

Computer security impaired by legal users.

Denis Besnard & Budi Arief

School of Computing Science
University of Newcastle upon Tyne
Newcastle upon Tyne NE1 7RU
United Kingdom

denis.besnard@ncl.ac.uk
l.b.arief@ncl.ac.uk

Abstract. Computer security has traditionally been assessed from a technical point of view. In this paper, we wish to adopt a cognitive standpoint and investigate some of the cognitive processes involved in computer security. One angle which is not considered very often is the active role played by legal users of systems in impairing the level of protection. In this paper, we thus attempt to highlight the cognitive processes underlying security impairments by legal users. This approach relies on the concept of trade-off. At the end of the paper, we propose a short usability-centered set of recommendations.

Keywords. Computer security, cognitive psychology, trade-offs.

1. INTRODUCTION

Our society is becoming more and more dependent on computer systems, which nowadays are used in everyday life, from business to banking, from entertainment to healthcare. Most of these systems are interconnected through the internet, which inherently is very open and vulnerable to cyber-attacks. Attacks on these systems cause a lot of inconvenience at the very least, and in some cases may lead to financial or safety consequences. Therefore, securing computer systems has become a very important part of system design, development and deployment.

Security has been very often addressed from the attackers' side. From this angle, the emphasis has classically been on the means used to break into systems. However productive this research area has been and still is, it tends to blur the exact role of the legal users¹ who are also actively involved in computer security. It seems to the authors that this other angle is worth exploring as well. Moreover, as part of the DIRC² research, we are interested in interdisciplinary aspects of computing. For these reasons, some aspects of the role of legal users regarding security are addressed in this paper.

Computer security is an area which has not been extensively investigated by cognitive scientists. It nonetheless offers an interesting aspect in the sense that there are conflicting objectives held by some of the actors of a single system, namely attackers and legal users. It follows that depending on the goal that a user is pursuing (attack or legal use), the use of a given computer system will differ dramatically. Whereas the role of attackers are pretty clear, those of a legal user regarding security are more subtle. Stemming from this assumption, we will examine some of the legal user's practices and shed some light on the mental processes involved. We will try to assess

¹ In this paper, legal users will refer to a large number of actors, including end-users, security officers, managers and designers. This list is not meant to be exhaustive. We only wish to build a simple category of actors who are concerned about security but are not attackers.

² DIRC (Dependability: an Interdisciplinary Research Collaboration) is a UK-based interdisciplinary research project on the dependability of computer-based systems. Visit DIRC at <http://www.dirc.org.uk>.

the extent to which computer security can be interpreted in terms of a balance between a number of factors on legal users' side. The approach we adopt here relies on trade-offs.

2. TRADE-OFFS IN THE WILD

Since Simon (1957) and his concept of bounded rationality, it is accepted that human actions do not reach perfection but instead seek an acceptable level of performance with respect to their goals and what the cognitive resources allow. The fact that the cognitive system never aims at handling all the data available in the environment is a central aspect of the cognitive resources saving strategy. As a consequence, cognitive acts are a trade-off between cost and some sort of satisfaction. An individual does not have an infinite amount of time or effort to allocate to a given goal. Instead, he or she seeks an answer which balances cost and efficiency in the best way. This strategy is put in place for the majority of human actions. Interestingly enough, the drawback with trade-offs in security is that they often introduce the opportunity for a threat. Before we consider computer security, it can be useful to have a look at a field example. Although the latter is quite remote from computing, we think it puts things clearly and shows how the parameters of a trade-off are manipulated by humans.

On December 30, 1999, in Tokaimura (Japan), a criticality accident³ occurred at the JCO nuclear fuel processing plant, causing the death of two workers (see Furuta *et al.*, 2000). The immediate cause of the accident was the pouring of approximately 15kg of uranium into a precipitation tank, a procedure requiring mass and volume control. The workers' task was to process seven batches of uranium in order to produce a uranium solution. The tank required to process this solution is called a buffer column. At JCO, its dimensions were 17.5 cm in diameter and 2.2 m in depth, owing to criticality safe geometry⁴. The inside of this tank was known to be difficult to cleanse. In addition, it was positioned only 10 cm above the floor, making it difficult to collect the uranium solution from the bottom of the column. Thus, workers illegally opted for using another (larger) tank called precipitation tank. Due to its dimensions, this latter tank was not geometrically safe but it was positioned 1 m above the floor. Moreover, it was equipped with a stir propeller making it easier to use for homogenising the uranium solution. The pouring of the seven batches at once triggered the criticality accident. Its causes were rooted in a complex combination of deviant organisational practices. Among these, pressures from the managerial team to increase the production without enough regard to safety implications and lack of crew training played a significant role.

