The (social) construction of information security

Wolter Pieters

DIES/IS group, Zilverling Building, Faculty of Electrical Engineering, Mathematics and Computer Science, University of Twente, P.O. Box 217, 7500 AE ENSCHEDE, The Netherlands, w.pieters@utwente.nl, www.cs.utwente.nl/~pietersw

Abstract

To solve societal problems of access to information, the scientific discipline of information security is vital. However, its implicit philosophy of physical containment is not practical anymore due to increased connectivity. To enable authorities to overcome their struggle with requirements and regulation, this paper provides a revision of these foundations, based on the system-theoretic notion of causal insulation, in comparison with existing technical approaches. From this perspective, it discusses differences and similarities between physical, digital and social protection mechanisms for information, defines basic concepts, and lays the foundations for a better understanding of the role of information security in society.

Keywords: causal insulation, constructionism, information security, security perimeters, system theory

1 Introduction

A major part of computer science research is now devoted to what is called information security. In this subdiscipline, the challenge is how to protect information systems against malicious users. This is quite different from research on e.g. programming paradigms or software engineering methods, primarily because security is concerned with what systems should *not* do, rather than what they should do. The philosophy of this research has not received much attention until now. Although security-related societal implications of information systems – especially in the area of privacy, see e.g. Nissenbaum (1998); Floridi (2005); Gutwirth and De Hert (2008) – have been discussed extensively, the fundamental definitions of the scientific endeavour were left largely untouched. This not only implies a philosophical problem. It also entails the impossibility of connecting the high-level privacy discussions to the technical possibilities that have been developed. Therefore, authorities struggle with regulation for private initiatives such as social networking services, as well as the security and privacy requirements of their own IT projects.

The main issue here is that the notion of privacy, in its informational meaning, is too narrow to describe the technical implementation of its own policies. Computer scientists speak of privacy, but they mean with it a special kind of information security, namely *confidentiality* of *personal* information. The repertoire of information security is much broader, and covers integrity and availability of information next to confidentiality, and business, military and government information next to personal information. For the computer scientist it does not matter what kind of information needs to be secured. For the policy maker, it does. Therefore, the technical solutions will never speak of privacy as it is used at policy level, and policy makers will never speak of information security as it is used in the technical domain.

Still, there seems to be throughout society an implicit philosophy[1] of what information security is based on the notion of containment, taken from physical analogies such as buildings and safes. In such a philosophy, the asset to be protected needs to be separated from the environment by one more or less homogenous security boundary, such as a fence, or a firewall. An analogy is often drawn to a fortress, with thick walls on the outside but weak on the inside. In the present text, I

[1] In this paper, I use the term "philosophy" to refer to an understanding of the foundations of a scientific discipline. This does not necessarily mean a systematic account, as such an understanding is often implicit and unarticulated.

argue that this implicit philosophy is unsatisfactory in the current age of increased connectivity, and provide an alternative foundation. I do so from a constructionist point of view, where the co-evolution of social and technical mechanisms is seen as the source of the security of an information system, rather than rational design choices only. The concept of causal insulation from system theory is employed in order to give an account of the fundamental characteristics of information security research. This generates definitions that can be used in discussing information security from a philosophical perspective, as well as in analysing security policies. Two themes are central in the analysis. First, information security is not merely a design problem, as external forces shape the threats to and protection of information systems. Secondly, people play a central role in the security of information systems, both in the role of attackers and in the role of defenders. Hence the title of this paper, and hence the claim that the vocabulary presented here enables discussing the (social) construction of information security.

In the following section, I describe in more detail what information security research involves, and why its implicit philosophy is inadequate. In section 3, I interpret this research in terms of the system theory of Niklas Luhmann, with the concept of causal insulation as a central theme. In section 4, the analysis is validated by showing how this interpretation matches existing (technical) approaches in modelling information security. Section 5 analyses the role of policies in the causal insulation approach. In section 6, I discuss from the system-theoretic perspective similarities and differences between information protection in the physical, digital and social domain. Section 7 provides new definitions that can guide future research on this topic. Based on these definitions, I focus on the mechanisms of (social) construction in section 8. Finally, in section 9, the analysis is applied to the example of electronic voting, and I discuss how the new views may inspire practices in information security modelling.

2 Information security

Information security aims at providing tools and mechanisms for protecting the confidentiality, integrity and availability of information in the face of attacks. Confidentiality protects against unauthorised reading, integrity against unauthorised writing, and availability against unauthorised deletion of information. These properties are associated with *risks* of the information systems involved: something may go wrong if the system is not designed properly. The term security denotes that there are enemies. Safety or correctness, by contrast, deals with such risks under “normal” circumstances, i.e. without an active adversary. For example, a safety property of a computer system is that it does not crash spontaneously; a security property may be that it is resistant to so-called denial-of-service attacks.

At first sight, information security seems to rely on a distinction between what needs to be protected and its environment. Confidential information should not get “out”, and unauthorised information should not get “in”. Following this intuition, the implementation of information security policies has often been based on a so-called security perimeter. An example is a firewall, a single device filtering all incoming and outgoing network traffic of an organisation, blocking potentially dangerous messages. The notion of perimeter makes an explicit distinction between inside and outside. What is inside is trusted, what is outside is not. Outside threats should not be allowed to reach the inside, whether it concerns confidentiality or integrity of information.

