*

Proof:

1. If we examine $f = \epsilon^x \bmod n$ 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:
   $((\epsilon^{x_1} \bmod n) * (\epsilon^{x_2} \bmod n)) \bmod n$
   $= (\epsilon^{x_1} * \epsilon^{x_2}) \bmod n$
   $= \epsilon^{x_1+x_2} \bmod n.$

◇

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, as noted above.

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.

6.4 Example implementation

A fictional example using the South African tax law. For this example we shall set $n = 19697446673$ and $\epsilon = 131$. 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.

6.4.1 Calculation

	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 6.4.2*
less	Travel expenses	-26136	847550577	A *From 6.4.3*
	Total taxable income	122229	6922828699	

Note that 18182403837∗2608534625∗7122548528∗9973525707∗847550577 mod 19697446673 = 6922828699 and that 122229 encrypted for addition also yields 6922828699.

6.4.2 Interest sub-calculation

	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	

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.

6.4.3 Travel expenses sub-calculation

	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

Note that the totals can be verified as above.

6.5 Security and other issues

Some issues need to be taken into consideration when applying this method. We mention the more prominent ones.

6.5.1 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. To demonstrate, let us assume that an encryption function g preserves ordering, and that $y = g(x)$ is an encrypted value. Since we have to supply the target organisation 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 the example, where the total interest calculation is disclosed, but the various sources of interest are hidden for a concrete example.

6.5.2 Exhaustive searches

A possible problem with this method is that the exhaustive searches required to decrypt $y_1, y_2, y_3, \ldots, y_t$ would not take too long using computers, since the values of $x_1, x_2, x_3, \ldots, x_t$ 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.

6.5.3 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).

6.5.4 Related work

Ahituv, Yeheskel and Neumann examined the use of homomorphic functions on encrypted data stored in a database [1]. As they were working in an environment where the encryption could be performed by the DBMS or secure operating system, they used solutions where either the key or the algorithm has to be kept secret. In the first case, for the addition of plaintext to ciphertext without decryption, the resulting system had the weakness that if any intermediate balance or the key became known, all the updates were exposed. In the second case for the addition of encrypted data to encrypted data they still had the key discovery weakness. Where they used homomorphic functions, these were based on block cyphers, and as such could be broken with a set of linear equations, as comparison was kept as a function. In the approach we have presented here, we could not use symmetric key encryption, and had to use public key encryption. As such, we do not have either the linear equation or key discovery through known plaintext vulnerabilities, at the cost of a reduced set of possible operations on the encrypted data.

Domingo-Ferrer and Herrera-Joancomatí presented a privacy homomorphism that had the advantage of being a field homomorphism [15]. Unfortunately it is not secure against a known plaintext attack, and it is a symmetric key encryption in that being able to encrypt values implies being able to decrypt them again, which would not work for our Internet privacy requirements.

6.6 Other problems to which this solution is applicable

This solution is not just suitable to protect the privacy of electronic tax forms. Many other examples exist, such as the ability to modify encrypted database values without decryption, secret e-mail voting with verifiable results, verifying directors' remuneration, etc. We describe the e-mail voting example in more detail.

6.6.1 E-mail voting

In a secret, verifiable e-mail vote, a trusted party (the vote counter) generates the encryption and decryption keys. The encryption keys are distributed to the voters. Each voter then submits his encrypted vote. We assume a yes or a no vote, represented by 1 and 0 respectively. The vote counter adds all the encrypted votes together, and decrypts the result. The vote counter then publishes the encrypted votes and the decrypted result. Everyone can now verify the calculations and ensure that their vote was counted, without compromising anyone else's vote. To prevent all the encrypted values from being the same when the vote is the same, we can increase the magnitude of the yes vote, and add a randomising factor to each vote.

In practice a voting system for i voters can be designed as follows: a yes vote can be defined as 1 followed by c 0's (the yes constant y) added to a random number m, where $m < i^2$, $c = \log i^4$. A no vote would just be a random number m. Based on the definition of c, the yes constant y would still be an order of a magnitude larger than the sum of all the no votes, even if every voter voted no and chose $m = i^2 - 1$ ($i * (i^2 - 1) = i^3 - i \ll i^4$ for positive i). The sum of all the votes, divided by the yes constant y will give the number of yes votes.

As a trivial example for 10 voters, the yes constant y would be 10000 and the randomising factor would be an integer less than 100, so that a yes vote would be a random number between 10000 and 10100, an a no vote would be a random number less than 100. If every voter voted no, the total would be less than 1000, which when divided by the yes constant y would clearly reveal that there were no yes votes. If (say) three people voted yes, the sum of the votes would be value between 30000 and 31000, which would reveal the three yes votes, when divided by the yes constant of 10000.

6.6.2 Other problems

Other problems such as modifying encrypted values in a database without decrypting them first and disclosing total directors' remuneration without revealing individual earnings, but whilst still being able to verify the calculation can also be addressed with this solution.

6.7 Summary

In this chapter we have presented a novel approach to the protection of electronic data privacy: that of algorithmic protection of privacy using encryption, which allowed us to use our private information (in this case details of income or an e-mail vote) but while still protecting our private information from prying eyes, including the party to which we sent our private information (at least in a simplified tax example and the e-mail vote). Typically the data we could protect would fall into class 1 of our classification. In the following chapters we shall look at protection mechanisms for the other classes.

Chapter 7

Method 2 — Retain control

In this chapter we present a method by which we can grant limited access to our private information, and thus enforce the terms of target organisations' privacy policies. Private information is also revealed at the last possible stage, further reducing the possibility of misuse. This safeguards private information in four of the five classes previously mentioned.