In hindsight, we speculate that the operators have traded-off productivity and practicality against risk. As their knowledge about critical uranium masses was poor, they were unaware that they were crossing a safety boundary. This case is an instance of how trade-offs can go wrong. With this example, we want to highlight the workarounds that operators often implement in order to perform daily actions in a less constrained manner (see Gasser, 1986). These workarounds can be put in place in a wild way, and depending on the level of knowledge and perceived risk, getting the work done sometimes overrides security concerns.

Following a cognitive approach, we believe that virtually every decision is a matter of trade-off. Humans do not try to produce perfect responses to the environment. Instead, they tend to accept good enough solutions. We think that this conception of human cognitive activities applies in computer security, for both attackers and legal users. The former attempt to design effective worms or denials of service, for instance. The latter, in turn, try to protect themselves as effectively as possible. But in both cases, there are not infinite amounts of resources (e.g. time,

³ There is a limited amount of uranium that can be put together without initiating fission. When this amount is exceeded, a chain reaction occurs, generating potentially lethal radiations.

⁴ Generally speaking, narrow tanks prevent unwanted fission reactions.

money or effort) to allocate to attacking or protecting. This is where human flexibility comes into play: people perform an intuitive trade-off between (some form of) cost and (some form of) benefits. We will consider some concrete computer examples after having briefly explained how the concept of trade-off translates in security.

3. SECURITY FROM A COGNITIVE PERSPECTIVE

To better illustrate how we use the concept of trade-off, we represent it graphically in Figure 1. The dark area at the lower right-hand corner represents the maximum efficiency where one reaches high benefits for low costs. The top left-hand corner, on the contrary, represents a poor efficiency where one spends a lot to gain little. Between these two extremities, there is clearly an entire continuum. The trade-off line represents a frontier between costs and benefits. Any activity above this line will cost more than it rewards. Conversely, any activity below this line will reward more than it costs.




Figure 1: Graphical representation of a costs/ benefits trade-off.

Let us now apply the graph to the simple example of a user wanting to configure his or her computer. Let us assume that the user has a routine task to do and is considering adding a button to a toolbar in order to access a function more quickly. If the button is going to be used only one time, the time cost of adding it may be higher than the expected time saved during the task. Therefore, it is likely that the button will not be added. On the other hand, if this function is to be used repeatedly, the expected benefits may be worth the cost. We are obviously not considering dimensions such as adding buttons for the sake of it or for a matter of investigating unknown functions. Nevertheless, with this simple example, we claim that humans intuitively, though implicitly, evaluate the efficiency of their decisions before they implement them.

There are cases where human actions are given explicit limits. Decisions can be benefit-driven or costs-driven (see Figure 2). In the first case, a course of actions is interrupted when some functional objective is met. In the other case, the target is set in terms of cost (money, time, etc.) and actions will stop when the limit is reached. In both cases, the course of actions never follows the straight trade-off line. Instead, we think it fluctuates with time depending on the given phases of the work (e.g. software development will hardly show a steady progress rate).




Figure 2: Graphical representations of benefits-driven (left) and costs-driven (right) strategies.

Because they constantly try to save cognitive resources while still attempting to reach some level of efficiency, humans often implement a *least effort rule*. In doing so, they attempt to reach an acceptable level of performance with the minimal mental effort. As a consequence, decision making can become a biased benefits-driven process. When applied to a security-usability trade-off, usability may come first, hence turning security into a side-issue. This will be discussed in section 5.2.

Classically, attackers are said to exploit security holes left open by legal users. In other words, the malicious intentions of the attackers are, to some extent, facilitated by the behaviour of some legal users. We obviously do not put the blame on them. Neither do we believe that the motives that some attackers promote (e.g. learning, curiosity, challenge, etc.) will ever justify any sort of damage caused to someone’s data, tool or service. Having said that, security is a two-way issue. Merely assessing it from the attacker’s point of view only captures half of the problem (see Arief & Besnard, 2003 for a survey on attackers’ techniques and motivations). The other half is about how we (legal users) use our computers.

