Analysis of affordance, time and adaptation in the assessment of industrial control system cybersecurity risk

J.S. Busby, Department of Management Science, Lancaster University, UK, LA1 4YX;
j.s.busby@lancaster.ac.uk (corresponding author)
B.Green, School of Computing and Communications, Lancaster University, UK, LA1 4YX;
b.green@lancaster.ac.uk
D.Hutchison, School of Computing and Communications, Lancaster University, UK, LA1 4YX;
d.hutchison@lancaster.ac.uk

Abstract

Industrial control systems increasingly use standard communication protocols and are increasingly connected to public networks – creating substantial cyber-security risks, especially when used in critical infrastructures such as electricity and water distribution systems. Methods of assessing risk in such systems have recognized for some time the way in which the strategies of potential adversaries and risk managers interact in defining the risk to which such systems are exposed. But it is also important to consider the adaptations of the systems' operators and other legitimate users to risk controls, adaptations which often appear to undermine these controls, or shift the risk from one part of a system to another. Unlike the case with adversarial risk analysis, the adaptations of system users are typically orthogonal to the objective of minimizing or maximizing risk in the system. We argue that this need to analyze potential adaptations to risk controls is true for risk problems more generally, and we develop a framework for incorporating such adaptations into an assessment process. The method is based on the principle of affordances, and we show how this can be incorporated in an iterative procedure based on raising the minimum period of risk materialization above some threshold. We apply the method in a case study of a small European utility provider, and discuss the observations arising from this.

KEYWORDS

Adaptation, affordance, industrial control systems, cybersecurity

1 INTRODUCTION

It has become widely recognized that industrial control systems are inherently vulnerable to the risk of cybersecurity failure as they become extensively connected to the public network. Modern societies rely heavily on such systems for their basic utilities and production processes[1], much of which constitutes 'critical national infrastructure'[2]. Yet 'decades of competition have led to infrastructure that is very lean and very fragile'[3]. The vulnerability was vividly illustrated by the Stuxnet and Flame events in 2010 and 2012 [4], a recent attack on the control systems of a German steelworks [5], and even more recently the attack on parts of the Ukrainian electricity distribution system[6] [7].

In some ways networked industrial control systems have an inherent resilience because an attacker intending a specific harm not only has to achieve privileged access to physical controllers, or the data coming from physical sensors, but needs specific engineering knowledge to achieve physical damage. However, in some ways they have an inherent vulnerability because they offer a virtual, remote pathway to physical damage that does not require physical proximity, physical intrusion or physical action. An attacker who can shut down a SCADA (Supervisory Control And Data Acquisition) server can cause considerable, non-specific disruption. And when such systems control basic infrastructures, such as water and electricity distribution, there is a large-scale risk at a societal level.

Assessing risk of this kind has much in common with assessing risk from deliberate, adversarial action in other types of system. Those working in this area have long recognized the need to analyse how defensive actions influence adversarial actions, and have used decision theoretic or game theoretic approaches to represent this adaptation. In this paper, however, we argue that there has been very little recognition of how the system at risk – in particular legitimate operators and organizations – also adapts to security measures and risk controls over time, sometimes nullifying those controls, and sometimes displacing rather than mitigating risk. We suggest that it is important to anticipate such adaptations, and to use some systematic theory to anticipate when and how adaptation occurs. We base our approach on the principle of affordances[8] [9] from ecological psychology, which appears especially suited to anticipating potential adaptations because of its emphasis on the fit

between an entity and its environment in terms of the action possibilities this fit creates. We show how the resulting framework can be applied in a case study of small utility serving a town in Europe. In the remainder of this paper we briefly review the relevant literature, describe the proposed method of analysis, show how it was applied, and discuss its contributions and limitations.

2 A BRIEF REVIEW OF THE LITERATURE

2.1 Risk and industrial control systems

Past work on analyzing risk in industrial control systems has tackled a wide variety of issues. It has developed methods for assessing threats based on general systems theory[10] [11] [12], assessing risks arising specifically from automated and robotic systems[13] and assessing vulnerability and survivability in complex industrial controls[14]. It has also produced more specific methods of risk assessment – for example using fuzzy logic for verifying safety criteria in nuclear power plants[15], applying hybridized probabilistic and AI-based approaches to real-time risk controls in systems such as gas compressors on offshore production platforms[16], and applying simulation approaches to the complex dynamics in chemical batch reactors[17].

Methods have been developed for analysing dependency in access control systems[18], the vulnerability of inter-dependent SCADA systems to cascade failures[19], reliability in the control room operations of process control systems[20], and risk in digital controls[21]. Some of our understanding is about the particular risks associated with specific industries – such as water distribution systems[22] [23]. But other work has been concerned with techniques to deal with the general problems of analysing risk in control systems, for example the use of hierarchical Bayesian approaches to address the sparsity of empirical data about rare failures[24], and petri net approaches to deal with state space explosion problems in analysing the safety of digital control systems[25].

