Obscured by Clouds

On the Privacy Implications of Cloud Computing

Joss Wright

joss.wright@oii.ox.ac.uk

Oxford Internet Institute, University of Oxford.
1 St. Giles, Oxford. OX1 3JS

Abstract

Cloud computing aims to make computational resources available to consumers as a third-party service, analogous to existing electricity grids or telecommunication networks. The resulting market of dedicated providers is intended to take advantage of economies of scale to provide effectively infinite levels of computing power to customers on demand, without the need for expensive hardware and software setup and maintenance costs.

In contrast to its obvious advantages, cloud computing raises a number of significant security issues that are quickly becoming the focus of active research[15]. Beyond this, however, cloud computing raises a wide range of severe privacy issues that may ultimately prove to have far broader day-to-day significance for end-users.

This paper explores the implications that widespread cloud computing will have for personal and organisational privacy, and suggests ways that privacy risks can be avoided or mitigated. By considering these issues at an early stage of the development of this new technology, we can improve the chance that users will be able to reap the benefits of cloud computing without sacrificing their right to control over their personal information.

Keywords: cloud computing, privacy, PETs

1 Introduction

> "Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."
>
> – NIST Definition of Cloud Computing[8]

Information technologies, and the communication infrastructure that has been built around them, are critical to many aspects of the global society. Governments, enterprises and individuals rely on computer networks for communication, storage and information processing.

Cloud computing aims to provide information technology services through dedicated third-parties, rather than having them developed and maintained internally by organisations. These dedicated computational service providers can exploit economies of scale in order to provide a level of computational resources infeasible through traditional approaches. Cloud computing can be viewed as a means of providing computational resources through a similar model to that of electrical grids and telecommunication networks.

Of particular note for our focus in this paper is that a user of a cloud computing service need not have any knowledge regarding the nature, location or even ownership of the underlying hardware, much as we are typically unaware of the power plant that generated the electricity powering our homes.

The "cloud", from this perspective, is therefore an abstracted pool of software and services provided by third parties that an organisation, or individual, can access remotely. In its idealised form, the user of a cloud computing service is detached from concerns about where and how resources are allocated; all actions occur automatically through a standardised interface that provides the services required, as required and when required. In practice this dream is likely to be more complex and, as we explore in this paper, will raise its own sets of problems and concerns.

1.1 Models of Cloud Computing

There are a number of approaches commonly taken towards realising cloud computing, corresponding to different levels of abstraction from the underlying virtual hardware. These approaches lend themselves to different types of use, and will be appropriate for different tasks as required by each organisation.

Cloud computing approaches range from the provision of a software application that runs remotely, to providing access to raw virtual machines that can run custom software or be combined to create a range of further services. We will briefly detail the major forms of cloud computing approaches here.

1.1.1 Software as a Service

At its highest level of abstraction, cloud computing can simply provide end-users with software applications that run on a remote host. Interaction with these applications typically takes place through a web browser, allowing users access to applications from any location with an Internet connection, and without requiring the user to have dedicated hardware.

This approach, generally referred to as *software as a service* (SaaS), is the form of cloud computing most likely to be directly experienced by end-users and is already in widespread use. To some extent, software as a service can be considered the logical extension of existing interactive Web technologies.

1.1.2 Platform as a Service

Platform as a service (PaaS) provides developers with lower-level access to resources in the cloud through components or services that can be used to build and deploy software. The interfaces used to develop applications allow for distribution and scalability of resources while abstracting away from the lower-level management of those resources.

These technologies are targeted at easy development of applications in the cloud, sacrificing some flexibility in the capabilities of the underlying platform for developer convenience.

1.1.3 Infrastructure as a Service

Infrastructure as a service (IaaS) provides low-level access to cloud computing resources. In this approach, the client has full access to *virtual machines* hosted remotely, and can directly develop and deploy software in a similar manner to traditional computing development. In effect, this form of cloud computing gives the client access to a large server farm on which custom software applications can be dynamically deployed as required.

One advantage of cloud-based infrastructure, besides the likely cost savings, is the capability of the cloud to provide dynamic, or *elastic*[7], resource provision. The virtualised nature of cloud infrastructure services allows virtual servers to be added to, or removed from, the client's resource pool as required. Enterprises therefore need only make use of the computing power required at a given moment, rather than maintaining the maximum required capacity.