This implicit philosophy seems to originate in an analogy with physical protection by means of e.g. safes and access control in buildings. In this form of protection, physical boundaries are created in which the assets are *contained*. The containing perimeter has a limited number of gates (such as doors), which also limit the traffic that can go through (using for example keys). When trying to protect information, it seems natural to interpret this kind of protection in similar terms. Consequently, the design of information systems has followed a similar pattern, and the associated concept of containment is often used in modelling information security (Scott 2004; Nunes Leal Franqueira et al. 2009). Also, the term *exposure* is used to describe what part of the “inside” is accessible from the “outside” (Dragovic and Crowcroft 2004).

This focus on containment, as expressed in the idea of perimeter-based security, has recently become controversial. First of all, the problem of insider threat, where persons inside the perimeter misuse their capabilities to disrupt the system, poses a challenge (Probst et al. 2007). Insiders are trusted by definition, and mechanisms to protect against insider threats may therefore be absent. Moreover, increasing demand for access to the organisation’s assets from outside the organisation’s physical boundaries, e.g. via virtual private networks (VPNs) and employee notebooks, has challenged the notion of a perimeter, as company networks now have to be accessible from outside the premises. The outsourcing of services to other companies is also a major drive for external access.

In information security modelling, we see this problem when multiple connections between entities need to be modelled. In the containment philosophy, the model of the connections is supposed to be a tree, where there is only a single path from one entity to another. When, in the "real" world, multiple paths exist, counterintuitive constructions need to be added to account for these features. For example, a building is then modelled as a tree, in which computers are contained in rooms, but additional connections between nodes of the tree are added to model wireless networks (see e.g. Nunes Leal Franqueira et al. (2009)). Why, then, is the basic model still conceived as a tree, and why is the philosophy one of containment? We might be better off with a different starting point.

According to the Jericho Forum (2005), protection of information should no longer be based on a single perimeter separating the organisation from its environment. In what is called de-perimeterisation, the boundaries of the information infrastructures of organisations dissolve. Where previously a firewall was used to separate the untrusted outside from the trusted inside, outsourcing of information management and mobility of employees make it impossible to rely on such a clearly located security perimeter. Nowadays, we hear increasingly about "cloud computing", where it becomes completely invisible to the user where the information is stored and processed, for example in Google Docs. It is argued that in such an environment, protection should therefore lie as close to the data as possible, i.e. "data level security".

The question has been raised whether this is really a paradigm shift, or just a relocation of the perimeter; whether it is de-perimeterisation or re-perimeterisation. After all, it is still necessary to protect the data; only the size of the trusted inside could be said to be reduced. Protection may no longer be based on the physical separation of networks through a firewall, but rather on digital separation of the data by means of encryption (e.g. sticky policies, Karjoth et al. (2003)). The relocation argument has a limited scope, though. Whereas the containment philosophy may still work for the encryption itself, the complex connections that allow access to the encrypted data cannot be modelled from such a perspective. Several people, possibly working in different organisations, will have access to the information, possibly based on different credentials and through different routes. In such a situation, the question is which concepts can describe the aims of different physical and digital protection mechanisms – and thereby the aims of the scientific information security community – appropriately.

3 Causal insulation

In order to answer this question, we may build upon existing research in philosophy of technology. In defining information security in the age of increased connectivity, we need to develop a theory that allows for dynamic and heterogeneous rather than fixed and homogeneous boundaries between what we wish to protect and the threats that endanger these assets. The inside then consists of things that work together, and the outside consists of things that work against the inside.

An analogy can be drawn here with the body, seen as protective mechanism of the genes. Where an intuitive perimeter seems to appear in the form of the skin, there are obviously protective mechanisms that operate within that perimeter (e.g. the immune system) as well as outside of that perimeter (changing the environment to offer better protection, e.g. in the form of building shelters). This is what Richard Dawkins calls "the extended phenotype" (Dawkins 1989), and we may speak similarly of "the extended security perimeter". Unlike the perimeter from the common sense philosophy, an extended perimeter is neither static nor homogeneous.

We should thus replace the common sense notion of boundary with something theoretically more sophisticated. In this paper, I use the distinction between a system and its environment, which forms the basics of system theory. We may also define the inside as an actor-network with a particular program of action, and the outside as an antiprogram. This actor-network theory based perspective is dealt with elsewhere (Pieters 2011b); in this paper, I focus on the system-theoretic point of view.

One of the most important researchers in 20th century system theory was the German sociologist Niklas Luhmann. In particular, his book *Risk* (Luhmann 1993) deals with matters of protection and security. Since we are interested in securing information technology, the chapter on technology is of particular interest. According to Luhmann, "what is called technology, is a *functional simplification in the medium of causality*" (p. 87). Although quite abstract at first sight, Luhmann's explanation of this definition provides the insight that "[t]he result of technicalization is thus the more or less successful insulation of causal relations [...]" (pp. 87–88). Creating technology is getting intended causes in and keeping unwanted causes out. "The form of technology [...] marks the boundary between enclosed and excluded (but just as real) causalities."