Cyber security risks to industrial control systems, specifically, have received increasing attention, especially when they form part of what is regarded as critical infrastructure. For example, Henry and Haimes[26] develop an extensive model of security risk in process control networks. Ralston

et al[27] provide a broad view of concerns and risk assessment methodologies for SCADA and Distributed Control Systems (DCS). Kundur et al[28] present a framework for assessing the impact of cyberattacks on smart grids, synthesising analysis of both the information and physical systems. Naedele[29] presents a review of commonly discussed SCADA security challenges, including the newly interconnected nature of control systems, and the use of commercial off-the-shelf products. Fleury et al[30] develop a taxonomy of attacks against energy control systems. And various studies use a range of different formalisms for risk and reliability assessments. For example Zhang et al[31] use Bayesian attack graph models to analyze attacks on SCADA elements in an electricity distribution system, while Ten et al[32] use petri nets for a similar purpose. Other instructive applications of similar methods include Henry et al's[33] [34] use of petri nets to represent the coupled development of attacks on a computer network and failure of physical control, and Ingols et al's[35] use of attack graphs to model both attacks and counter-measures, and to automate the process of risk assessment.

However we argue below that there is a problem which has not yet received much attention in this literature: the way in which a complex sociotechnical system inevitably adapts in some way to any attempt to protect it – and the need to recognize this in the risk assessment and management process. Henry and Haimes[26] (2009) took the step of building into their framework the analysis of candidate risk controls. We take the further step of building in the analysis of adaptations of the system to those controls. As we discuss in detail later, this involves anticipating how such controls provoke legitimate system operators and users to change their behavior. This then requires iterating through the risk assessment and control identification process, recognizing such adaptive behavior, and continuing until risks are judged acceptable. This inherent iteration is not intended to reduce uncertainty in some general sense, as in certain other approaches[36], but to find the state in which a system exhibits acceptable risk even after it has adapted to newly-applied risk controls.

2.2 Risk of deliberate harm

The general problem of how to assess risk of harm that is inflicted deliberately, by some adversary, has received considerable attention. There are surveys and comparative views of the different methods – both decision theoretic and game theoretic, and combinations of these – that can be used for risk

assessment with intelligent adversaries[37] [38]. There are assessment frameworks dealing with specific problems, such as the identification of critical assets[39], and incomplete information about early moves in adversarial games[40]. There are also empirical models of terrorism risk estimated from regional characteristics[41], and of cyber security incidents and their effects on internet traffic[42].

Terrorism risk has attracted particular interest. Special journal issues show a range of concerns and approaches, including the effects of terrorist attacks on utilities[43], the use of historical data on disturbances to analyze utility risk[44], and methods for assessing the vulnerability of critical infrastructures[45]. Recent work provides assessment methods for terrorist threats to airports specifically[46], and methods that recognize the interaction of counter-measures in a portfolio-based analysis[47]. A central aspect of much of this work is the mutual adaptation between attackers and defenders. McGill et al[39], for instance, deals with attacker preference changes in response to security measures. Merrick and Parnell[37] explicitly refer to the attacker adaptations to defender decisions and advocate further development of adaptive adversary risk assessment, and Rios and Rios Insua[38] (2012) similarly deal with strategic adversaries. Bier's[48] work shows how game theoretic models capture significant aspects of the decision problem faced by risk managers anticipating, with uncertainty, the responses of adversaries – such as the centralization of defensive choices, and choices to leave targets undefended. Research such as that of Chatterjee et al[47] similarly shows how we can deal with the interactions between risk controls that seem inevitable in systems of any complexity.

Much of this work has come to be called Adversarial Risk Analysis and is characterized by having two decision makers with opposing interests rather than the single decision maker of traditional, probabilistic risk analysis[49] . Its main limitation, as a general approach to analysis, is that it deals with the way adversaries adapt to risk controls but not how non-adversaries – such as a system's operators – also adapt to such controls. Our aim in this paper is to argue that this is a significant problem, and to propose a way of dealing with it. We also argue that this problem, in general terms, has a key difference from that of adversarial risk analysis. In adversarial analysis, it is typically reasonable to portray the attacker as having the opposite motivation to the defender – to maximize rather than minimize the risk. But it would be unreasonable to portray legitimate system operators and users in this way. The reason they adapt to risk controls is not in order to nullify or

avoid those controls but to pursue some orthogonal objective, such as reducing time and intellectual or physical effort. Therefore to predict their adaptations we need another theory, one not primarily concerned with risk. Our approach is to base this on the idea of affordances, as explained and justified below.

2.3 The adaptation problem

It is not hard to anticipate how attempts to make passwords more secure by forcing them to be more complex can lead to individuals writing them down, using the same passwords on multiple systems, sharing them with colleagues, and adopting other insecure practices. It would not be surprising if people forced to go through extensive authentication processes leave computers logged-on when they leave their desks unattended. It is quite likely that people having to exchange emergency information in a crisis will break rules that prohibit the sharing of confidential information. More generally, any attempt to intervene in a complex adaptive system, in order to control risk, is likely to produce an unexpected and perhaps self-defeating response[50].

This principle that risks are defined as much by a system's adaptations as by exogenous threats is central to some theories of disaster, such as Vaughan's[51] idea of 'normalized deviance'. It occurs in Reason's[52] model of organizations operating in a protection-production space, in which gains in protection are typically exchanged for gains in production. And it emerges in observations of risk migration, when attempts to manage risk simply transform it from one type to another[53]. Phenomena such as risk compensation and risk homeostasis[54] indicate that people adapt to risk controls even when they realise this adaptation can undermine those controls. Similarly, work on the violation of information security policies intended to control risk[55] shows how readily people can rationalise such violations.