Each of the approaches to cloud computing described above provide different levels of convenience to the user, different potentials for the provider and different security and privacy concerns in terms of storage, sharing and access to the personal information of users. In this paper we focus on broader privacy concerns across the various forms of cloud computing, and examine ways in which organisations can seek to exploit the benefits of the approach whilst preserving user privacy.

2 Privacy

Privacy is studied in many fields, and the interpretation of the term varies accordingly. In computer science, much research into privacy arose from computer security, focusing on control of access to information and the tools that achieve this. Even in this restricted context, however, privacy is a term that is far from having an accepted definition. It is therefore useful to define our particular interpretation of privacy before exploring its implications in cloud computing.

Without attempting to present a formal definition, we consider privacy in this paper from the viewpoint of *informational self-determination*[14], focusing on the ability of an individual to maintain control over

data describing its uniquely identifying characteristics and the history of its actions. This interpretation implies a relatively broad view of private data that includes social and business relationships, the sequence of actions engaged in by an individual and the range of identities that individual presents in their interactions, as well as more direct personal characteristics.

This notion of privacy focuses almost exclusively on individuals. In this paper we will not consider the potential for information regarding organisations to be considered private, except in the context of an individual's interaction with that organisation. The notion of private information should also not be conflated with *confidential* information, relating to all forms of information that an entity wishes to remain secret regardless of its nature.

For the purposes of this paper, therefore, privacy in cloud computing concerned with enabling individuals to control data describing their identities and actions in the cloud. To best understand how this control can be maximised we now explore particular risks to privacy that are inherent in the dream of cloud computing, and suggest ways in which these privacy risks can be alleviated.

3 Privacy Implications of Cloud Computing

The majority of the fundamental privacy issues in cloud computing stem from data being stored outside of the control of a single organisation. This lack of direct control opens channels through which information can be released, intentionally or not, to third parties.

A significant and less obvious addition to shared data is that the *actions* of a user in the cloud can be observed and logged by any entity involved in carrying out the user's actions, such as the user's Internet service provider and the providers of cloud services with whom they interact. This results in an explosion of data regarding an entity's actions in the cloud, various forms of which can persist for some time in log files and transaction records. Over time, this stored data can reveal underlying trends and behaviours of individuals that can be exploited in ways that are difficult to foresee.

Current computing approaches typically store and process sensitive information within an entity's direct sphere of control. For individuals this will typically be their home PC, while for a larger organisation this will typically take the form of servers and workstations physically located within the organisation.

This provides organisations the opportunity to restrict access within a very clearly defined zone of trust. In cloud computing, where storage and processing of customer data no longer occurs exclusively within the organisation, this zone of trust must be extended to one or more third parties.

This sharing of data is largely a one-way process. Data, once shared, is extremely difficult to retract verifiably. Even when legal, policy-based or technical solutions are employed, the ease and undetectability of copying data require only a single malicious or careless employee for data to be lost.

Outsourcing of computational resources is certainly not a new phenomenon. Many individuals and small businesses host their websites through dedicated third-party hosting providers, small companies use third-party services to handle payment processing and order tracking and remote backup and storage systems are increasingly common. Cloud computing is not fundamentally different to these existing services, but takes them to the logical conclusion in which the majority of services are remotely located, even for major companies.

Remote storage of customer data, however, is merely the simplest and most obvious of the privacy implications raised by cloud computing approaches. We now examine more complex implications that the cloud computing approach can have for individuals.

3.1 Observable Activities

A home user typically runs software that is largely stored and processed on their local machine, accessing the Internet in order to retrieve data and communicate with the outside world. The moment-to-moment activities of that user, however, largely take place on the local system, even when they require interaction with the Internet. While editing documents and images, playing games and listening to music increasingly require an Internet connection, this requirement is mainly concerned with transmission of content, with processing and storage occurring on the local machine.

Cloud computing seeks to break down the barrier between the user's machine and the cloud of remote service providers. Documents, images and videos are not only stored and made available online, the software used to create and edit them is also located and run remotely.