4. SECURITY TRADE-OFFS BY LEGAL USERS

In our view, legal users consider the actions required for securing the system from an economic point of view and they trade-off security against usability. The increasing number and length of passwords, the tedious installation and updating of anti-virus software or the straightforward nature of opening email attachments are threats that are introduced in systems and that attackers exploit. It does not imply that they can be easily removed but understanding where trade-offs lie allow security engineers and software designers to think more about the interest of making security products and policies compatible with some intuitive notion of usability. This is not only an issue about “comfort of use” but, as we will see in the following sub-sections, is a problem that directly impacts security.

4.1.1. Passwords: a memory issue

Passwords can be eight characters long or more and because we are making computing more and more ubiquitous, we seem to need an increasing number of them. This tendency probably

originates from a desire to control accesses more tightly with the hidden assumption that it will increase privacy of data. It may also derive from an erroneous belief that long passwords are more difficult to crack. When one actually looks at what happens at the workplace, human cognitive limitations become obvious: users cannot remember their passwords and need external memories (e.g. sticky notes on monitors). This is an instance of a nice paradox where by increasing the complexity and number of passwords, the level of privacy actually decreases (Weirich & Sasse, 2002).

User login and passwords to computer accounts are used very often and are the main method to get access to systems. In this case, frequency compensates for complexity: the password is used often enough to be remembered (Sasse, Brostoff & Weirich, 2001). But how many counter-examples are there where people have to write down passwords? To cope with this problem, operating systems offer to remember passwords. Again, for the sake of usability, a user may be tempted to use such storage features (cookies). It is a useful feature but it comes down to the user's judgement as to whether a service or a piece of data is trivial enough so that its password can be stored on a computer. Another drawback is that users who rely on cookies may not spend effort in remembering their passwords anymore. As a consequence, if somehow the cookies are emptied (e.g. when the system is rebuilt), they may lose their password.

4.1.2. Anti-virus software updates: a risk issue

Anti-virus protections are useful barriers but only when they are up-to-date. They need some attention in this respect. But maintaining, updating and upgrading them has a cost that can conflict with end-user's main task, thereby impacting security. So we think anti-virus protections always leave a hole open, the size of which depending, among others, on the frequency of updates. Hence risk, which is typically perceived inaccurately by humans, comes into the equation at this level. Users or security officers have to accept a certain level of (perceived) vulnerability regarding their system. Automatic updates have been felt to tackle this problem by lowering the likelihood of holes in the anti-virus protection. Unfortunately, the corruption of this service puts high threats on IT systems since it can offer an attacker to automate the installation of backdoors or the downloading of harmful contents. For this reason, automatic updates are not a panacea. Despite this negative state of facts, large organisations use this feature extensively, implicitly assuming that the gain in usability and the regularity of updates together will provide benefits that are worth taking the risk.

4.1.3. Email attachments: a trust issue

Email attachments have been used very widely for spreading malicious code. Typically, the code is added as an attachment to a seemingly benign email. When the attachment is opened, the code is executed on the machine, exploiting security holes in the email program. Harmful email attachments pose two problems. First, they very often come from trusted third parties who were infected themselves. Thus, due to this trust relationship, the degree of suspicion regarding the decision to open the document is already low. Second, email attachments are used so widely for legitimate purposes that opening them has become as automatic as picking up the phone. Of course, a rule such as "*do not open attachments you are not expecting or from people you do not know*" is unworkable since it generates too many false positives: many valid attachments would have to be left unread. Again, automation has been felt to tackle this problem: many mail servers now include anti-virus scanners. However, this is not the optimum solution since these scanners a) sometimes detect false positives and b) open a discussion on whether or not scanning might invade privacy.

5. DISCUSSION

After having considered some concrete examples based on the use of computers, it seems necessary to adopt a broader view and address some more general issues. Among these, accepted losses and risks have to be mentioned as security policies are not meant to protect every single piece of data or service. Furthermore, threats can be discussed from the standpoint of an antagonism between the roles of attackers and legal users. The discrepancy in their motives is where security holes lie. Lastly, we will address organisational issues by describing a multi-layered notion of systems' security.