Adaptation to technical risk controls – such as aircraft collision avoidance technology – has been implicated in large-scale risk events such as the Überlingen disaster[56]. But adaptation is especially likely in connection with social controls and the control of social elements of system vulnerability. Much of what we know about counter-productive adaptation to risk controls involves the use of organizational rules in particular[51] [57] [58] [59] [60]. In the security context, many modes of

attack are largely technical, involving vulnerabilities and associated exploits in specific software and hardware components. Yet many modes of attack involve 'social engineering'[61], for which social controls are often the most important and sometimes the only feasible form of risk mitigation. It is these that are the most liable to provoke adaptations when implemented.

3 DEVELOPMENT OF THE METHOD

3.1 A synopsis

We first summarize the basic principles of the proposed approach, before explaining the detail. Figure 1 shows the essential ecology. There is an interaction between risk manager and adversary as they create and encounter protective measures in the system, with broadly opposing objectives of trying to minimize or maximize risk. There is also an interaction between risk manager and legitimate users of the system (which includes those operating and maintaining it), through the protective measures created by risk managers and negotiated, confronted, accepted or bypassed by the users. The users broadly have the objective of achieving some level of convenience, minimizing effort, or generally being able to engage in some task without impedance.

Figure 1 here

The analysis is based on the principle of *affordances* – the action possibilities created by the fit between an actor and its environment. Its starting point is a graph that represents how risks materialize over time during attacks on a system. The edges in this graph are the affordances presented to the adversary by the vulnerabilities of a system. The materialization of a risk involves spanning the graph from a starting node to some node at which harm takes place. Edges in the graph are weighted by the uncertain time taken to exploit these affordances, and the first aim of the assessment is to find the path of least duration spanning the graph from start to end nodes. This is a heuristic intended to minimize the effort needed for the subsequent analysis. It assumes that if all paths can be raised above some

minimum duration, the system is resilient to a required degree because intrusion along all such paths can be detected and intercepted. If the shortest path falls below this minimum, candidate risk controls have to be identified, and an attempt is made to anticipate how these controls change the affordances experienced by legitimate users of the system. This predicts how such users might exploit alternative affordances. The final step, before iterating, is to assess how this then changes the original adversary affordance graph. The process repeats until the minimum path duration is acceptable.

In this approach risk is not measured directly, and the test of acceptable risk is not that a direct measure of risk falls below some risk acceptance threshold. Instead the test is that the time for a risk to materialize rises above some minimum tolerable handling time. This avoids the need to quantify the probability of an adversary choosing a specific method of attack – an inherently problematic requirement[3]. And it more clearly concerns the capacity of the system to withstand attack. Figure 2 outlines the method and in the remainder of this section we explain and justify the method in detail, discussing separately the main components of the analysis. In Section 4 we apply the method to a case study.

Figure 2 here

3.2 The affordance principle

Affordances are the possibilities for action in an environment that fit an actor's capabilities[8] [9]. It is this idea of fit between actor and environment, expressed in terms of potential actions, that makes the affordance principle so promising as a basis for identifying adaptations. Affordances often advertise themselves in some sense – being action possibilities that actors notice and are motivated to exploit[62] [63]. But they emerge in the relationship between actor and environment, and do not inhere in either separately[64].

There are various ways in which the affordance concept has been formalized, and Sahin et al[65] provide a brief survey. Stoffregen's[64] approach is to express affordances as a 'higher-order' property of the actor-environment system in the form of a relation between 'lower-order' properties of

actors and environments within that system. Chemero[66], who reviews earlier attempts at formulating a theory of affordances, sees them as relations between 'features' of an environment and 'abilities' of an organism. Steedman's[67] approach is to stress the way in which an affordance provides a linear implication from some pre-condition to a post-condition. And Sahin et al[65] incorporate the effect of an affordance into its definition. We follow this approach, broadly, considering the effect as a cost that can be used to predict how an actor might choose among competing affordances[63] [68] [69]. Thus we generally represent an affordance $f \in F$ as a tuple $f = <\phi, \chi, \pi, \kappa, \sigma_s, \sigma_d>$ in which $\phi \in \Phi$ is a relevant feature of the environment, $\chi \in X$ is a relevant property of an actor, $\pi \in \Pi$ is the action possibility, $\kappa \in K$ is some cost of exploiting this, $\sigma_s \in \Sigma$ is a source state in which the action possibility is available (that is, a pre-condition) and $\sigma_d \in \Sigma$ a destination state which exploitation of the action possibility produces (the post-condition). This is illustrated in Figure 3, showing how an affordance consists of an action possibility produced in the relationship between a feature of an actor's environment and its own properties. This action possibility, when exploited, makes a transition from source to destination states at some cost.

Figure 3 here

In practice, large parts of the affordance space can remain unspecified. In the analysis of the attack risk, as we explain below, the actor properties are irrelevant because no attempt is made to predict attacker decisions. Their properties (such as skill levels) are regarded as being random, in order to avoid making fragile assumptions about people and groups that analysts may have little knowledge about. On the other hand, in the analysis of user adaptations to risk controls, where the pre- and post-conditions are not informative, it is these that are left unspecified.