From this perspective, designing technology involves decisions that on the one hand specify which causes are allowed to pass in and out, and on the other hand which causes are not allowed to pass. The latter are the safety and security properties of the technology. What does this mean for information security? I take what is often called a socio-technical view here, in which physical, digital and social elements can be part of the technology (system) under investigation. From this point of view, we can address information security of computer systems as well as information security in organisations.

In traditional security in organisations, we had a physical perimeter separating the inside of the organisation from the outside. Digital data had to pass this physical perimeter in order to move into or out of the organisation. The Jericho approach is interpreted to stand for data-level security, where the physical perimeter is replaced by security of the data itself, by means of cryptographic techniques. From the causal insulation point of view, both are different mechanisms to achieve a causal insulation of the data from the environment.

In both cases, the confidential data inside the organisation is not supposed to cause changes in the environment of the organisation: if it would, then the environment could be using the confidential data for some purpose. Conversely, the organisation wishes to protect its sensitive information from outside influence; because the data is important, outsiders should not have control over what the information tells the organisation. Thus, there are different ways in which we can implement the causal insulation for the socio-technical system under investigation. Contrary to the perimeter perspective, these mechanisms need not be static or homogeneous.

We should keep in mind that Luhmann is primarily speaking of the *safety* of technology, that is, the keeping out of *unintended* external causes. When we move to security of technology, and thus face adversaries, we wish to keep out *intended* causes, i.e. malicious acts of an attacker. This by itself leads to philosophical considerations, but these have been discussed elsewhere (Pieters 2010). Here it suffices to say that the enemies are determined to *make happen* those causes that match their intentions. If we focus on information, we need to include *only* this type of causes. The leakage of information is only dangerous if enemies will make use of it. That is, the leakage of information *by itself* will not be harmful, but only when it is used by an agent.[2] This means that in information security, the causal insulation of information always assumes adversaries. By contrast, causal insulation of other

[2] In privacy research, there is a similar discussion between privacy as opacity and privacy as transparency, where in the latter, the *use* of private information is regulated (Gutwirth and De Hert, 2008).

technologies may include harmful effects to e.g. health or the environment that do not need human mediation in order to occur. Information, by itself, does not have such effects.

This analysis can be compared to Luciano Floridi's work on "ontological friction" (Floridi 2005), which also deals with a type of "resistance" that exists in what he calls the "infosphere": "the environment constituted by the totality of information entities – including all agents – processes, their proprieties and mutual relations" (Floridi 1999). This perspective provides a useful abstraction for understanding how information technology changes the flow of information, in particular in relation to ethical questions about privacy. However, the system-theoretic perspective of causal insulation and perimeters has a number of advantages. First, it does not rely on the acceptability of claims on ontological changes that information technology induces in the infosphere, and rather provides a pragmatic modelling perspective in terms of systems. Second, it thereby emphasises the possibilities for achieving insulation in design, including a multi-level view on extended security perimeters, where causal insulation can even run through agents, as we will see later..

The basic understanding of technology by means of causal insulation thus provides us with a new way of considering perimeters: they are not necessarily about physical boundaries, but about limiting the possibilities of information influencing other information. Physical boundaries are a *specific type* of causal insulations. In the past, physical boundaries were a good way of causal insulation, but it is precisely the process of de-perimeterisation that challenges this success. Later, we will see what kinds of boundaries are characteristic of the new situation.

4 Non-interference

To validate the definition of information security in terms of causal insulation, a comparison can be made to existing computer science research. In information security, a particular view on the protection of information takes the perspective of information flow. The question then becomes which information can influence other information.

Based on this research, it can be argued that a notion of causal insulation has already been developed that is specific to *information*. This notion has been called *non-interference* (Sabelfeld and Myers 2003). From the perspective of information flow, non-interference means that high-security information cannot flow to low-security environments (confidentiality), or that low-security information cannot flow to high-security environments (integrity). For example, privacy-sensitive information cannot end up on a publicly accessible web page. Or, conversely, information that was entered on a website by an unknown user cannot end up in a critical file.

One of the possible definitions of confidentiality from the perspective of non-interference is found in Jacobs et al. (2005). In this definition, the basic assumption is that, if a partition of the world is not influenced by information from outside this partition within a given period, then the final state of the partition should be independent from outside causes. That is, if in two different states of the world the projection of the state on the partition is the same, the projection of the resulting states of the world on the partition should still be the same after the indicated period. This holds for integrity of the partition. For confidentiality, the situation is the other way around. Then, the resulting state of the world *outside* the partition should be independent from what happens inside the partition, so that information from the partition cannot be leaked.

In this approach, the focus is on computer programs, and the world consists of a computer memory. This memory is partitioned according to security levels. The notion of non-interference thus provides an informational point of view on causal insulation. If a partition of the memory is properly protected, this means that information cannot pass its "boundary" without conforming to its policy. Such policies may in practice be enforced by encryption: only with the right credentials one can access the information.