'Tis our true policy to steer clear of permanent alliances, with any portion of the foreign world.
George Washington
Farewell Address [September 17, 1796]

7.1 Introduction

In most cases where a subject is granted access to confidential 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.

More formally, 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.

In this chapter we present a protocol 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 shall require the classification of private information developed in Chapter 5 as well as tickets, public key encryption and a trusted third party, described in Chapter 3.

7.2 Overview

To effectively present this method, we shall start by giving an example step by step application of the proposed method, and then we shall specify it in more detail.

7.2.1 Example

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 ticket granting ticket (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 to 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).

7.3 Implementation

The general case is stated, as well as certain properties of the tickets to implement the protocol correctly. A ticket granting ticket (TGT) will be issued by the trusted third party S to the target organisation o for each transaction between the individual i and o. Such a TGT will grant o access to such private information m about i stored at S described by the intersection of o and i's privacy policies. A graphical example of the protocol in action is depicted in Figure 7.1.

7.3.1 General Case

When an individual i wants to send some private data m to a target organisation o, o's privacy policy P is checked against i's preferences using P3P. This policy (if valid) 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 privacy preferences 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, another organisation who actually requires the access 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.

7.3.2 More on Tickets

Tickets are used to access the actual private information. 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 as defined in Chapter 5.

Validation and calculation

Although not directly addressed by this chapter, this method can be used to calculate aggregates and averages on a set of different users' private information. If allowed to do so by the privacy policy, o can request S to calculate such aggregates.

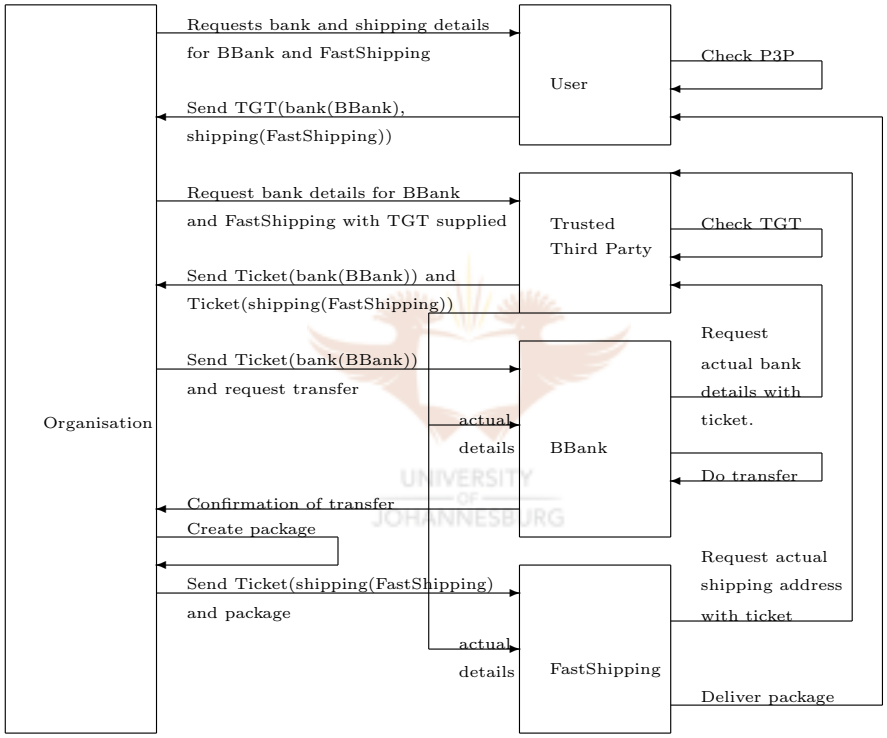

Figure 7.1: Graphical example of the protocol

Third party requirement

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

Future use

In our third class 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.

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.

Actual contents

In the fifth class where actual access to m is required to actually deliver a package (used by the shipping company), 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).

7.3.3 Properties of Tickets

In order to function as described above, tickets need certain properties, similar to the Kerberos implementation, discussed in Section 3.5 on page 29.

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.

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 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.

Format of the tickets

The ticket granting ticket should implement the DTD specified in Figure 7.2, while the ticket should implement the DTD in Figure 7.3. Tickets can not be reused. Sample XML implementations can be found in Figure 7.4 and Figure 7.5.

7.4 Discussion

Careful study of the proposed protocol raises a couple of discussion points, which we briefly investigate here, namely the trustworthiness of the third party, possible collusion between the participants in the transaction, assurances for both parties, and some possible modifications to prevent misuse of such privacy by criminal elements.

7.4.1 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 chosen for a prior implementation of the P3P protocol, as the private information was to be stored by the browser. The latest revision of P3P has removed this requirement in order

```
<!--This DTD describes the format of a valid TGT.xml file.-->
<!ELEMENT TGT
(Originator,Recipient,TrustedParty+,PrivateDataDefinition+)>
<!ATTLIST TGT
 Id ID #REQUIRED
 IssueDate CDATA #REQUIRED
 ExpiryDate CDATA #REQUIRED
>
<!--A trusted third party element, which describes the source
where the information pertaining to this TGT can be retrieved
from.--> <!ELEMENT TrustedParty EMPTY> <!ATTLIST TrustedParty
 URL CDATA #REQUIRED
 PublicKey CDATA #REQUIRED
>
<!--The owner of the private information to which access is
granted through this ticket--> <!ELEMENT Originator EMPTY>
<!ATTLIST Originator
 Id ID #REQUIRED
 PublicKey CDATA #REQUIRED
>
<!--The recipient being given access to the private information
through this TGT.--> <!ELEMENT Recipient EMPTY> <!ATTLIST
Recipient
 URL CDATA #REQUIRED
 PublicKey CDATA #REQUIRED
>
<!--A description of the private information being given access
to.--> <!ELEMENT PrivateDataDefinition
(Recipient+,TrustedParty?,PrivateData)> <!ATTLIST
PrivateDataDefinition
 ExpiryDate CDATA #REQUIRED
 UsageCount CDATA #IMPLIED
>
<!ELEMENT PrivateData EMPTY>
```