3.3 The attack affordance graph

Stoffregen[70] points out that affordances are often nested, and this is clearly the case for an attacker, for whom one affordance's destination state is a subsequent affordance's source state. For example,

contemporary accounts of the recent cyber attack on Ukrainian electricity distribution[6] [7] indicated that malware was introduced through an email attachment, infiltration was from the operators' business networks into their control networks, and the attackers could then directly control breakers in the distribution system. Prior work in this area similarly represents attacks as developing progressively. It is thus commonly based on attack trees or graphs, Markov processes, state space models, Bayesian networks and so on[71] [72] [31] [73], in which the central feature is that being in one state gives access to another.

To represent this progressive materialization through the sequential exploitation of affordances, our approach also uses a directed, acyclic graph $R = < N, E >$ of nodes N and edges E. This is illustrated in Figure 4. Each node $n = <y, v>$ is a state in which an attacker has access to a particular device $y \in Y$ (for example a computer, router, programmable logic controller (PLC) and so on), at a particular level of privilege $v \in V$. This is similar to prior work, such as McQueen et al's[74] 'compromise graph'. Each edge $e \in E$ is associated with an affordance $f = <\phi, \pi, \kappa, \sigma_s, \sigma_d>$ where σ_s, σ_d (the source and destination states) are its source and destination nodes $n_s = <y_s, v_s>$ and $n_d = <y_d, v_d>$. The environmental feature ϕ is typically some vulnerability in device y_s, and the action possibility π is some exploit available to someone having access at level v_s to y_s. As described in section 3.4, the graph is weighted by a measure of time, so it is natural to associate this time with the cost of the affordance although, as we explain, this is *not* used to predict an attacker's decisions.

Figure 4 here

What is of interest are paths P from a specific starting access state n_{start}, to which no edges point, to a state of access to harm n_{harm} from which no edges emerge. A path essentially defines a 'risk': some route from the status quo to a harm, with an associated probability and consequence. The materialization of a risk involves traversing the associated path, and this generally defines a trajectory of increasing privilege to devices with increasing proximity to physical processes. A risk control $c \in C$ is a modification of one edge e that modifies its weight $t(e)$, as described below.

3.4 The time metric

The basic principle is to estimate the uncertain exploitation time that it would take to traverse the edges in the affordance graph, and find the minimum time to span the graph from start node to harm node. This focus on the shortest path in the network does *not* assume that attackers will always take the shortest path, but *does* assume that if an attack on the shortest path can be intercepted (because the shortest path is greater than some threshold time) then any attack on a longer path can also be intercepted. 'Interception' could mean preventing the further development of any attack beyond a certain node through a periodic software patching cycle, or could mean the detection of a specific attack through periodic audit and monitoring activity. This focus on the shortest path is a heuristic that makes the process substantially more economical, as will become evident in Section 3.4, but it is just a heuristic and it is possible to imagine circumstances in which it will not work. These are discussed further in Section 5.2.

A range of methods is available for eliciting the probability distributions over the estimated edge times (see for example the survey by Garthwaite et al[75]. The choice needs to be based on what, in a particular context, is a suitable statistic for a subject area expert to judge and whether, for example, a judge is better able to estimate absolute bounds or quantiles[76]. We choose to use van Dorp's[76] procedure based on a family of two-sided power distributions, which has a bounded domain and straightforward definition that subject area experts are likely to find intuitive:

$$f(t \mid t_{min}, t_{mode}, t_{max}, n) = \begin{cases} \lambda\,((t - t_{min})/(t_{mode} - t_{min}))^{n-1} & \text{for } t_{min} < t < t_{mode} \\ \lambda\,((t_{max} - t)/(t_{max} - t_{mode}))^{n-1} & \text{for } t_{mode} \leq t < t_{max} \end{cases}$$

where $\lambda = n/(t_{max} - t_{min})$

Figure 5 illustrates $f()$ for a fixed choice of t_{min}, t_{mode}, t_{max}, with varying n. It is uniform for $n = 1$. Its use requires expert estimation of an exponent n as well as the bounding and modal values of t, for example by choosing from distributions expressed pictorially. Van Dorp[76] also describes a method that avoids eliciting the bounding values, instead eliciting quantiles and deriving the bounding values

from these. But in a context such as ours it appears no easier to elicit quantiles than bounds, and the procedure for deriving the latter from the former is complex, approximate and opaque.

Figure 5 here

Prior methods such as McQueen et al's[74] find the expected duration of paths spanning the graph from start to end states, adding means of relevant edges. But there may be an appreciable probability that the shortest mean path is not the shortest path, and that the materialization time for a risk is substantially less than the mean. So it makes sense to estimate, for each path, the duration for which there is some probability Q it will be exceeded by an actual duration – for example to find the path durations $t_{0.95}(p_i)$ that are exceeded with $Q = 0.95$, and to iterate through the process of identifying risk controls and adaptations until the shortest path duration $Min_i\,[t_Q(p_i)]$ is greater than the threshold duration T (which, in the case study below, is 24 days). We use a simple Monte Carlo procedure, assuming edge durations are independent.

It is difficult to specify a sensitivity analysis because, as explained further below, the procedure iterates through qualitative judgments of appropriate risk controls, anticipated adaptations to the controls, and expected ramifications of the adaptations. The most important outcome of the quantitative analysis at each iteration is the selection of the shortest path through the graph for comparison with the acceptable minimal path duration. It therefore becomes necessary to know whether the selection of the shortest path is marginal, and sensitive to small changes in the elicited probability distributions. So we measure the probability that the duration of any path could be less than that of the shortest, $q_{error}\,(p_i) = Pr[t_Q(p_i) < t_Q(p_{min})]$, again using a Monte Carlo procedure. Then q_{error} is made available for a decision about whether to revisit the estimation of the durations associated with the edges in the shorter paths through the graph.