This adds a layer of convenience to the user, who no longer requires the purchase and maintenance of a powerful machine. From a privacy point of view, however, the provider of a cloud computing service now gains a detailed view into the user's activities. The provider no longer sees only uploaded documents, images and videos, but is involved with all aspects of the creation process, including early drafts and minor edits.

Beyond this, usage profiles and access times reveal when a user worked on a document, when they stopped to eat and when they went to sleep. [17] has demonstrated that monitoring the home electricity supply can reveal a surprising amount of information regarding the habits of an individual; when extended to their Internet usage and detailed online activities, the amount of information revealed has the potential to seriously compromise their privacy.

3.2 Observable Interactions

Increased scrutiny of user activities by a cloud provider can cause a serious privacy risk. From an outsider's point of view, however, even a user's relations with different cloud services can potentially reveal a great deal of information about them to an observer.

In cloud computing, users develop an increasing number of "critical" relationships with service providers as they rely on them not only to provide feeds of information or entertainment, but also with software that we require in day-to-day life. This software is not, as in a traditional model, downloaded to a local machine and then run, it relies on a constant interaction between the provider and the consumer.

The effect of this, where not only data but also software are located remotely, is that an observer of the flow of traffic to and from a user's computer can gather information about the type of software installed on the user's computer and when that software is accessed. While this information may not always be critical, knowledge of a user's interaction with medical or financial services can be considered a serious breach of privacy.

3.3 Service Aggregation

In making use of a cloud computing service, an end-user accepts that their customer information and data is stored on the service provider's system. The user accepts, too, that the service provider gains a finer-grained view into their activities and habits, and that an observer with the capability to view their Internet connection may be able to determine some aspects of their behaviour.

A service provider, however, is unlikely to provide a single service or product. Larger-scale providers will offer suites of software that perform interrelated tasks and, as with existing suites of software, using a single provider's products is likely to be more convenient in terms of interoperability than selecting from a range of competing products.

A user selecting a range of software from a single provider leads to that service provider gaining information not only on the user's activities in a single application, but on the whole range of applications that it supplies. As with many privacy issues, the value of such aggregated data is significantly greater than the sum of its parts.

3.4 Infrastructure Aggregation

Aggregation of services allows a cloud software provider access to the data and usage patterns of user across its range of services, potentially revealing far greater amounts of information than the user might expect. This risk can become even more significant, however, when large cloud computing entities act not as providers of software to users, but as providers of resources to organisations.

The model of *infrastructure as a service* is designed to allow organisations direct access to low-level computing services such as virtual machines[11]. These can be used by organisations to deploy their own software services, including higher-level cloud applications.

Infrastructure providers in this scenario are in a position to observe the behaviour of customers of those organisations that build on their infrastructure. These customers will typically not be aware of the infrastructure providers underlying the services that they access, nor will they have a direct legal relationship with them. A particular concern in this scenario is that an individual accessing two entirely separate cloud

application providers may unwittingly be communicating with the same underlying cloud *infrastructure provider*.

The privacy of customer data in these scenarios will almost certainly be protected by service-level agreements, and an infrastructure provider is unlikely to have legal permission to make use of this data. Despite this, both service and infrastructure aggregation create centralised points for observing and tracking users that provide tempting targets for hackers, or simply untrustworthy employees, seeking to gather large amounts of private information.

3.5 Personal Information Economy

The model of computation as a service is attractive in many ways, with significant advantages for both consumers and providers. As a logical extension of existing Internet services, however, it is likely that some service providers may choose not to charge users directly but instead to derive revenue from targeted advertising and sale of customer data.

Treating personal information as a currency for access to services is tempting to end users, who typically have relatively low concerns for privacy violations, and to providers who benefit from a seemingly "free" service.

One of the more troubling outcomes of this trend is that privacy-conscious users can find it extremely difficult to opt out of popular services that require them to submit personal information. When particular services become near-universal, users can find it difficult to interact in business or in personal life without being a member of those services, as is already seen in suites of office software and social networks. Cloud computing can exacerbate this situation by extending the reach of payment through personal information much more deeply into users' lives.

These *natural monopolies* can be reduced by promoting open software standards that allow interoperability between different services and providers, and thus create a market of alternatives to any given product. Unfortunately, this is not typically considered by a service provider to be in their best interest. History has shown that many large software providers prefer to maintain their own proprietary approaches.