Figure 7.2: DTD describing the format of the Ticket Granting Ticket

```
<!--This DTD describes the format of a valid ticket.xml file.-->
<!ELEMENT Ticket
(Originator,Recipient,TrustedParty+,PrivateData+)> <!ATTLIST
Ticket
 Id ID #REQUIRED
 IssueDate CDATA #REQUIRED
 ExpiryDate CDATA #REQUIRED
>
<!--A trusted third party element, which describes the source
where the information pertaining to this ticket can be retrieved
from.--> <!ELEMENT TrustedParty EMPTY> <!ATTLIST TrustedParty
 URL CDATA #REQUIRED
 PublicKey CDATA #REQUIRED
>
<!--The owner of the private information to which access is
granted through this ticket--> <!ELEMENT Originator EMPTY>
<!ATTLIST Originator
 Id ID #REQUIRED
 PublicKey CDATA #REQUIRED
>
<!--The recipient being given access to the private information
through this ticket.--> <!ELEMENT Recipient EMPTY> <!ATTLIST
Recipient
 URL CDATA #REQUIRED
 PublicKey CDATA #REQUIRED
>
<!--A description of the private information being given access
to.--> <!ELEMENT PrivateData EMPTY>
```

Figure 7.3: DTD describing the format of the Ticket

```
<?xml version="1.0" encoding="UTF-8"?> <!DOCTYPE TGT SYSTEM
"TGTDTD.dtd"> <TGT Id = "example"
     IssueDate ="2001/01/01"
     ExpiryDate = "2002/01/01">
    <Originator Id = "Alice"
                PublicKey = "------ Begin PGP key block
                             fgfgf fgfdgdf
                             ----- End PGP key block"/>
    <Recipient URL = "http://www.onlineshop.com"
               PublicKey = "------ Begin PGP key block
                           fhytjmhgjh ghjkghjmcv
                           ----- End PGP key block"/>
    <TrustedParty URL = "http://www.megabank.com/trustserv"
                  PublicKey = "------ Begin PGP key block
                              jkghkg ghkgfcf
                              ----- End PGP key block"/>
    <PrivateDataDefinition ExpiryDate = "2001/12/01"
                           UsageCount = "1">
        <Recipient URL = "http://www.onlineshop.com"
                   PublicKey = "------ Begin PGP key block
                               fhytjmhgjh ghjkghjmcv
                               ----- End PGP key block"/>
        <PrivateData/>
    </PrivateDataDefinition>
</TGT>
```

Figure 7.4: Sample XML for a Ticket Granting Ticket

```
<?xml version="1.0" encoding="UTF-8"?> <!DOCTYPE Ticket SYSTEM
"TicketDTD.dtd"> <Ticket Id = "example"
        IssueDate = "2001/01/01"
        ExpiryDate = "2002/01/01">
    <Originator Id = "Alice"
                PublicKey = "------ Begin PGP key block
                             fgfgf fgfdgdf
                             ----- End PGP key block"/>
    <Recipient URL = "http://www.onlineshop.com"
               PublicKey = "------ Begin PGP key block
                            fhytjmhgjh ghjkghjmcv
                            ----- End PGP key block"/>
    <TrustedParty URL = "http://www.megabank.com/trustserv"
                  PublicKey = "------ Begin PGP key block
                               jkghkg ghkgfcf
                               ----- End PGP key block"/>
    <PrivateData/>
    <PrivateData/>
</Ticket>
```

Figure 7.5: Sample XML for a Ticket

to simplify P3P and speed up its adoption. 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 assured of availability, for instance if someone orders books from an online vendor, and disconnects from the Internet, making his 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 or long distance carrier 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. The security of this option could be increased by spreading the information over more semi-trusted third parties. This increases both the complexity and the trust of the solution.

7.4.2 Collusion Between Participants

Another important question to consider is that of collusion between some of the parties, as already mentioned in Section 7.2 — 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 organisation to risk).

7.4.3 Assurances

When making use of this protocol, 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 exists in the real

world when not using this protocol. 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.

7.4.4 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.

7.5 Related Work

Namesafe.com [40] and iPrivacy.com [26] mentioned in Chapter 4 attempt to act a go-betweens to protect the private information of a user. SET [53] is only applicable for credit card details, and does not protect any other aspect of an individual's privacy.

A quick comparison between these services and our solution can be seen in Figure 7.6.

Our protocol 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

Product	Namesafe	iPrivacy	SET	Our solution
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 7.6: Comparison of privacy solutions

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 the first four classes of information. 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.

7.6 Summary

We have created a protocol to protect private information in several of our classifications defined in Chapter 5.

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.

UNIVERSITY
OF
JOHANNESBURG

Chapter 8

Method 3 — Selective disclosure

In this chapter we propose a conceptual method to extend P3P in order to add more flexibility to the decision on whether or not a given item of private information will be supplied to a target organisation by using the Chinese Wall security policy. This will enable a user to not only define rules as to which items of private information she would disclose, but also to define what collection of private information any given organisation would be able to build about her.