There are various alternatives to these assumptions about appropriate distributions, elicitation procedures and acceptance criteria. The evidence on expert elicitation seems to suggest that mode estimates are reasonably accurate, but that estimates of variances and distribution tails are less so[75] (Garthwaite et al, 2005), so how judgments about probabilistic time should be expressed remains in

need of further exploration. It is not the main component of our method, and not its particular contribution, so we pass over these issues here.

3.5 The user adaptation problem

Our central concern is to anticipate how legitimate actors such as users adapt to potential risk controls *C*, when these modify the affordances they currently experience, F_{legit}, and how this in turn modifies the attack affordance graph *R*. It would be an extremely onerous task to identify F_{legit} completely. However, an analyst only needs to identify controls (perhaps only one control) for what is currently the shortest path in the attack affordance graph, and the effect of a new control might only be to change the costs of existing affordances rather than to add completely new affordances or completely withdraw old ones.

Having identified the way a proposed control *c* impedes some prior affordance $f = (\phi, \chi, \pi, \kappa)$ by modifying its cost to κ_{post}, for example, the analyst needs to anticipate how such a modification might provoke an adaptation. This involves the user exploiting some other affordance *f*' in place of *f*, if *f*' achieves the same goal as *f*, and does so at a lower cost. In this case, the relevant user property in the affordance definition, χ, is the user's goal. So we expect *f*' to substitute for *f* if both $\chi' = \chi$, and $\kappa' < \kappa_{post}$. An example we give in Section 4 is where attacks via malicious emails suggest controls that prohibit file exchange by email. This kind of control might provoke legitimate users into making file exchanges by USB (Universal Serial Bus) devices – dropping the affordance *f* of exchanging a file (the goal χ) by email when they have access to a device (some ϕ), because the cost κ_{post} involves the possibility of disciplinary action. Instead the affordance *f*' of exchanging a file using a USB drive is now taken up.

Finally, the analyst has to anticipate the effects of exploiting *f*' on *R*. In particular, it can modify the traversal times for the edges, making certain attack paths longer or (more likely) shorter. In the example in Section 4, the routine use of USB drives is expected to make illicit use of USB ports less noticeable, discourage system managers from disabling USB ports on devices in public places (like

some of the utility's commercial offices), and lead users to be less suspicious of USB drives left in such places. This modifies a number of edges in *R*, including one on the shortest path.

The process in practice relies heavily on expert judgment, requiring deep social understanding of the goals that legitimate users are pursuing, the alternative ways they have to meet them, and how these can make attacker moves more or less difficult. The procedure moves from the domain of an attacker's affordances to that of users, and then back again. These two domains have some common ground in terms of shared features of the environment (the states of access on various devices), but the affordances are directed to orthogonal purposes. But the analyst's task is limited: there is no need to identify user affordances F_{legit} completely, only the affordances *f* modified by introducing a candidate control, and the alternative affordances *f'* that share a goal with the modified *f*. There is also no need to identify the space of user goals X completely, only the shared goal $\chi(f, f')$ for the modified *f*.

3.6 An orientation process

As the process is founded on a series of expert judgments it is important to ensure these are well-informed, and this is not just a matter of developing a coherent analytical framework but also of having a systematic process of engagement between expert and system under analysis. Assessment of social engineering attack time distributions, in particular, is contextual and judgmental. It has to be informed by a social process in which the analyst engages.

Gollman's[77] basic taxonomy of an attack – under the headings of 'impact' (states of harm), 'exposure' (actors intending some harm), and 'vulnerability' (states of the system facilitating the harm) – provides a way of organizing this process of engagement. This suggests the following main elements:

1) 'Impact' requires identification of the primary attack target(s) within the industrial control system. These are states of harm, which the system offers affordances to reach.
2) 'Exposure' requires identification of actors who illegitimately experience these affordances. Identifying such affordances requires understanding of the potential motives and capacities of these actors.

3) ‘Vulnerability’ requires identification of the affordances offering direct and derived access to the targets. This will typically require a) identification of device interconnections and perimeter devices; b) search for poor configuration and lack of security patching; c) assessment of social norms, cues and responses; d) discovery of responsibilities and competences for protection.

This process is *ad hoc* and reliant on the artfulness of the analyst. But the principle of affordances remains the guiding idea. It is the inherent fit between attacker capabilities and system functions that create the risk, and the fit between legitimate user capabilities and system functions that create adaptations when risk controls are introduced. These need to be the focus of the analyst’s familiarization with, and orientation to, the organization and system under assessment.

4 APPLICATION OF THE METHOD

4.1 The organization and its systems

To illustrate the use of the method, it was applied in a small utility providing the final distribution of a commodity to consumers in a specific region of a European state. For reasons of confidentiality and security the organization has been anonymized, and some of the analysis simplified and changed. The orientation process described in Section 3.6 involved an intensive two-day’s work by a researcher who, until recently, had worked as a computer security professional for five years in the specific field of industrial control systems. He carried out interviews, worked through documentation, and undertook some basic vulnerability analysis at the organization’s site. The primary observations from the orientation process were as follows:

1) The primary target was a single SCADA server with affordances to various types of harm, most obviously the disruption of service to domestic or commercial consumers. An attacker would be able delay the restoration of service by excluding access for legitimate users, typically by changing all account passwords.
2) The candidate sources of an attack were multiple, and included disenchanted former employees. The different types of attacker were associated with distinct affordances. For example, former

employees had greater affordances for certain elements of ‘social engineering’, whereas remote attackers with technical capabilities were associated with the affordances presented by technical vulnerabilities in certain systems.