5.1. Accepted losses, risk perception and systems' protection

Although we have reasoned so far under the reductive assumption that all the data have to be protected, we now want to highlight a somewhat different picture according to which there exist some acceptable losses that humans implicitly take into account in setting up protections. The trade-off here involves the cost of protection and the cost of loss (See Flechais & Sasse, **year**). Data or services that can be easily replaced, disclosed or lost without serious consequences will probably have a relatively low level of protection. The underlying evaluation of the required level of protection is believed to be done in an intuitive manner most of the time. We also think it guides, to some extent, the security policies adopted by organisations. As each single piece of data cannot be equally protected, some of these data are inevitably left vulnerable to attacks. This may be a sensible decision if, as aforementioned, loss or disclosure of data is accepted. In this view, it seems important to highlight that protecting a system is a sort of dialogue between security officers and attackers. We think security policies define the nature of this dialogue before the occurrence of any attack. To some extent, this conception goes against the widespread belief according to which "attackers play first".

IT security shares some similarities with safety that help highlighting the implicit side of risk taking. Let us take the following example. A car is safer when it is immobile than when one is driving it. But for a car to deliver its service (transporting people and goods), the driver and the passengers are forced to expose themselves to risks. These may be reduced down to some acceptable level if the driver is careful and experienced. But there will always be a number of factors he or she will not be able to influence (other drivers, mechanical incidents, etc.) that will impact the level of safety of this situation. The same argument holds with, for instance, a server. Not plugging it into a network is a secure condition but the service will not be delivered⁵. Therefore, some risk has to be accepted for virtually any piece of equipment to fulfil its function.

Another problematic situation is one where a threat has been identified by e.g. a software developer but found too costly to fix or too unlikely to care about. This applies to an extremely wide range of cases. For example, in everyday's life, we tend not to wear our safety belt when driving on very short distances (e.g. parking the car into the garage). This behaviour has a simple explanation. We, as individuals, adapt our protection level to the environment as a function of the perceived risk. The smaller the perceived risk, the lower the level of protection. However, this intuitive and implicit risk analysis is problematic. As humans are typically biased at perceiving actual levels of risk and never have an absolutely exhaustive knowledge of the systems they interact with, it follows that the impact of a given practice over the security of a system is unlikely to be accurately assessed by a user. It implies that these heuristic risk assessments do not accurately capture the criticality of certain threats, therefore leaving breaches identification and compensation subject to subjective decisions.

⁵ Actually, even disconnecting a machine does not solve all the problems. Social engineering techniques allow the attacker to get in through other means, including by getting physical access to the target system.

Last but not least, there is a human tendency to “slide on the risk slope”. Large security incidents or industrial accidents are not caused by a sudden change in security or safety policy. Departure from a reasonable level of risk does not happen in one day. It is an accumulation of a number of small insecure increments that progressively deteriorate the level of protection, each of which being seen as acceptable *per se*. This is classic in large industrial system’s safety: large-scale accidents are made of a concatenation of small failures (Mancini, 1987).

5.2. Antagonism among security actors

It seems plausible that attackers, just as legal users do, perform trade-offs in the way they use their own computers. They may tend to intuitively and implicitly compare the costs of their actions to the expected benefits and then take decisions on the basis of this evaluation. The rule-of-thumb states that if costs are perceived as worth the expected benefits, then some action is likely to be performed. However, because attacker’s and legal users’ motivations are fundamentally different, we think that their respective trade-offs are different in nature. Attackers attack because they like it and/or get a reward of some sort (self-satisfaction, peer-recognition, money, etc.). Legal users protect themselves because they need to. As far trade-offs are concerned, these motives bring a consequence that attackers may care less about costs than legal users do. This discrepancy of motivations may create some room for an increased likelihood and/or success of attacks.

As depicted in Figure 3, legal users may prioritise usability with little concern about security. This is a common case of usability-driven behaviour. On the attackers’ side, it may be that one is focused on the damages expected from his work with little concern about the costs involved. We have represented this discrepancy as a dotted line in the legal users’ trade-off strategy graph in Figure 3. Thus, in our conception, the gap lying between attackers’ and legal users’ trade-off strategies gives some advantage to attackers. According to this view, the larger the gap, the most successful the attack could be.




Figure 3: Graphical representations of two different trade-off strategies.

From our point of view, a successful attack can therefore be described in terms of a malicious action whose degree of refinement is higher than the degree of protection of the target system.