Honesty's the best policy.
Miguel de Cervantes
Don Quixote de la Mancha [1605 1615], pt. II [1615], bk. III, ch. 33, p. 666

8.1 Introduction

In this chapter we shall develop a method to protect private information that is based on the Chinese Wall security model [8], allowing free access to all items of private information at first, but then restricting what can be accessed by a particular target organisation as they find out more and more about an individual.

One of the limitations of P3P is that there is no mechanism to control the total set of information disclosed to an organisation over time. We could allow an organisation access to our salary information, but only if we can remain anonymous. If any of our personal details have been supplied to the site, we might no longer wish to allow access to the salary information. This can currently not be represented using P3P, and is in fact an example ideally suited for the Chinese Wall security model. In this chapter APPEL is used to implement the Chinese Wall security policy in P3P.

This solution will be particularly well suited to protect information in class 5 of our classification, as presented in Chapter 5, since no additional steps are required by the target organisations and the private information can still be used for whatever purposes they are required, as no restrictions are placed on these.

8.2 A description of the Chinese Wall security policy

In the Chinese Wall security policy entities to be secured are grouped according to some common denominator and these groups are then placed into different conflict of interest classes. At the start of an exchange any entity in any group may be accessed by the target organisation. As soon as an entity in a particular group has been accessed, access to entities in other groups belonging to the same conflict of interest class is disallowed, and only entities in the same group as the first entity accessed can further be requested. All other entities belonging to groups in other conflict of interest classes can still be accessed. Further restrictions are then added as subjects access entities, until at some point a subject will only have access to those entities already accessed previously.

As an example, this model could be used to govern access to documents by a consultant in a consulting firm. In principle, a consultant has access to

all documents in the firm. As soon as the consultant has read a document on say, insurance firm A, access to all documents relating to insurance firms other than A is revoked. This is done to prevent possible conflicts of interest and decisions based on confidential information about competitors of firm A who may also be clients of the consulting firm. The consultant should still have access to other documents about firm A, and of course, all documents that do not contain information about insurance firms. In this example, the groups are defined by individual firms, and the conflict of interest classes are their sector of business. Some firms could belong to more than one conflict of interest class if their activities span several business sectors — if a consultant accesses documents about a firm B that is both an insurance firm and a bank, access to other insurance firms and other banking firms is revoked for that consultant.

In the context of privacy protection, the entities are a particular individual's private information items, and the groups are logical groupings such as identifying information, financial information and demographical information. Conflict of interest classes are arbitrarily defined by the individual to prevent any given organisation from learning too much about her.

8.3 Description of the method

Using the Chinese Wall security model to control access to private information items is very well suited to integration with P3P. Several aspects of P3P make such an integration relatively simple to implement. The grouping of private information required can be mapped directly to the P3P categories described below as defined in section 3.4 of the P3P specification [63]:

Physical contact information <physical/> Items such as a delivery address or phone number that make it possible to physically locate the user.

Online contact information <online/> Items such as email addresses or home pages that make it possible to contact the user on the Internet — this does not include computer specific information such as IP addresses.

Unique identifiers <uniqueid/> Non-government items generally issued by the Web site in order to uniquely and consistently identify the user, such as a username and password.

Purchase information <purchase/> Items generated through the purchase of products or services such as a credit card number or the method

of payment.

Financial information <financial/> Items such as a credit card number or cheque account number, as well as the balances and payment and transaction history. Information about individual purchases fall under the previous category.

Computer information <computer/> Items relating to the physical device being used, such as the IP address of the computer, operating system or browser type.

Navigation and click-stream data <navigation/> Items passively generated by browsing a website such as the referrer page and time spent visiting each page.

Interactive data <interactive/> Items actively generated through explicit interactions with a server such as server logs or shopping history that define the past interaction with the target organisation.

Demographic and socioeconomic data <demographic/> Items such as the gender, age or income of the user.

Content <content/> The words and expression contained in the body of a communication, such as the body of an email or IRC chat transcript.

State management mechanisms <state/> Items such as cookies used for maintaining a stateful session with the user.

Political information <political/> Information about a user's religious and political affiliations, as well as trade union and professional organisation memberships.

Health Information <health/> Items about a user's state of health or enquiries into health related services and health related purchases.

Preference data <preference/> Items about a user's individual preferences, such as favourite colour or hobbies.

Location data <location/> Items such as GPS position information that can determine a user's current location and track changes.

Government-issued identifiers <government/> Items such as a social security or passport number issued by a government to consistently identify an individual.

Other <other-category> Items not covered by any of the above-mentioned categories.

The only remaining set to define is to determine which conflict of interest classes these groups will be divided into. This would ideally be left to the user, as such a decision would vary from one person to another. An example of such a partition is given in Table 8.1.

8.4 Example partitioning

In this example, the user is willing to divulge computer information, navigation and click-stream data, interactive data, state and preference data regardless of which other categories of information is required, so each of these categories is placed into their own conflict of interest class (Allow01 – Allow05). The user would not want any site to access both purchase information and her financial information, but access to any one of these categories is allowed, therefore they are placed in the same conflict of interest class (Conflict01). Health and political information as well as demographics and socioeconomic data can be requested, as long as she can not be identified; however, sites that can identify her are allowed to have her physical and online contact information, as well as unique identifiers. In order to model this, conflict of interest classes Limited01 through to Limited33 are used. These conflict of interest classes are now added to the health, political and demographics categories to prevent simultaneous access to this information with identifying information. Access to content, location and government issued identifiers is not allowed under any circumstances, so normal APPEL rules can be used to implement this, and they do not have to belong to any conflict of interest class.