3.6 Required Participation

While individuals can find it highly inconvenient to avoid dominant products, there exist services in which users have no choice but to take part. This is particularly true of government services, which may also be major repositories of highly sensitive personal information.

When governments choose to outsource their infrastructure to the cloud, which has already begun to occur, the associated privacy and security risks are transmitted to citizens. Clearly, government-level services are an especially tempting target for hackers due to the importance of the information that they hold.

As with all software and hardware in use by governments, cloud computing providers considered for governmental use must be expected to meet rigorous criteria for security and privacy. The development of these criteria must take into account not only the information stored on systems, but also their interactive behaviour in the cloud.

3.7 Cloud Security

One of the fundamental shifts that cloud computing promotes comes from the use of software that is entirely installed and configured by the cloud provider on their systems, without relying on less experienced home users. While, as has been discussed, this reveals far more about computer usage and content of personal data to the provider itself, it also has the potential to decrease the risks of misconfigured or insecure software running on the user's computer. While viruses, trojan horses and spyware all remain possible in cloud computing, the specific focus of these risks shifts towards targeting providers rather than end users.

The most significant *negative* aspect of this reliance on server-side software is that any breach of the software provider will reveal vast amounts of personal information. It is therefore possible that, as cloud computing becomes more widespread, privacy breaches will become fewer in number, but more dramatic in scope.

4 Model-specific Issues

In the previous section we examined the broad-scale privacy implications of cloud computing. We now briefly consider the privacy concerns that arise in each of the major approaches to cloud computing.

4.1 Software as a Service

Software as a service is the most visible form of cloud computing to end users, providing direct access to software applications in the cloud. This has the potential to partially replace the existing model of users installing software on their machine in favour of connecting to remote applications.

Users have little control over how their information is stored and shared by providers in this model, beyond examining stated privacy policies of the service. Adherence to these standards can be extremely difficult to verify as releases of private information are rarely detectable. Despite this, users of online services in recent years have demonstrated increasing concern for well-designed privacy policies. Protests by users have proven sufficiently powerful to effect changes in the policies of large organisations.

The nature of this model of software provision, and the relative lack of power and scrutiny that users have over the software provider, does intrinsically require a high level of trust from users. In order for this trust to be in some way justified, mechanisms to ensure compliance with policies and legal mechanisms must be developed, along with strong incentives for software interoperability to prevent user lock-in.

4.2 Platform as a Service

Platform as a service systems allow clients to develop cloud computing software at a high level of abstraction, without requiring underlying access to infrastructure details.

This model presents an interesting set of design problems to privacy-conscious developers. The greatest concern in this model is that low-level details of where and how information is stored are abstracted away from the developer's direct control, and can thus leave open security or privacy flaws in lower-level code on which the software is built. This can take the form of unnecessary or observable communication between hosts, or insecure storage or replication of data.

In developing software at this level, the most significant principle should be to adhere to *data minimisation*[13], in which no private data is stored unless absolutely necessary. By removing unnecessary information, a developer can greatly mitigate unforeseen risks. While this approach should be considered fundamental to the design of all privacy-conscious software, the lack of direct control over underlying mechanisms makes it particularly critical in this case.

4.3 Infrastructure as a Service

Infrastructure as a service is, to a large extent, invisible to the normal cloud computing user. Through having direct access to low-level infrastructure details, a developer in this model of cloud computing gains fine-grained control over many aspects of how data is transmitted and stored. This allows for greater flexibility in the design of privacy-friendly applications and architectures.

Despite this potential, infrastructure as service does not necessarily provide a developer with full control over the location of their resources, or the specific details of how the virtual resources are allocated. As with all approaches to cloud computing, resources are allocated by a third party who has full access to all data and computation performed on the service. This leaves the ultimate responsibility for data security, and therefore data privacy, with the underlying cloud provider.

5 Legal Considerations

This paper focuses mainly on *technical* aspects of privacy in cloud computing, however it would be wrong to ignore entirely the range of *legal* issues as well as applicable legal protections that affect cloud computing. We will briefly examine these here.

The individual right to privacy is widely recognised in societies across the world, and is enshrined in the United Nations Universal Declaration of Human Rights[10]. Specific recognition of privacy in law is, however, relatively recent in many Western countries[18], and its specific application is still the subject of much debate.