Thus, the perspective of causal insulation corresponds to information flow analyses in information security research. In particular, such methods analyse the situation where there is no *physical* boundary between pieces of information, and we still wish to keep them separate in terms of influence. In the above example, the analysis focused on information flow *within* a computer program. However, apart from the complications of moving from a formal to a natural domain, there is no reason why the idea could not be applied to a broader setting of information security, where flows between physical systems and people can be included. This, however, is not the aim of the present analysis.

With respect to our goal, providing a philosophical foundation for information security, the comparison to information flow shows that the analysis of information security in terms of causal insulation is a valid one in principle. It also shows that causal insulation for information security means a *specific* kind of causal insulation, namely one between information items. As such, the causal insulation aimed for is causal insulation in the area of information and meaning (infosphere in Floridi's terms), rather than in the spatial-physical world.

5 Policies

Requirements for causal insulation of information can be described in terms of *policies*. A policy denotes under which conditions causes can pass the causal insulation. Policies are ascribed to the world by agents, and the only function agents have is ascribing policies. I thus do *not* see the access relation of one object to another as an inherent agent-to-object relation. Rather, these are relations between information objects (where an information object can be a human), and agents ascribe policies to these access relations (where an agent can again be a human). Policies for granting access can be represented in terms of the access that an entity already has to other (information) objects. If the actor then wants to be granted access, she needs to either conform to the policy or have the policy changed. For example, a door may be entered by using a key (conforming to the policy) or by breaking the lock or the door (changing the policy). In a digital setting, one may guess a password (conforming to the policy), or change the access rights of the file one wants (changing the policy).

From this perspective, it does not matter how causal insulation of information is implemented, since it only concerns the (dis)connection between different pieces of information. The physical layout of a building is only an implementation of a particular information access policy, and is only relevant *as* implementation of this policy. Therefore, it is not relevant whether in the physical world room 2 is adjacent to room 1 and only reachable through room 1; it is only relevant that there is a policy stating that access to room 2 is limited to entities already having access to room 1, which is implemented with a certain strength, and can be modified by entities capable of interacting with the policy (the room might be tempted to change its policy in interaction with dynamite).

This analysis of the role of policies can move our attention from the physical analogy of containment to a more general foundation of information security. Still, the intricacies of the different possible forms of *implementation* of such policies deserve a more detailed analysis. This will be our focus for the next section.

6 Physical, social and digital protection

We have seen that, in order for technologies to function, they need to "decide" which causes they let in or out. This is what Luhmann calls causal insulation. Causal insulation properties for information can be specified in terms of policies, in which it is specified which access is needed to gain more access. Intuitively, causal insulation in the infosphere may be realised by physical, digital, or social mechanisms, depending on the type of agents involved. We may build a wall, separate information flows, or tell people not to give away their passwords. How do these different types of mechanisms fit into the causal insulation perspective?

First of all, we can distinguish between passive and active causal insulation. In passive insulation, the insulation is implicitly realised by "common" physical properties. In active insulation, a special mechanism is included in the design that is supposed to take care of the protection. A piece of paper is in principle not accessible, unless you have the paper in your hands (the so-called "air gap"). A file on the Internet is in principle accessible, unless it is actively protected (e.g. by encryption).

As an example, consider the difference between barcodes and RFID (radio-frequency identification) chips on consumer products. The information in the former cannot easily be captured from a distance, since the products mostly reside inside shopping carts and bags. By contrast, the information in RFID chips can be read, unless there are protective measures in place. This makes the security of the RFID information dependent on the adequacy of the security protection mechanism. Such differences also apply when boundaries fade with de-perimeterisation and converging technologies: there is a shift from passive causal insulation to active causal insulation due to increased connectivity.

Active protection, in contrast to passive protection, is by definition based on design decisions. This means that, in Luhmann's terminology, the possibility of failure is always one of risk instead of danger: one could have made a different design decision, which is not the case with passive protection by physical separation of technologies. Moreover, how the protection works can no longer be understood without specialist knowledge. It is easier to convince the public that barcodes cannot be read from a distance than to achieve the same result for RFID, even when experts find the protection adequate. This means that trust becomes increasingly important. Instead of unconsciously relying on the physical separation of systems, we have to decide consciously whether we trust a security measure to protect our assets.

Simultaneously, increased connectivity often amounts to a shift from causal insulation based on physical separation to causal insulation based on informational separation (non-interference). Whereas a traditional pill relies on chemical properties to release its contents, subject to local causes only, a digital pill may be steered from outside the body. This requires again active protection, which is typically based on informational properties rather than physical properties (e.g. authentication and encryption).

When insulation is insufficient, as is often the case when connectivity increases, an alternative or complementary approach is to detect when a technology is being misused. In information technology, this is called intrusion detection (Bolzoni and Etalle 2008). When everything is connected in the information domain ("Internet of things"), lack of protection may lead to, for example, digital pills being "hacked", even when we *think* that adequate protection is in place. In such a case, pills need to be suspicious about the instructions given to them: if they get a strange sequence of instructions, they may decide not to execute them and generate a warning instead. Moreover, this security mechanism will itself rely on information about the use of the device, which also needs to be protected. We could decide to call this *causal exile*, which is complementary to causal insulation.