Note again that this partitioning defined in Table 8.1 is an example only, and that each user is likely to have a similar, but possibly different policy.

No changes to the server implementation of P3P is required, as target organisations still only have to specify their privacy policies based on which items of private information they would like to access, and what purpose they intend to use it for. The user agent would then have to implement the required checks to verify that private information requested would not break the Chinese Wall security model.

8.4.1 User Agent Implementation

The Chinese Wall security policy is implemented on the user agent side by making use of rules specified in APPEL, as described in Section 4.4.3 on page

Table 8.1: An example of a conflict of interest class definition

Private information group description	P3P Element	Conflict class
Physical contact information	<physical/>	Limited11, Limited21, Limited31
Online contact information	<online/>	Limited12, Limited22, Limited32
Unique identifiers	<uniqueid/>	Limited13, Limited23, Limited33
Purchase information	<purchase/>	Conflict01
Financial information	<financial/>	Conflict01
Computer information	<computer/>	Allow01
Navigation and click-stream data	<navigation/>	Allow02
Interactive data	<interactive/>	Allow03
Demographics and socioeconomic data	<demographic/>	Limited11, Limited12, Limited13
Content	<content/>	N/A
State management mechanisms	<state/>	Allow04
Political information	<political/>	Limited21, Limited22, Limited23
Health Information	<health/>	Limited31, Limited32, Limited33
Preference data	<preference/>	Allow05
Location data	<location/>	N/A
Government-issued identifiers	<government/>	N/A
Other	<other-category>	N/A

38. As APPEL has deliberately been kept simple in its current form (version 1.0), it is not possible to implement sophisticated rules. This means that the Chinese Wall security policy has to be implemented over several rules, and can not be done in a single complex rule.

The first step to implement a given partitioning of categories into conflict of interest classes in APPEL is to expand the possible combinations where a particular category is present in more than one conflict of interest class. Each possible invalid combination is then represented with a corresponding APPEL rule. The final step is then to add the Chinese Wall APPEL rules to the ruleset used by the user agent doing the P3P policy verification. We demonstrate this process by reusing our example partitioning from the previous section. The rules we wish to implement in APPEL in the second step of the process are:

1. Block access if both <physical/> and <demographic/> is required.
2. Block access if both <online/> and <demographic/> is required.
3. Block access if both <uniqueid/> and <demographic/> is required.
4. Block access if both <physical/> and <political/> is required.
5. Block access if both <online/> and <political/> is required.
6. Block access if both <uniqueid/> and <political/> is required.
7. Block access if both <physical/> and <health/> is required.
8. Block access if both <online/> and <health/> is required.
9. Block access if both <uniqueid/> and <health/> is required.
10. Block access if both <purchase/> and <financial/> is required.

We supply the actual XML for the first rule as an example in Figure 8.1. Lines 01 – 04 contain the header information. The first rule starts on line 05. Line 14 specifies the APPEL connective for the categories in the policy; in this case a normal "and" rule, implying that access to pages whose policies require both physical and demographical information be blocked. Other connectives defined in APPEL are "or" (if one of the listed items occur in the policy, evaluates to true); "non-and" (evaluates to true unless all the listed items occur in the policy); "non-or" (evaluates to true unless all the listed items occur in the policy); "and-exact" (evaluates to true if all the listed items occur in the policy, and the policy contains no other additional items from

```
01 <appel:RULESET xmlns:appel=
02 "http://www.w3.org/2002/04/APPELv1"
03  xmlns:p3p=
04  "http://www.w3.org/2000/12/P3Pv1">
05  <appel:RULE behavior="block"
06   description="rule 1 blocks simul-
07     taneous access to both physical
08     and demographical information">
09   <p3p:POLICY>
10    <p3p:STATEMENT>
11     <p3p:DATA-GROUP>
12      <p3p:DATA>
13       <p3p:CATEGORIES
14            appel:connective="and">
15        <p3p:physical/>
16        <p3p:demographic/>
17       </p3p:CATEGORIES>
18      </p3p:DATA>
19     </p3p:DATA-GROUP>
20     <p3p:RECIPIENT
21          appel:connective="or">
22      <p3p:ours/>
23     </p3p:RECIPIENT>
24    </p3p:STATEMENT>
25   </p3p:POLICY>
26  </appel:RULE>
27   <appel:RULE behavior="block"
28    description="rule 2">
29  </appel:RULE>
    :
    :
95  </appel:RULE>
96   <appel:RULE behavior="block"
97    description="rule n">
98  </appel:RULE>
99 </appel:RULESET>
```

Figure 8.1: APPEL XML rule to block simultaneous access to both physical and demographical information

that list); and "or-exact" (evaluates to true if one of the listed items occurs in the policy, and the policy contains no other additional items from that list). Line 26 ends the rule, and from lines 27 onwards further rules could be defined, until the ruleset is terminated at line 99.

Note that according to the APPEL specification, rules are evaluated in order, so if the user wishes to have the Chinese Wall rules override her normal P3P rule settings, they have to be placed at the start of the rule file; similarly if the user wishes the Chinese Wall rules only to fire if she did not specify other explicit P3P rules, they should be placed at the end of her P3P rule file.

One of the current limitations is that the collection of data over several visits with policy changes between visits could violate the integrity of the Chinese Wall. Some of the future enhancements to APPEL (extensible behaviours and matching external schemas) could be used to build external schemas containing a history of information required by a site's past policies, and this could then be used to implement the Chinese Wall in a persistent and consistent way across policy changes.