5.3. Beyond the individual picture...

So far, we have been concentrating on an individual perspective where cognitive factors are thought to play a determinant role. Beyond this picture, we want to acknowledge the collective

dimension of security in large distributed computer-based systems. More precisely, we think that Reason's model (1990; see Figure 4) adopts a useful view on organisations in the sense that they are described in terms of multi-layered systems. Applied to the field of security, this view can help describe a computer-based system as one composed of threats, actors and protection layers. With this model, security is described as a multi-layered process where a variety of actors (e.g. developers, security officers, end-users...) have a role to play. Each of these actors impact on security. Legal end-users, for instance, by not making updates for their anti-virus software, leave holes open for attacks. This type of sub-optimal behaviour may exist at any given layer of the organisation, for any role. It creates latent security breaches that, combined with each other, can defeat an entire system's protections. In the context of this paper, these breaches are interpreted in terms of trade-offs whereby actors simply wish to reach good enough solutions. When applied to security, Reason's model can describe, from a system point of view, the impairments made to security by legal users and attackers. The latter attempt to propagate attacks through security holes in order to reach an objective such as data, a service or causing disruption in the functioning of the system.

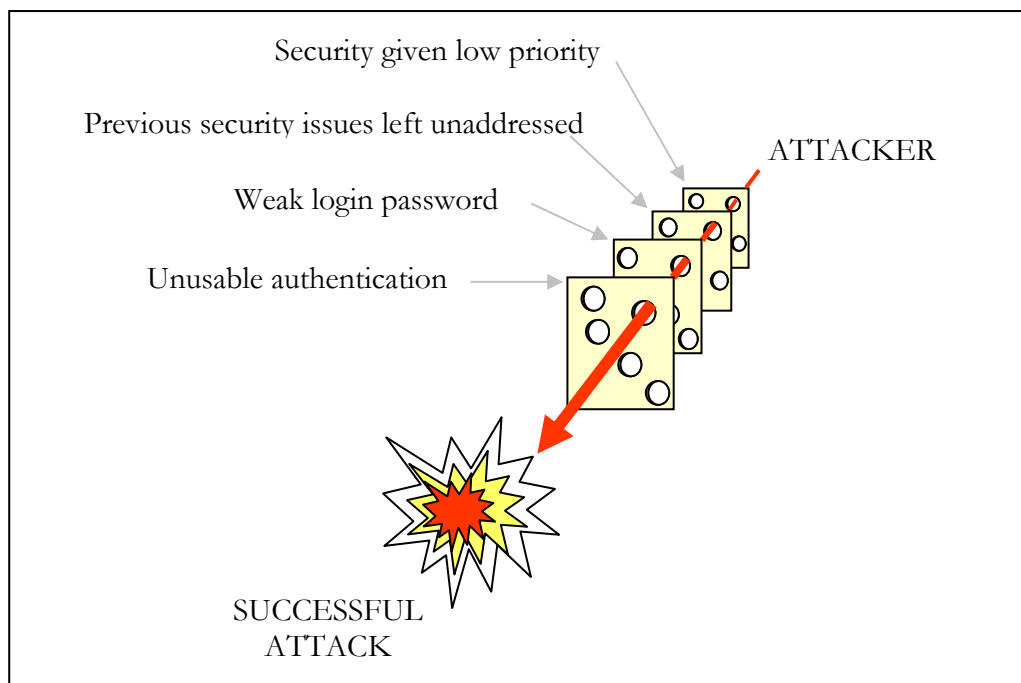


Figure 4: Successful attacks propagate through several protection layers (adapted from Brostoff & Sasse, 2001).

By quoting this model, our intention is to highlight the combination of factors needed for an attack to succeed. According to the popular belief, attacks occur because some malicious people exploit security holes. We wish to promote a somewhat different view according to which successful attacks are a combination of weak protections and malicious intentions. This idea can be stretched even further. As we cannot eliminate attacks altogether, the most productive approach may be to regard attacks as the outcomes of flawed policies and/or practices. Their origins are deeply rooted within early design decisions or within organisations. For instance, due to productivity constraints, the manager of a small company may misjudge the importance of protecting IT activities. This can take the form of a backlog of security actions waiting to be done. Such a *laissez-faire* policy may propagate through the various stakeholders of the organisational process, leading to e.g. unprotected data or weak passwords. There could obviously be an infinite number of examples that would follow the same pattern of a multiple,

intricate set of causes (see Brostoff & Sasse, 2001 for more complete views on Reason's model application to security).