3) The vulnerabilities were both social and technical. Interviews suggested a susceptibility to malicious emails as employees appeared willing to open messages and attachments from known senders. Firewalls had been configured to allow direct remote access for subcontractors who interacted with certain subsystems from remote locations over the public Internet.
4) There was public accessibility to machines with exposed, active USB ports. And there was a strong culture of treating clients as clients, not as vulnerabilities or threats. This helped the organization serve its community in a principled way, but suggested a possible reluctance to see clients as potential sources of vulnerability or threat.

The ways in which specific utilities become vulnerable to specific risks is documented and modelled in the literature: for example Zechman[23] on water distribution systems, and Zhang et al[31] on electrical distribution systems. But this understanding of local conditions is an important element in a risk assessment of a particular system in a particular social organization.

4.2 The first iteration

Figure 6 shows the initial graph of attack affordances in the organization. Moving from left to right moves to access on devices that are increasingly close to physical controls with increasing privilege. There is a single start node, denoting an undifferentiated state in which any attack begins. There is a single end node denoting a state in which the attacker causes physical harm. The remaining nodes define access to some entity (for example ‘Eng WS’, denoting an engineering workstation and ‘File SRV’ denoting a file server) at some level of privilege (for example ‘@ Root’, denoting root or administrator level). The system under analysis was a very small one. Its computing network had three perimeter nodes, one accessible locally via Wi-Fi, one accessible over the public Internet, and one configured within a so-called DMZ (‘De-militarized zone’) on a gateway router. These define the early states of access in the graph, directly linked to the ‘Start’ node. Access with root privilege on a SCADA (supervisory control and data acquisition) server was the primary target in the sense that it

enabled a user to alter physical processes and states of operational sites. Henry and Haimes[26] argue that it is naïve to specify targets, as these emerge in practice as an attack progresses. But in such a simple system as this one, it seems reasonable to do so. There were also programmable logic controllers linked to the network but these were excluded from the analysis by request of the operating organization. The graph has been somewhat modified to avoid undue disclosure, and simplified to aid exposition. Some edges were feasible, in the analyst's judgment, but did not form part of a coherent failure path and so were excluded. In general, the shorter paths through the affordance network (for example those in which there is a direct connection from the 'Start' node to later nodes such as root-level access on an engineering work station), involve some kind of social engineering.

Figure 6 here

Table I shows the analyst's exploitation time estimates, in days, for each edge. As McQueen et al[74] argue, the expertise of the attacker is clearly relevant to exploitation time (or 'time to compromise' in their work), but it is also relevant to the frequency of attacks – and in this case more frequent attacks were expected from less expert attackers. The probability distribution over exploitation times expresses the analyst's judgments over all sources of exploitation for each edge. Most edges have a common minimum of 1 day reflecting the analyst's assumption that an attack would not proceed over *N* affordances in less than *N* days. Many edges in this particular analysis share a common maximum of 21 days, particularly those associated with technical rather than socially-engineered modes of attack. McQueen et al suggest this as a general value where there are no known vulnerabilities and exploits and the attacker has an expert level of capability. We make no assumptions about capabilities, so this estimate is conservative. The exponent value, common to all entries, gives a distribution with substantial weight in the tails, indicating a relatively large uncertainty for the analyst.

Table I here

Table II shows the paths p_i in this graph, together with associated values of duration, $t_{0.90}(p_i)$, the durations for which the probability of exceedance is 0.90. The table also shows the mean durations $E[t(p_i)]$ and the probabilities $q_{error}\ (p_i) = Pr[t_{0.90}(p_i) < t_{0.90}(p_{min})]$ that the paths may be shorter than the identified shortest path.

Table II here

The analysis shows a primarily technical path, path 14, to be the shortest. The control proposed at this stage was a technical one, reconfiguring the WAN/LAN (Wide Area Network/Local Area Network) gateway such that a port formerly open to the public Internet is brought within a firewall. There were no affordances that were modified by the introduction of this control in the analyst's judgment.

4.3 The second iteration

The second graph is a trivial adaptation of the first, modifying the exploitation time distribution for two edges (edges 16 and 17). On re-analysis, path 6 was then the shortest according to $t_{0.90}(p_6)$. But this duration remained below the acceptance threshold of 24 days. The path involved a spear-phishing attack giving direct root access to the engineering work station, for example via a PDF (Portable Document Format) attachment that is vulnerable to an exploit in a PDF viewer, or via an executable file disguised with a .pdf extension. The analyst had discovered, unsurprisingly, that most employees in the organization would open an email attachment with high probability if apparently sent from another employee, and if it appeared to be technical in nature they would open it anyway – almost regardless of its source (for example if it appeared to be a consumer emailing them with a supposed photograph of a defective installation).