8.5 Summary

In this chapter we have presented a method to extend P3P in order to control the collection of private information each target organisation is allowed to access by using the Chinese Wall security model. This makes it possible to prevent any given organisation from learning too much about an individual user, whilst not needlessly restricting the information new organisations could access. It is possible to seamlessly implement this using the current specification of P3P and APPEL.

Using the Chinese Wall security model and P3P, we have devised a privacy protection mechanism that protects anonymity using conflict of interest classes, automates privacy policy verification using APPEL and P3P and prevents disclosure of information where such disclosure would violate the conflict of interest classes.

The biggest advantage to this method is that no changes are required on the server side to implement this functionality. All changes required are made on the client side. This has the added advantage that no protocols need to change, and no more bandwidth is required than for cases where only P3P is used.

Even on the client side, the changes required are minimal, as we have implemented this method using APPEL rules — any user agent capable of implementing P3P should be able to implement this method by just adding

the required rules to the ruleset.

UNIVERSITY
OF
JOHANNESBURG

Chapter 9

Classification revisited

In this chapter we take another look at our classification based on the solutions we have presented in the previous three chapters.

Do not turn back when you are just at the goal.
Publilius Syrus
Maxim 580

9.1 Introduction

In the previous three chapters we have developed different methods for protecting private information. We now revisit our classification and where these methods fit into it. We specifically compare what we have achieved to the goals we were aiming for in Section 5.7.

9.2 Protection evaluation

In this section we compare the goals from Chapter 5 with what is possible from Chapters 6–8.

9.2.1 Class 1 — Validation

Goal

In the case where private information is only required to validate a calculation or to generate statistics, we would ideally like to keep the information secret (secrecy) since the individual terms of the calculation are not directly required. Only the result of the calculation would have to be divulged (accessibility). Despite the individual terms being kept secret, the party to which the result is sent must be able to verify the correctness of the calculation (integrity).

Achievement

Using non-disclosure from Chapter 6 we are able to validate simple calculations without revealing the individual terms of the calculation, whilst being able to demonstrate the correctness of said calculation. For aggregates, the trusted third party from Chapter 7 could be used to supply these for all the individuals it is keeping private information of, without revealing the information used to generate the aggregates. On an individual level, a user could use a Chinese Wall setup from Chapter 8 to ensure that sites requiring private information for statistical purposes can request such such information, but without identifying the user, thereby protecting the user's privacy.

Whether through non-disclosure or by retaining control, the secrecy of the information has been improved. When using selective disclosure, the secrecy has improved slightly in the sense that the user would not be identified

without his explicit permission through the Chinese Wall. The accessibility and integrity of the information have not been affected.

9.2.2 Class 2 — Third party use

Goal

For our second class, where information is requested in order to pass it on to a third party which requires it, we would also like to hide the information from the organisation (secrecy) since it does not need the information itself. The third party requiring the information should be able to access it (accessibility). Updates to the information should also filter through to the third party directly, without the organisation being involved (integrity).

Achievement

With the method developed in Chapter 7, we would only allow the actual third party requiring the information to access it. The organisation transacting with the user would be issued a TGT that only allowed it to request a ticket for the third party, which would then have to get the private information directly from the trusted third party, without the target organisation ever having had access to it. If such information were to be updated after the TGT (or even the ticket) has been issued, but before the actual access, the eventual access would be the latest version.

The secrecy of the private information has been improved in that fewer organisations have direct access to the data. The accessibility has been unaffected, whilst the integrity has been improved — a single update (at the trusted third party) would allow every party with access to get the latest version of the private information.

9.2.3 Class 3 — Future use

Goal

In cases where private information is requested by an organisation for future use, we would like to reveal the information only when it is actually required, and not ahead of time (secrecy). When the information is then actually required, the organisation should have guaranteed access to it (accessibility). Such information should be up to date, in case it changed since access was granted but before access is required (integrity).

Achievement

Using the third party from Chapter 7 and a valid TGT, an organisation would be able to access the data when required, and not need to have it right away. As long as the TGT has not expired, the organisation would be able to request tickets from the trusted third party, and then use these tickets for access.

The secrecy of the private information has been improved as the organisation requiring the data has access to it only when required, and not ahead of time. The accessibility has not been affected, but the integrity has significantly improved, as the organisation would automatically benefit from any interim updates to the information from when access is requested until it is eventually accessed.

9.2.4 Class 4 — Identification

Goal

Where private information is required for identification purposes only, it would be best not reveal the information at all (secrecy) but to establish some other mechanism for identification that is reliable (accessibility) but that does not require knowledge of private information about the individual. Identification should be possible without false matches or invalid rejections (integrity).

Achievement

If the method of retaining control from Chapter 7 is used, identification of a user could be based on knowledge of her private key, which is a very safe way of identifying her, as no private information is required, and as long as her private key is safe, impersonation is very hard. The Chinese Wall from Chapter 8 could also be used, if the user so wishes, to ensure that a site could only access unique identifiers (<uniqueid/>) and nothing else (for say, information only sites).

If the method of retaining control is used, secrecy is improved in that the information used to identify a user can not be reused and has no significance in itself (typically a random number and the time encrypted with her private key). Accessibility and integrity are also improved in that the probability of a user forgetting her unique identifier, or using another one is significantly reduced. Where the selective disclosure method is used, security could be improved depending on the user's Chinese Wall setup, but accessibility and integrity would be unaffected.