In the case of the physical perimeter in an organisational context, the causal insulation is achieved by separating the causal mechanisms inside and outside the organisation. This separation is physically represented by, for example, a firewall, which is the only connection between the network of the organisation and the outside and untrusted Internet. Other sources of data flowing into or out of the organisation should be controlled in a similar way, e.g. by disabling USB ports and other ways for employees to take away or insert data. However, what the employees *know* is still moving outside the organisation. Employees have to work with the data, making it necessary to give them the information in such a way that they can do so, i.e. unencrypted. Since people cannot be asked to give up their private life, they inevitably operate in both trusted and untrusted environments, and are therefore "part of the security perimeter". Next to physical and digital protection, the social factor is thus crucial in protecting the information of the organisation.

Many researchers have investigated this social side of information security. Where both the physical and digital parts of the perimeter can be controlled by technology, causal insulation of data that is present in people in the form of knowledge cannot be protected in such a way. Here, the causal insulation is achieved by training and law. An important question is whether we can represent social separation in a similar way.

In digital and physical protection, the protection mechanism has to decide whether or not it will let certain causes in or out. This is usually based on something else that the “cause” has access to, such as a key or a password. As said before, one can then gain access either by conforming to the policy, or by changing the policy. Does this also work in social settings?

The answer seems to be yes. Again, there are basically two ways for an actor to convince someone else to give her something she should not be given, for example a password. The first is to present some credential that according to the other’s policy gives her the right to have the password. The second is to make her opponent change his policy, such that the request and the policy are compatible. This is not so different from the methods to gain access to a building or an IT system. In terms of causal insulation, the first method is to change the environment to conform to the causal insulation while still reaching the goal, and the second is to change the system’s causal insulation.

It may be argued that the notion of roles makes the social domain fundamentally different from the physical and digital domain. However, roles can be modelled in terms of policies and credentials. If I wish to impersonate an employee of an organisation, I can either obtain a credential such as an employee card, or make someone change his policy in order to grant me access without such a card. In both cases, I may be said to have successfully impersonated an employee.

Still, it may be objected that in the second case, the impersonation is based on trust rather than credentials, which would then be something specific to the social domain. Again, I would reply that trust is a matter of what one would or would not do in an interaction with a person. If I trust you, I am more likely to delegate an important task (and the necessary credentials) to you. But we can also reverse the definition: if I am more likely to delegate goals or authorisations to you, then I can be said to trust you more. Trust is then defined as intention to delegate.[3]

The most important difference between the social domain and the physical and digital domains seems to be that the implementation of policies is not deterministic. A door will always, or with very high probability, let someone in who has the key, and keep someone out

[3] For more about definitions of trust, see Nickel (2012) and Pieters (2006).

who does not have the key. By contrast, a person may act differently in different circumstances, and she may only conform to the policy, say, 60 % of the times. Whether this is a matter of free will or of circumstances is not something to be addressed here. Even if people's behaviour may be expressed by deterministic-but-very-complicated policies, depending on many circumstances, for all practical purposes the behaviour will need to be understood probabilistically.

In all cases – whether it concerns physical, digital or social implementations – changing the policies should be difficult, as it can be a very powerful way to get any type of access to a system. Thus, this subsystem should have its own causal insulation, which is usually stricter than the overall one. Still, system administrators often have a lot of power, making the insulation dependent on their goodwill alone.

I conclude that, although some aspects are different, physical, digital as well as social aspects of information security can be modelled in terms of causal insulation. In all cases, the causal insulation is realised by means of access policies. Causal insulation can – and should – be complemented by what I called "causal exile", i.e. intrusion detection. To bypass causal insulation, one either needs to conform to the policy or have the policy changed. Changing policies may again require special causal insulation, to prevent giving too much power to administrators.

7 Containment revisited

Based on the analysis in the previous sections, I argue that information security is best modelled by the possible interactions between information entities, based on the causal insulation between them. In such a model, the primary question is what can access what, and how this may change over time.

When we wish to investigate security, we can abstract from the mechanism that implements causal insulation, and focus instead on the level of resistance that a certain mechanism gives to unwanted causes trying to break the insulation. In such a model, each entity has a policy of keeping in, keeping out, letting in and letting out. This policy is enforced with a certain strength. Whether the policy is actually enforced depends both on the value of the asset to be protected and the force that the environment can apply to break in.

Existing approaches often focus on containment as the fundamental security relation. However, this seems to lead to arbitrary choices for the direction of the relation. For example, does a firewall "contain" a network? The choice to represent one network as "inside" and the other as “outside”, as in Nunes Leal Franqueira et al. (2009), will depend on the location of the assets, but cannot be meaningfully deduced from the structure of the world only. If the asset were on the other side of the firewall, the containment would be reversed. The representation of the structure of the world is then dependent on the value assigned to the entities. It seems that, rather than being a fundamental property, containment is *derived* from what is being protected against what. Intuitively, we may use entities and connections between entities to model these relations. Entities can then access each other if they are connected.