Although an individual approach is useful to assess individual contributions or impairments to security, these have to be brought back into a broader picture where security breaches are caused, facilitated or maintained by the combination of a variety of causes rather than by mere, isolated end-users' actions.

6. WHAT CAN WE DO?

Telling people what to do about security is one option. But one lesson that can be drawn from violations in systems is that one should not expect humans to always act as prescribed. Within industrial settings, procedures themselves do not rule the human behaviour (Fujita, 2000) and there are many ways in which humans can configure a system and use it in unexpected and/or unprotected modes, even if it implies implementing a violation (Adams & Sasse, 1999). The motivation for doing so may be based on an intuitive cost/benefit evaluation. This is typical for passwords that are written down or passed on to colleagues. It is also true for harmful email attachments that happen to hit computer-aware staff in academic departments every now and then. Generally speaking, if the perceived risk attached to an illegal action (e.g. lending a password) is seen as lower than the expected benefits (e.g. gain in time), then a violation will be put in place. This is extremely common practice and goes well beyond computer security. In this trade-off, factors such as security culture and risk perception are key notions. And whether or not the user has a relevant knowledge of the potential consequences of his/her actions is what partly determines the level of risk involved and the final security of the system. Thus educating users, although it will not solve all the problems, gives them the opportunity to better understand the consequences of their actions, hence making them less prone or susceptible to security attacks.

6.1. Summary

If it is true that cognition is a matter of trade-offs, we then have to modify the balance of the factors that compose them. As far as legal users are concerned, one may want to act so as to increase security without impairing usability. We think that adopting a trade-off approach allows one to look at security as a balance. Identifying which are the factors involved and which are their respective weights in the decisions taken is hoped to constitute a progress towards better protections.

Here are the points the we have defended in this paper:

- Trade-offs are sometimes implemented in a wild, uncontrolled manner. Legal users sometimes prioritise immediate benefits to the detriment of long-term security.
- Passwords, anti-virus updates and email attachments respectively raise such issues as memory limitations, risk and trust.
- Security does not imply protecting everything since some losses are acceptable. However, because risk perception by humans is highly biased, valuable data could be under threats.
- Trade-offs by legal users differ in nature from the ones performed by attackers. The resulting gap creates or maintains security breaches.
- Computer security is an organisational matter.

6.2. Recommendations

To put things simply, humans obey least-effort rules because they are cognitive machines that attempt to cheaply reach flexible objectives rather than to act perfectly towards fixed targets. As a consequence, each time an opportunity to do so arises, cognitive resources are saved. This rule is sometimes applied even to the detriment of performance or accuracy. From our point of view,

anything that can push legal users into trading off security against anything else is harmful. In our opinion, this leads to some simple recommendations.

- *Educate staff.* As we stated earlier, education will not solve all the problems but will at least allow users to be aware of the consequences of their actions. Tell staff how some email attachments may contain harmful code. Explain how intrusions are performed. Highlight the importance of memorising passwords.
- *Security must be user-centred* (Zurko & Simon, 1996). Following the previous point, passwords must be, at least, easy to remember and reduced in number as much as possible. As far as end-users are concerned, the ideal number of passwords is zero. It may seem an unworkable view to security officers but the reason why security policies have to be enforced to humans is because they require an effort from them. To this respect, any measure getting closer to an *effortless* security policy is a step forward.
- *Security is not end-users' task.* How secure a system is partly depends on how high security is set on the scale of objectives. If security is a relatively obvious goal for a security officer, it is not the case for an end-user. Solutions have to be thought of in order to make security transparent for whom it is not a primary objective.
- *Do not set contradictory objectives.* Asking staff to carry out their duty and spend time on updates and/or scanning files cannot be done at the same time and have to be traded-off against each other. Do not impose contradictory objectives. Security does not come first in end-users' mind (See Sasse, Brostoff & Weirich, 2001).

As a summary of these short recommendations, end-users will always have something else to do than think about security. It seems to us that the idea of a *user-centred security* for end-users is a useful policy driver. Any measure going in this direction will improve systems' security.