9.2.5 Class 5 — Direct usage

Goal

In the fifth class where private information is directly required for some purpose, the secrecy requirement can not be fulfilled, implying that the information would have no enforceable security. Secrecy in these cases would be achieved through privacy policies, legislation and trust. Accessibility is good if the organisation stores its own copy of the information, but this compromises secrecy even further and destroys integrity. If the organisation does not store a copy but requests the information every time it is required, accessibility is compromised, although integrity is good.

Achievement

In cases where the retaining control method from Chapter 7 is being used, the organisation requiring direct access might not have a direct interest in the transaction (such as the shipping company) thereby improving the secrecy of the data slightly, whilst not affecting the accessibility. Even if the organisation requiring the access is the one involved in the transaction, the integrity of the data is improved, as any updates would automatically filter through.

In this class the selective disclosure of private information from Chapter 8 is most useful, as a user could ensure that no organisation can build up a collection of data that would violate his Chinese Wall setup. In this way the secrecy of some private information items would be improved, but this would rather lead to an overall improvement in the secrecy of the user profile. Accessibility and integrity would not be affected.

9.2.6 Class 6 — Other

Class 6 is mentioned for completeness only, as no protection mechanisms was developed for this class.

9.3 Summary

In this chapter we have revisited the classification developed in Chapter 5 and evaluated the protection mechanisms developed in this thesis for each class. We have measured these protection mechanisms against the properties of secure data as described in Chapter 3. The protection mechanisms we have

developed have significantly enhanced the control and security (and therefore the enforceability) of the privacy of a user's private information.

UNIVERSITY
OF
JOHANNESBURG

Chapter 10

Conclusion

In this final chapter we review the state of the work, what we have achieved, and possible future enhancements.

A whole is that which has beginning, middle, and end.
Aristotle
Poetics, ch. 7

10.1 Introduction

This is the last chapter of this thesis, and as such reflects on the results we have achieved, how the problem we have tried to solve has been addressed and we comment on areas for future research.

10.2 Results

In this thesis we aimed to remedy the lack of enforceability regarding privacy protection, and to reach a workable compromise between total privacy (and an impossibility to transact on the Internet) on the one hand, and total disclosure of private information (and a total loss of privacy).

For a more effective compromise between these extremes, we have developed a classification of private information based on the purpose it has been acquired for. With this classification, we have developed three mechanisms to improve existing privacy protection, namely nondisclosure of private information, retaining control over private information that is disclosed as far as possible, and selectively disclosing private information in order to prevent any one organisation from learning too much about one individual.

The results and methods in this work have been summarised in a number of papers. The following papers have been published:

- "Enforcing Privacy by Withholding Private Information", based on Chapter 6 [32].

- "On Granting Limited Access to Private Information", based on Chapter 7 [33].

The following paper has been accepted for publication:

- "A Chinese Wall Approach to Privacy Policies for the Web", based on Chapter 8.

The following paper has been submitted for publication:

- "PrivGuard: A Model To Protect Private Information Based On Its Usage", based on Chapter 6, Chapter 7 and Chapter 5.

10.3 Future research

Each of the methods we have developed in this work could be enhanced further. We take a quick look at the possibilities for enhancement.

10.3.1 Nondisclosure

More problems where this type of solution is applicable could be researched, and the solution could then be expanded to include more mathematical functions.

It would also be quite useful if a single encryption function that is homomorphic for both addition and multiplication could be designed. This does not have to be a public key function; a one way encryption function could also be suitable. Such a function would transform this solution from a theoretical solution into a practical, easily implementable solution, and could make a significant contribution to electronic privacy.

10.3.2 Retain control

Further study could possibly integrate this approach with P3P, in order to automate this process and make it totally transparent to the end user. A suitable method to integrate the trusted third party communication seamlessly into the browser just like SSL certificate verification would also be a great step forward.

10.3.3 Selective disclosure

The next step would be to create a utility that will help a user to partition her groups into conflict of interest classes, and to automatically then generate the rules for embedding into the user agent. The Chinese Wall could also be made more fine grained, such as using P3P data items instead of just the categories when building the Wall.

10.4 Final Words

As long as private information has to be supplied by users, privacy on the Internet will not be complete, and areas for privacy research will exist.

Exiting new ways to discover information about individuals, such as timing attacks to determine whether or not a particular (unrelated) page has

been visited by a user recently [19] open new vulnerabilities for which there are as yet no protection.

A researcher's job is never done ...

UNIVERSITY
OF
JOHANNESBURG

List of Figures

List of Tables

UNIVERSITY
OF
JOHANNESBURG

Bibliography

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

[2] Paola Benassi. TRUSTe: an online privacy seal program. *Communications of the ACM*, 42(2):56–59, 1999.

[3] Tim Berners-Lee, Robert Cailliau, Ari Luotonen, Henrik Frystyk Nielsen, and Arthur Secret. The world-wide web. *Communications of the ACM*, 37(8):76–82, 1994.

[4] Tim Berners-Lee et al. Uniform resource locators (URL). RFC 1738, December 1994. RFCs available on http://www.rfc-editor.org.

[5] Tim Berners-Lee et al. Uniform resource identifiers (URI): Generic syntax. RFC 2396, August 1998. RFCs available on http://www.rfc-editor.org.

[6] Marjory S. Blumenthal and David D. Clark. Rethinking the design of the Internet: The end-to-end arguments vs. the brave new world. *ACM Transactions on Internet Technology*, 1(1):70–109, 2001.

[7] Ephraim J. Borowski and Jonathan M. Borwein. *Dictionary of Mathematics*. Collins, 1989.