Definition 1 *a is* informationally contained *in b to the extent that its connection with b can prevent events in the world from causing informational changes in a (integrity), and/or can prevent a from causing informational changes in the environment (confidentiality). a is* completely contained *in b if a can exchange information with the environment only through its connection with b.*

For example, a computer may be (partly) contained in a room. If it furthermore has a wireless network connection, it is also (partly) contained in the wireless network. If the computer is stand-alone, it is fully contained in the room.[4]

Definition 2 *An* informational perimeter *of a is a set of entities that together can prevent events in the world from causing informational changes in a (integrity), and/or can prevent a from causing informational changes in the environment (confidentiality)*

Note that if {*b*} is a perimeter of *a*, then *a* is fully contained in *b*. In the previous example, a room plus a wireless network may form a perimeter of a computer. This composite system may have its own perimeter again, say, in the form of a building plus a firewall. The building may have a perimeter in terms of the people who can go in and out (taking information with them).

These definitions show us that the notions of containment and perimeter are still relevant, *but not a priori*. Instead, containment and perimeters are *derived* concepts, and they are derived from a model of the world in which all possible interactions between information items are incorporated. Since this model concerns the infosphere, spatial or physical arrangements are only relevant to the extent in which they represent causal insulation in the infosphere. Such a philosophy is more

[4] Obviously, these examples depend on the chosen level of abstraction in the model of the world.

suitable in the current age of complex informational networks, since it does not limit the acceptable types of causal insulation on forehand.

In many cases, it is not sufficient that causal insulation is in place, in the sense that it inhibits information flow. Often, it must be assessable by parties involved that this insulation is indeed in place, i.e. the information *about* the insulation must not be insulated. I call this *observable insulation*. (In Floridi's terms, we may speak of "visible friction".) Such visibility depends on the capabilities of the observer. Typically, physical insulation is more visible than digital or social insulation, as human observers are better equipped for / trained in physical observation. Why a ballot box constitutes insulation is so trivial that explanation is often unnecessary.[5] Therefore, forcing the information flow through the physical world is often a way to improve observable insulation.[6]

8 Double contingency

Not only are humans, and thereby social aspects, part of the security perimeter. For all that information adds to the complexity, such analyses still apply to safety issues as well. What is different for security, is that *attackers are also part of the perimeter*. When attackers decide not to attack, they are effectively contributing to the security of the system: they reduce the probability that the desired system properties fail. Moreover, what defenders do and say influences the attacker decisions, which again influences what defenders do. As both attackers and defenders are aware of the contingency of the other's actions, and therefore find themselves in a situation of *double contingency* (Luhmann 1995).

This situation has interesting self-reinforcing properties for the perception of security of both attackers and defenders. When attackers attack a system in a specific way, the focus of both the attacker and the defender community is drawn to the specific problem that is exploited, leading on the one hand to more attacks and on the other hand to better

[5] The role of the concept of explanation is dealt with elsewhere (Pieters, 2011a).

[6] The proposal of a Dutch committee on the future of the voting process was exactly this: people can vote electronically, but the ballot must be forced through the physical world (i.e. printed) (Election Process Advisory Commission, 2007). In the US, the notion of a Voter Verified Paper Audit Trail (VVPAT, Mercuri (2002)) does not actually force the information flow through the physical world, but creates a physical backup for detection of problems. The physically separate devices used in online banking systems in the Netherlands are another example. Here, the codes for access and signing have to be manually entered, so that digital threats such as viruses cannot seize power over them, *and* it can be observed that this is the case.

defences. Both of these can again reinforce the attention that is being paid, and therefore reduce or improve security, depending on whether the attackers' or the defenders' efforts are more successful. In any case, an arms race is constituted about the specific problem, and similar problems are likely to appear in the near future, as attackers will try variations of the same trick, before the defenders think of said variations.

This also means that in security – and this is a key claim in this article – that *probabilities of attack are dependent on security perception*. In the words of those who like to distinguish between actual and perceived security, *actual security is dependent on perceived security*. Therefore, what is often called actual security is necessarily socially constructed, or, rather, constructed in a socio-technical constitution of artifacts and humans. We cannot speak of the security of an electronic voting machine by itself; the probabilities of attack depend on what is perceived about its security, and are therefore context-dependent. A report about vulnerabilities in the machine not only changes security perception, it also changes the probability of attack, and therefore the actual security of the device.

This is not to say that technical models or measurements of a machine's security are meaningless. The point is that if security is understood as probability of damage (or probability times damage), then these technical methods do not measure security. They measure security as if the security perimeter *is* the device, which is not true in any practical situation.[7]

9 Example: electronic voting

To illustrate how this new philosophy of information security would work in a practical situation, and to show how it can contribute to political discussions and policy on information security, I discuss the example of electronic voting.