The interaction between the largely global Internet and local laws is a perpetual source of concern for new technologies, and cloud computing is certainly no exception to this.

Perhaps the most significant legal concern for cloud computing is that private data entering the cloud has the potential to pass between legal jurisdictions that present vastly different protections for privacy. It may even be the case, with the abstraction of computing resources made possible by cloud technologies, that a developer is unaware of where the data gathered by their application may be stored. In the case of European privacy law, which states that private information cannot be transmitted to countries with incompatible privacy protection laws, a cloud computing application may experience great difficulty in fulfilling their legal obligations[16].

Legal mechanisms are therefore of critical importance to enable cloud services to function effectively. Further, these mechanisms must be developed alongside the cloud technologies themselves in order to provide an adequate framework in which these technologies can function effectively.

6 Privacy Solutions

We now briefly examine principles and techniques that can be used to mitigate privacy issues in cloud computing.

A fundamental principle of privacy, and arguably the most effective, is that of *data minimisation*; data should not be gathered, stored, accessed or shared unless it is strictly necessary for the functioning of the service. It has been repeatedly shown[6, 9] that even seemingly trivial storage of data can result in serious privacy breaches when data is aggregated and analysed.

A second principle that is key to preserving privacy in cloud computing is that once data has been shared with a third party, the data cannot feasibly be kept secret. Laws and privacy policies that require providers to delete data after a certain period provide useful guidelines, and allow for organisations to be penalised if data is not handled appropriately; such methods, however, can only be enforced once sharing of data has occurred, and only then when the sharing has been detected.

Technical solutions, on the other hand, aim to prevent the *possibility* of unauthorised data sharing. A range of technologies are available that, in isolation, provide strong guarantees that data *cannot* be misused. Unfortunately, many theoretical techniques are fragile, inconvenient, computationally expensive and difficult to combine into a functional service.

6.1 Homomorphic Encryption

Encryption, the mathematical transformation of information in order to prevent its being read by any party not possessing a secret key, is a fundamental building block of the modern Internet. Encryption is used in the transmission of data, and has limited use in cloud computing for purposes such as storage of static information.

In 2009, Gentry proposed a *fully homomorphic* encryption scheme[4], positively resolving a long-standing fundamental challenge in modern cryptography. Gentry's scheme has received a great deal of attention as the solution to a wide range of security and privacy issues. We will briefly examine its potential, and its limitations, here.

Simply put, a fully homomorphic encryption scheme allows a third party, such as a cloud computing provider, to perform arbitrary operations on encrypted data without the need to decrypt that data. The advantages of this for cloud computing are clear: a third party can be contracted to process data for a customer without ever learning the content of the data. This is a great advantage to businesses that are uncomfortable sharing their sensitive intellectual property with third parties.

Homomorphic encryption is by no means a universal solution for the protection of user privacy. With the assumption that such a scheme becomes practical, in contrast to the best known approach which is many orders of magnitude too inefficient to be considered for real-world use, homomorphic encryption solves only one of the many ways in which private information regarding users can be compromised.

As detailed above, user privacy extends beyond the content of customer databases and remote back-ups. Encrypting content still allows observers, and cloud providers, to learn communication partners, transaction histories, usage profiles and many other aspects of a user's actions*DiffieL08*. Homomorphic

encryption therefore, if and when it becomes a practical technology, will be only one part of the larger approach towards protecting privacy in the cloud.

Aside from the various forms of data encryption, highly privacy-sensitive cloud computing services may take advantage of a range of recent developments from cryptographic research such as anonymised communication systems[2, 3, 1], that hide the relationships between consumers and providers; private information retrieval[12], that allows clients to retrieve information from databases without the database owner learning which records were accessed; and multi-party computation, which allows clients to spread their computational resources over a number of providers, none of whom learn the full details of what was computed[5]. These technologies are largely theoretical constructions at present, but the demands of cloud computing may see their increasing use in practical systems.

7 Conclusions

Cloud computing is an exciting technology that provides a number of practical benefits for end users and large enterprises alike. The increasing trend towards remotely hosted services means that cloud computing is likely to play a large role in the future of computing, but as with any new technology its final form is unlikely to resemble our current view of it.