[8] David F.C. Brewer and Michael J. Nash. The Chinese Wall security policy. In *IEEE Symposium on research in security and privacy*, pages 206–214, 1989.

[9] Duncan Campbell. Global surveillance: the evidence for Echelon. In *Computers, Freedom and Privacy: challenging the assumptions*, pages 149–154. ACM Press, 2000.

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

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

[12] Crowds home page. http://www.research.att.com/projects/crowds/index.html.

[13] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, IT-22(6):644–654, 1976.

[14] Hans Dobbertin, Antoon Bosselaers, and Bart Preneel. RIPEMD-160: A strengthened version of RIPEMD. In *Fast Software Encryption*, pages 71–82, 1996.

[15] Josep Domingo-Ferrer and Jordi Herrera-Joancomatí. A privacy homomorphism allowing field operations on encrypted data. *I Jornades de Matemàtica Discreta i Algorísmica*, 1998.

[16] Doubleclick home page. http://www.doubleclick.com.

[17] Echelon watch. http://www.echelonwatch.org/.

[18] Payment card chargeback when paying over the Internet, July 2000. http://europa.eu.int/comm/internal_market/en/ecommerce/chargeback.pdf.

[19] Edward W. Felten and Michael A. Schneider. Timing attacks on web privacy. In *Proceedings of the 7th ACM conference on Computer and communications security*, pages 25–32. ACM Press, 2000.

[20] Roy T. Fielding et al. Hypertext transfer protocol – HTTP/1.1. RFC 2616, June 1999. RFCs available on http://www.rfc-editor.org.

[21] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, 1979.

[22] Ian Goldberg and David Wagner. Randomness and the netscape browser. *Dr Dobbs Journal*, January 1996.

[23] Rüdiger Grimm and Alexander Rossnagel. Can P3P help to protect privacy worldwide? In *Proceedings of the 2000 ACM workshops on Multimedia*, pages 157–160. ACM Press, 2000.

[24] David Harel. *Algorithmics: the spirit of computing*. Addison-Wesley Publishers Limited, 1987.

[25] Michael Holt. *The Pan Pocket Puzzler*. Pan Books Ltd., 1985.

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

[27] Internet domain survey, January 2001. http://www.isc.org/ds/.

[28] John Kelsey, Bruce Schneier, David Wagner, and Chris Hall. Cryptanalytic attacks on pseudorandom number generators. *Lecture Notes in Computer Science*, 1372:168–188, 1998.

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

[30] Dave Kristol and Lou Montulli. Http state management mechanism. RFC 2965, Octber 2000. RFCs available on http://www.rfc-editor.org.

[31] Frans A. Lategan. Complexity aspects of certain graphical parameters. Master's thesis, Randse Afrikaanse Universiteit, 1995.

[32] 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.

[33] 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.

[34] Wayne Madsen. IRS fails to provide adequate security, privacy controls. *Computer Fraud & Security*, page 8, February 1999.

[35] George Marsaglia. A current view of random number generators, 1984.

[36] George Marsaglia and Arif Zaman. Monkey tests for random number generators, 1993.

[37] Meg McGinity. Staying connected: surfing the turf. *Communications of the ACM*, 43(4):19–21, 2000.

[38] Ryan Moats. URN syntax. RFC 2141, May 1997. RFCs available on http://www.rfc-editor.org.

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

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

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

[42] George Orwell. *1984*. Harcourt Brace Jovanovich, Inc., 1949.

[43] Steven Pemberton et al. XHTML 1.0: The extensible hypertext markup language, January 2000. http://www.w3.org/TR/xhtml1/.

[44] Charles P. Pfleeger. *Security in Computing*. Prentice-Hall, 1989.

[45] David Raggett, Arnaud Le Hors, and Ian Jacobs. HTML 4.01 specification, December 1999. http://www.w3.org/TR/html401/.

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

[47] Michael K. Reiter and Aviel D. Rubin. Anonymous web transactions with crowds. *Communications of the ACM*, 42(2):32–48, 1999.

[48] Ronald L. Rivest, Adi Shamir, and Leonard Adleman. A method for obtaining digital signatures and public key cryptosystems. *Communications of the ACM*, 21(2):120–126, February 1978.

[49] Bruce Schneier. The doghouse: Bungled SSL, 1999.

[50] Bruce Schneier. Echelon technology, 1999.

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

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

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

[54] Barbara Simons. Building Big Brother. *Communications of the ACM*, 43(1):31–32, January 2000.

[55] Simon Singh. *The Code Book*. Fourth Estate Limited, 1999.

[56] 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.

[57] TRUSTe home page. http://www.truste.org.

[58] Versisign home page. http://www.verisign.com/.

[59] Voila translations. http://trans.voila.fr/textrad/.

[60] Eugene Volokh. Personalization and privacy. *Communications of the ACM*, 43(8):84–88, August 2000.

[61] Extensible markup language (XML). http://www.w3.org/XML/.

[62] A P3P preference exchange language 1.0 (APPEL1.0). http://www.w3.org/TR/P3P-preferences.

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

[64] Web accessibility initiative (WAI). http://www.w3.org/WAI/.

[65] World-wide character sets, languages, and writing systems. http://www.w3.org/International/.

[66] David Wagner and Bruce Schneier. Analysis of the SSL 3.0 protocol. In *The Second USENIX Workshop on Electronic Commerce Proceedings*, pages 29–40. USENIX Press, 1996.

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

[68] Robert H. Zakon. Hobbes' Internet timeline. RFC 2235, November 1997. RFCs available on http://www.rfc-editor.org.

Index