Traditionally, security in the voting process in an election relied on two types of containment. One was the voting booth, in which a voter could cast her vote without pressures from the outside world (e.g. vote buying or coercion). The other was the ballot box, assuring that only legitimate ballots would end up in the count. This arrangement seems to support the idea of security as containment. However, the voting booth and the ballot box are by themselves not sufficient to safeguard

[7] This would also mean that the notions of threat, vulnerability and impact, used in security risk assessment, would have to be redefined in terms of causal insulation. I leave this for future work.

the properties they seem to provide. For example, voters leave fingerprints on their ballots, in principle allowing others to assess which vote is theirs. Such “electoral traces” (Pieters 2009) may break the secrecy of the ballot. Also, ballot boxes may not be empty at the start of an election, allowing so-called “ballot stuffing”.

Additional procedural measures are therefore part of the perimeter. These include the public nature of counting and the destruction of the ballots (making it impossible to take fingerprints from them), and the checking of the integrity of the ballot box before the start of the election. Here, the security perimeter runs through the people who observe the procedural measures: they decide whether undesirable informational causes can pass through. The adequacy of such measures heavily depends on whether attackers are actually interested in, say, taking fingerprints from ballots. Therefore, they also form part of the perimeter.

In electronic voting, the situation is different. Most voting machines do, for example, have a feature to assure that the count is zero at the beginning of the election. However, it is impossible for the poll workers to verify that this procedure is adequately implemented in the software. Therefore, the perimeter will now include the people and the places involved in programming the machine. If the machine can be re-programmed, also the storage facilities are places where unintended informational causes may intervene. Again, potential attackers within the organisations involved are part of the perimeter. If they see benefits in manipulating the software, they can cause damage. If they are not, they actually protect the system information-wise. Here, even ethical codes or moral values can be containers of information.

In Internet voting, the perimeter is extended even further. The integrity of the individual vote is then often dependent on the integrity of the computer the voter uses to cast her vote. As we know, many personal computers are infected by viruses and spyware.

It seems to be a general tendency that with trends of automation, virtualisation, convergence of technologies, cloud computing, and many more, an extension occurs of the security perimeter. This means that in the new version of the associated procedures, more people and places become involved in the setup of the procedure, and thereby also more people and places become part of the security perimeter, aimed at safeguarding the process against unwanted interference. Moreover, such new versions often offer fewer possibilities for intrusion detection, because they lack the necessary transparency. For intrusion detection (causal exile), in the sense of being able to find out if parts of the perimeter fail, openness is needed, whereas closure is often seen as needed for security (causal insulation), especially in connection with commercial interests of companies. Thus, when companies are part of the security perimeter, they may provide security, but this cannot be verified, and neither can it be observed (from the outside) when incidents happen and need to be responded to.

The debate on openness versus obscurity still runs within the information security community, and will continue to do so for the foreseeable future, precisely because of these two conflicting requirements (cf. Federspiel and Brincker (2010)). A general direction to look for solutions is data classification. By providing transparency for unclassified data (e.g. system design and encryption algorithms) and secrecy for classified data (e.g. encryption keys), a combination of openness and closure may be achieved. However, business interests often make it impossible to provide the required openness. The renewed definitions of containment and security perimeters at least make it possible to cast new light on this debate, and continue it in a more informed way.

10 Conclusions

In this paper, I analysed the philosophical foundations of the scientific discipline of information security. I argued that information security can be interpreted and explained in terms of causal insulation, based on Luhmann's system theory. I showed that this interpretation is consistent with existing research paradigms in information security. Based on this analysis, I discussed the relation between physical, digital and social aspects of information security, and provided definitions for fundamental concepts in the area. The definitions provided are more flexible than they would be in a philosophy of (physical) containment. In particular, they allow for security perimeters partly running through the *social* world, both in the cultures of defenders and of attackers, which is essential for understanding the social origins of information security, and its transformation by new technologies and new ways of organising businesses and society.

By connecting the technical and policy discourses on information security and privacy, this analysis can form the basis for a better understanding of their relations in current and future developments. This does not only hold for electronic voting, as shown in the example, but also for public transport payment systems, road pricing, electronic patient records, and many more. In all of these cases, technical perimeters as such are overrun by the many connections needed, but perimeters in terms of causal insulation, running through computers, organisations, buildings and people, can provide the necessary understanding of how security is constructed, and in the end enable better judgements on what is more secure than what.

This is not to say that the analysis is complete, or without challenges. In future work, I aim at comparing the system-theoretic approach to a second analysis in terms of Latour's actor-network theory (Latour 2005; Pieters 2011b). I expect this analysis to strengthen the arguments for moving away from a containment-based philosophy of information security to a "flat" ontology consisting of different actors that connect or disconnect from each other. However, the comparison between the two may also reveal possible weaknesses in both of them, and contribute to further improving the conceptual framework. Then, it could be operationalised for decision support in policy contexts. I would also like to address the question how information security contributes to realising the moral laws in information ethics (Floridi 1999; Ess 2009), as well as how ethics itself can improve our security perimeters. For if people constitute (part of) the boundary in information security, their own policies are fundamental to improving our ways of dealing with the infosphere.