Despite its advantages, however, cloud computing presents a large number of extremely serious concerns for individual privacy. Remote storage of data is the most obvious of these concerns but, as we have described, there are many others. Sharing data with third parties is intrinsically opposed to privacy; with enterprises that view personal information as source of revenue, privacy is under an even greater threat.

Further, it is unlikely in the long run that we will be able to opt out of cloud computing. The providers of our services, from a commercial and a governmental point of view, will increasingly use the technology at all levels. It is therefore critical that cloud-focused privacy-enhancing technologies be developed alongside other cloud computing technologies, and are considered a fundamental part of their design.

From a user point of view, it is important that we minimise the amount of information that we release. It is highly unlikely, based on existing examples of network services, that users will consider privacy a sufficiently important factor to opt out of convenient services. As such, there must be legal requirements on cloud providers to ensure sufficiently stringent default privacy policies for their services, which must be backed up by mechanisms to ensure compliance as well as means to ensure enforcement of these laws.

Cloud computing provides many opportunities and many risks, and the full implications of the technology are hard to predict. It is important, however, that consideration for the importance of privacy be given at the birth of this technology, rather than allowing them to be neglected until the technology matures. In this way, the benefits of cloud computing will be felt not only by businesses and entrepreneurs, but also by the people with whom they interact.

Refernces

[1] Jan Camenisch and Els Van Herreweghen. Design and implementation of the idemix anonymous credential system. In *CCS '02: Proceedings of the 9th ACM conference on Computer and communications security*, pages 21–30, New York, NY, USA, 2002. ACM.

[2] David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 4(2), February 1981.

[3] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th USENIX Security Symposium*, August 2004.

[4] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178, 2009.

[5] Oded Goldreich. Secure multi-party computation. Working Draft, 2000.

[6] Philippe Golle. Revisiting the uniqueness of simple demographics in the us population. In *WPES '06: Proceedings of the 5th ACM workshop on Privacy in electronic society*, pages 77–80, New York, NY, USA, 2006. ACM.

[7] Amazon Inc. *Amazon Elastic Compute Cloud (Amazon EC2).* Amazon Inc., http://aws.amazon.com/ec2/, 2008.

[8] Peter Mell and Tim Grance. The nist definition of cloud computing. Technical report, National Institute of Standards and Technology, Information Technology Laboratory, July 2009.

[9] Arvind Narayanan and Vitaly Shmatikov. How to break anonymity of the netflix prize dataset. *CoRR*, abs/cs/0610105, 2006.

[10] United Nations. *Universal declaration of human rights*. Australian National Committee for United Nations, Melbourne :, 1949.

[11] Daniel Nurmi, Rich Wolski, Chris Grzegorczyk, Graziano Obertelli, Sunil Soman, Lamia Youseff, and Dmitrii Zagorodnov. The eucalyptus open-source cloud-computing system. In *Proceedings of Cloud Computing and Its Applications*, October 2008.

[12] Rafail Ostrovsky and William E. Skeith III. A survey of single-database private information retrieval: Techniques and applications. In *Public Key Cryptography*, pages 393–411, 2007.

[13] Andreas Pfitzmann and Marit Hansen. A terminology for talking about privacy by data minimization: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management, December 2009. v0.32.

[14] Gebhard M. Rehm. Just Judicial Activism? Privacy and Informational Self-Determination in U.S. and German Constitutional Law. *SSRN eLibrary*, 2000.

[15] E. Eugene Schultz. A futuristic look at cloud computing security. Forthcoming. (2010 Information Security Summit, Prague.), 2010.

[16] Staff. Eu assesses adequacy of us safe harbor privacy compliance. *IWAYS*, 28(1):27–33, 2005.

[17] M. Stringer, G. Fitzpatrick, D. Chalmers, E. Harris, R. Krishna, and M. Haarlander. Kuckuck exploring ways of sensing and displaying energy consumption information in the home. In *Proceedings of Workshop on Ubiquitous Sustainability: Technologies for Green Values (at UbiComp 2007)*, 2007.

[18] Samuel D. Warren and Louis D. Brandeis. The right to privacy. *Harward Law Review*, 4(5):193–220, December 1890.