7. LIMITS

One of the issues this paper does not address is how cognitive flexibility eventually contributes to security. In previous sections, we have seen how trade-offs between e.g. usability and security could impair a system's level of protection. This surely accounts for the success of some intrusions. But this aspect of human functioning can be seen in a more positive way as not complying to the rules can also generate beneficial behaviours (Besnard & Greathead, 2002). According to this view, violations are reconsidered under the angle of ad-hoc contributions to security, happening under exceptional circumstances and outside the frame of any clearly identified procedure. Several examples can be mentioned. One is unplugging the network cable from a connected machine when a suspicious behaviour is noticed. It may not preclude any damage on this specific computer but it will prevent the attack to spread to other machines. This is the type of actions that designers probably do not expect users to take but that can nonetheless be implemented on-the-fly, thanks to human's intrinsic flexibility. This kind of unexpected contribution to security is hardly ever addressed in computer security but nevertheless deserves some attention.

8. CONCLUSION

One conclusion we can reach in this paper is that when legal users impair the level of protection of a system, they do not always do it because of a pure lack of knowledge or because they explicitly disregard security. It simply is that legal users who are not security-aware (e.g. researchers, clerical staff, managers) have other tasks to perform than spending their work time on securing their system. If asked to do so, they might consider security as getting in the way to the completion of their main tasks and will probably not bother scanning data for viruses if this allows them to work faster. Harmful usability-driven trade-offs are then put in place and this creates holes in systems' protections. Interestingly enough, the same argument applies to

attackers but in a slightly different way. They care about the time and efforts they allocate to a given attack with regards to the expected benefits. Thus, when attacks or intrusions fail, it is not always because attackers are incompetent. Instead, they may be faced with e.g. technical problems during an intrusion that would be too long to surmount given their level of competence, available time, or expected reward.

Understanding where trade-offs lie can allow a better understanding of the mental processes involved in security practices. In the case of legal users, we speculate that security is impaired because it is traded-off against usability or efficiency.

9. ACKNOWLEDGEMENTS

This paper was written at the University of Newcastle upon Tyne within the DIRC project (<http://www.dirc.org.uk>) on dependability of computer-based systems. The authors wish to thank Peter Ryan and Jeremy Bryans as well as anonymous reviewers for useful comments. The authors are also grateful to the sponsor EPSRC for funding this research.

10. REFERENCES

- Adams, A. & Sasse, M. A. Users are not the enemy. *Communications of the ACM*, 42, pp. 41-46 (1999).
- Arief, B. & Besnard, D. Technical and Human Issues in Computer-Based Systems Security. *Technical Report CS-TR-790*, School of Computing Science, University of Newcastle upon Tyne, UK (2003).
- Besnard, D. & Greathead, D. A cognitive approach to safe violations. *Technical Report CS-TR-791*, University of Newcastle, UK.
- Brostoff, S. & Sasse, M. A. Safe and sound: a safety-critical approach to security. Proceedings of the *New Security Paradigms Workshop*, Cloudcroft, NM, pp. 41-50 (2001)
- Flechais, I. & Sasse, M. A. Developing secure and usable software. To be presented at *OT2003* (March 30th-April 2nd 2003).
- Fujita, Y. Actualities need to be captured. *Cognition, Technology & Work*, 2, pp. 212-214 (2000).
- Furuta, K., Sasou, K., Kubota, R., Ujita, H., Shuto, Y. & Yagi, E. Human factor analysis of JCO criticality accident. *Cognition, Technology & Work*, 2, pp. 182-203 (2000).
- Gasser, L. The integration of computing and routine work. *ACM Transactions on Office Information systems*, 4, pp. 205-225 (1986)
- Mancini, G. Commentary: Models of the decision maker in unforeseen accidents. *International Journal of Man-Machine Studies*, 27, pp. 631-639 (1987).
- Reason, J. *Human error*. Cambridge University Press (1990).
- Sasse, M. A., Brostoff, S. & Weirich, D. Transforming the weakest link – a human computer interaction approach to usable effective security. *BT Technological Journal*, 19, pp. 122-131 (2001).
- Simon, H. A.. *Models of man*. New York, Wiley (1957).
- Weirich, D. & Sasse, M. A. Pretty good persuasion: A first step towards effective password security in the real world. Proceedings of the *New Security Paradigms Workshop*, Cloudcroft, NM, pp. 137-144 (2002).
- Zurko, M. E. & Simon, R. T. User-centred security. Proceedings of the workshop on *New Security Paradigms*, Lake Arrowhead, CA, pp. 27-33 (1996).