Acknowledgements

This research is supported by the research program Sentinels (`www.sentinels.nl`). Sentinels is being financed by Technology Foundation STW, the Netherlands Organization for Scientific Research (NWO), and the Dutch Ministry of Economic Affairs. The author wishes to thank (in alphabetical order) André van Cleeff, Johnny Søraker and Roel Wieringa for interesting brainstorm sessions and useful comments. Part of this paper is based on an extended abstract presented at the 2009 conference of the Society for Philosophy and Technology (SPT).

References

Bolzoni, D., and S. Etalle 2008. Approaches in anomaly-based network intrusion detection systems. In *Intrusion Detection Systems*, pp. 1-15. Advances in Information Security 38. Berlin: Springer.

Dawkins, R. 1989. *The extended phenotype*. Oxford: Oxford University Press.

Dragovic, B., and J. Crowcroft 2004. Information exposure control through data manipulation for ubiquitous computing. In *NSPW'04:*

Proceedings of the 2004 workshop on New security paradigms, pp. 57-64. New York, NY: ACM.

Election Process Advisory Commission 2007. Voting with confidence. http://www.kiesraad.nl/nl/Overige_Content/Bestanden/pdf_thema/Pdf_voor_Engelse_site/Voting_with_confidence.pdf (accessed June 17, 2011).

Ess, C. 2009. Floridi's philosophy of information and information ethics: Current perspectives, future directions. *The Information Society* 25:159-168.

Federspiel, S. B., and B. Brincker 2010. Software as risk: Introduction of open standards in the Danish public sector. *The Information Society* 26:38-47.

Floridi, L. 1999. Information ethics: on the philosophical foundation of computer ethics. *Ethics and Information Technology* 1:37-56.

Floridi, L. 2005. The ontological interpretation of informational privacy. *Ethics and Information Technology* 7:185-200.

Gutwirth, S., and P. De Hert 2008. Regulating Profiling in a Democratic Constitutional State. In *Profiling the European Citizen: Cross-disciplinary Perspectives*, eds. M. Hildebrandt and S. Gutwirth, pp. 271-302. Berlin: Springer.

Jacobs, B., W. Pieters, and M. Warnier 2005. Statically checking confidentiality via dynamic labels. In *WITS '05: Proceedings of the 2005 workshop on Issues in the theory of security*, pp. 50-56. New York, NY: ACM.

Jericho Forum 2005. Jericho whitepaper. https://www.opengroup.org/jericho/Business_Case_for_DP_v1.0.pdf （accessed June 17, 2011).

Karjoth, G., M. Schunter, and M. Waidner 2003. The platform for enterprise privacy practices: privacy-enabled management of customer data. In *2nd Workshop on Privacy Enhancing Technologies (PET 2002)*, pp. 69-84. Lecture Notes in Computer Science 2482. Berlin: Springer.

Latour, B. 2005. *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford: Oxford University Press.

Luhmann, N. 1995. *Social Systems*. Stanford, CA: Stanford University Press.

Luhmann, N. 2005 [1993]. *Risk: a sociological theory*. New Brunswick: Transaction Publishers.

Mercuri, R. 2002. A better ballot box? *IEEE Spectrum* 39:26-50.

Nickel, P. J. 2012. Trust in Technological Systems. In: *Norms and the Artificial: Moral and Non-Moral Norms in Technology*, eds. M. de Vries, S. Hansson, and A. Meijers. Berlin: Springer. Forthcoming.

Nissenbaum, H. 1998. Protecting privacy in an information age: The problem of privacy in public. *Law and Philosophy* 17:559–596.

Nunes Leal Franqueira, V., R. H. C. Lopes, and P. van Eck 2009. Multi-step attack modelling and simulation (MsAMS) framework based on mobile ambients. In *Proceedings of the 24th Annual ACM Symposium on Applied Computing, SAC'2009, Honolulu, Hawaii, USA*, pp. 66-73. New York, NY: ACM.

Pieters, W. 2006. Acceptance of voting technology: between confidence and trust. In *Proceedings of iTrust 2006*, eds. K. Stølen et al., pp. 283-297. Lecture Notes in Computer Science 3986. Berlin: Springer.

Pieters, W. 2009. Combatting electoral traces: the Dutch tempest discussion and beyond. In *E-Voting and Identity: Second International Conference, VOTE-ID 2009*, eds. P. Ryan and B. Schoenmakers, pp. 172-190. Lecture Notes in Computer Science 5767. Berlin: Springer.

Pieters, W. 2010. Reve{a,i}ling the risks: a phenomenology of information security. *Techné: Research in Philosophy and Technology* 14:176-188.

Pieters, W. 2011a. Explanation and trust: what to tell the user in security and AI. *Ethics and information technology* 13:53-64.

Pieters, W. (2011b). Representing humans in system security models: An actor-network approach. *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications* 2:75-92.

Probst, C. W., R. R. Hansen, and F. Nielson 2007. Where can an insider attack? In *Workshop on formal aspects in security and trust (FAST2006)*, pp. 127-142. Lecture Notes in Computer Science 4691. Berlin: Springer.

Sabelfeld, A., and A. C. Myers 2003. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications* 21:5-19.

Scott, D. J. 2004. Abstracting Application-Level Security Policy for Ubiquitous Computing. PhD diss., University of Cambridge.