The obvious risk control was to prohibit email attachments, introducing an organizational rule that forbids the opening of received attachments and forbids file attachments on sent emails. But fairly clearly this reduces the affordance for legitimate users – the organization's employees – to exchange files *via* email attachments and the analyst judged it highly likely that employees would simply violate

the rule. This would restore the legitimate users' affordance but nullify the effect of the control, leaving the attack affordance graph virtually unchanged.

4.4 The third iteration

The analyst therefore proposed the automatic removal of all file attachments from emails received in the organization. Again this undermines the affordance provided by the email system to share files, but this affordance cannot be restored by simple rule violation. Users of the system are therefore likely to substitute competing affordances for achieving the same goal. The most obvious competing affordance was provided by USB drives. In the analyst's judgment, the result would be that the use of USB drives would become normal and habitual, and that employees would not find it suspicious when an unidentified USB drive was observed in a USB port. Moreover, there were several office work stations in a public area with exposed, enabled USB ports – allowing a visitor to insert surreptitiously a USB device with key logging or wireless access. This led to a modification of the exploitation times attached to edge 2 in another path (path 7) in the affordance network – a path that involved socially engineered strategies using USB drives. The analysis of the modified graph made path 7 the shortest, with a $t_{0.90}(p_7)$ of 22 days, which remained below the 24 day acceptance threshold.

4.5 The fourth iteration

The final iteration involved finding a control on this second path – involving the blocking of USB ports in public spaces, and the provision to staff of distinctively-marked USB drives for all file sharing activity. The analyst could identify no relevant affordances that were materially affected, and the control brought the shortest path just above the criterion threshold.

For such a small example this kind of iteration is likely to be intuitive and regarded as part of what we would expect a competent organization to do as a matter of course. But some of the steps laid out here had *not* been obvious to the organization in question. It had no systematic knowledge of the possible routes of attack laid out in the attack affordance graph, it had no specific understanding of which paths were the ones that presented the greatest vulnerability given their possible traversal times, and it had no planned series of risk controls. There was no apparent appreciation of the way in which,

if it did have to introduce additional controls, this could lead to adaptations that might create or amplify further risks. The outcome of the analysis was therefore not only a better-protected system, but also a deepening of the organization's understanding of its own systems and their vulnerabilities, and of its own responses to these vulnerabilities.

5 DISCUSSION AND CONCLUSION

5.1 Contributions of the method

Although this study is about a specific class of risk – deliberate attacks on industrial control systems – its basic premise and method are more general. Any system which receives some intervention to control risk can adapt to that intervention and change the risk in a way that was not the intention behind the control. Therefore assessment should naturally be an iterative process – and our aim was to propose a method for carrying it out. This method makes the iteration relatively economical, by concentrating attention on the shortest known risk materialization path. And it makes the iteration relatively systematic, by using a consistent principle (that of affordances) to anticipate adaptations to risk controls. The iteration reflects the fundamental idea that risk is reflexive, and that knowing about a risk almost inevitably changes it in some way.

The principle that these adaptations take place because risk controls change affordances appears to be a general one. And the affordances principle has a big advantage. Because affordances concern the relationship between actors and environments in terms of actions, they help predict in a direct way how a change in behavior might follow a change to system. We do not have to predict how actors build their representations of the system and work out their plans for using it[(69)], because actors 'pick up' affordances in a direct and effortless way[(78)] . This directness makes affordances a plausible explanation of what happens and also a more practical way of predicting what *can* happen. They offer 'considerable heuristic guidance'[(64)] when we search for possible adaptations. There are various developments of the affordance principle that on face value appear to suit risk analysis – notably the idea of affordances as probabilistic functions[(79)]. But to adopt such ideas would involve quantifying an

attackers' probability of acting in a particular way, a step fraught with the difficulties of estimating how such probabilities will change as attackers discover information about the system they are targeting[3]. So it is the basic idea of an affordance that appears to have most relevant to risk assessment generally.

We have also suggested there has to be a clear 'orientation' process. This is not part of the formalized structure of the assessment procedure but it is an essential part of putting it into practice. In our case study, for example, the analyst found that to predict how system users would react required provoking them into responding – confronting them with typical malicious emails and getting them to describe what actions they would take. The analyst also only found out through physical engagement that an important vulnerability was created by a system that had been configured to give remote access to a camera server used by an external organization to monitor the organization's protected sites. This orientation process is especially important in an iterative assessment that takes account of possible adaptations, because the analyst needs not just a record of the system as it currently exists but also an understanding of the affordances that people in the system currently exploit – and what they might do if those affordances were modified.

5.2 Limitations and further work

Because our intention was to analyze the adaptations of the legitimate system and its users we have neglected a range of issues addressed in earlier work. We do not deal with the adaptations of adversaries to risk controls[39] [37] [38], which includes the possibility that risk controls may inadvertently create new affordances for an attacker. We do not address the interactions between different risk controls[47]. We do not deal with the effects on dependent infrastructures[80], and we ignore the costs of risk measures, which form an important component of other methods[46]. We also neglect the problem of multiple actors facing independent choices about protective actions that are nonetheless inter-dependent[81]. And we do not address the analysis of threat agents, their capacities, incentives and motivations. Further work is therefore needed to synthesize all these elements in some way, although the complexities of the individual methods suggest a single, over-arching method may

not be feasible. Perhaps what we should be looking for is a way of defining individual methods in such a way that they can be selected and combined for the purposes of any specific risk assessment.

Probably the greatest weakness of the method as it stands is the working assumption that if all path durations can be raised above some threshold with a given confidence level then risk is acceptable. As a reviewer of an earlier version of this article emphasized, this ignores the variation in detectability of the paths, and the attraction to potential attackers of the less detectable paths. In the current method, the times assigned to the edges in these paths are the times it takes the attacker to negotiate them, not the times that the defender has in order to respond to the attack. An edge with a long negotiation time may offer very little response time if its detectability is low in some way. While attackers are negotiating some edge, they have access to some device at some level defined by the start node of that edge. So, to deal with the problem of variable detectability, an analyst would need to quantify the detectability δ_n of every node n in R that starts an edge. This detectability expresses for what fraction of time an attacker occupying a state is detectable in that state, which nominally starts at the time the state is attained and ends at the time a transition is made out of it. This is the time associated with an outgoing edge from that node. For example, if an attacker having root status to an office workstation has a 0.25 detectability level, offering an affordance to gaining user status on a SCADA workstation after an expected period of 5 days, or to root status after 10 days, the detectable period of these affordances becomes 1.25 and 2.5 days respectively. It is the detectability-adjusted path durations in the graph that then have to be compared with the time threshold. Of course detectabilities may themselves be uncertain, and the Monte Carlo procedure for finding shortest path durations to some confidence level would also then sample from probability distributions over detectability values.

Another complexity is that legitimate users' adaptations may respond not just to changes in affordances but also to their evaluations of security. If users know a particular risk measure is intended to control a risk on a short compromise path then they might not undermine it through a self-serving adaptation. This contingency between adaptation and threat perception reinforces the need for the analyst to make judgments. Our approach is just a framework that helps prompt and organize these judgments, but it does not avoid the need to make them and emphasizes the importance of what we

have called the ‘orientation process’ of properly familiarizing oneself, as an analyst, with social conditions in the organization under analysis.

The study also neglected the problem of scalability. The organization used in the case study, and its systems, were small in scale. This made it suitable for a research study, but raises questions about how well the method will scale up. Assessing a specific system necessarily requires a detailed representation of that system, and an assessment based on identifying, testing and mitigating a minimum path through a failure graph necessarily requires a comprehensive identification of all possible paths that could be minimal. But how the judgmental and computational burden increases with the scale of the system depends on the structure of the system, and its decomposability. The economizing property of the method is its focus on the minimal path spanning the graph, which means that identification of risk controls, and identification of potential adaptations, is limited to this path. But – as we indicated earlier – this rests on the assumption that the path of minimal duration is always the worst case with some known probability. When this cannot be known with confidence, the analyst has to explore more paths and the method becomes less economical, and much more sensitive to increasing scale.

We also make assumptions about the best approach to eliciting and expressing judgments about uncertainty and time, generally assuming such distributions are uni-modal, and assuming the distributions over different edges in the graph are independent. These assumptions need examining in any particular context. It is plausible, for example, that attacker capabilities could cluster around expert, highly-skilled actors and unskilled actors who only use existing, download-able exploits. Such clustering could produce bi-modal distributions. However, the core of the method is the identification of affordances, and the prediction of adaptations. Other distributions, and joint distributions, could be substituted for those we have used, and the use of simulation rather than an analytical approach means that substitution is relatively straightforward. Overall, the basic principle that risk analysis in contexts like this should not stop at analyzing the status quo, but also analyze candidate risk controls and potential adaptations to them, appears robust. And the use of the affordances principle to do this appears productive.

ACKNOWLEDGMENTS

Many thanks are due to the members of the organization in which the case study took place. We are also very grateful to the anonymous reviewers and the editor for their perceptive and constructive advice. This study was partly funded by an EU Seventh Framework Programme grant, no. 608090, on 'Hybrid Risk Management for Utility Networks'.

REFERENCES

1. Stouffer K, Falco J and Scarfone K. Guide to Industrial Control Systems (ICS) Security. Gaithersburg, MD : National Institute of Standards and Technology, 2013; Special Publication 800-82.
2. CPNI. The National Infrastructure. UK Centre for the Protection of National Infrastructure, 2014. Retrieved March 11, 2014 at http://www.cpni.gov.uk/about/cni/.
3. Brown GG and Cox LA. How probabilistic risk assessment can mislead terrorism risk analysts. Risk Analysis, 2011; 31, 196-204.
4. Miller B, and Rowe D. A survey of SCADA and critical infrastructure incidents. Proc. 1st Annual Conference on Research in Information Technology, Calgary, October 10-13th 2012, pp. 51–56.
5. Pauli D. Phishing proves too hot for plant: Hackers pop German steel mill, wreck furnace. The Register. Posted 22nd December 2014 at http://www.theregister.co.uk/2014/12/22/hackers_pop_german_steel_mill_wreck_furnace.
6. Pultarova T. Ukraine grid hack is wake-up call for network operators. Engineering and Technology, February 2016, 12-13.
7. Zetter K. Everything we know about Ukraine's power plant hack. Wired, 20th January. Last accessed on 26th February 2016 at http://www.wired.com/2016/01/everything-we-know-about-ukraines-power-plant-hack/.
8. Gibson JJ. The theory of affordances. In Shaw RE and Bransford J (eds.), Perceiving, Acting, and Knowing: Toward an Ecological Psychology, Hillsdale, NJ: Lawrence Erlbaum Associates, 1977, pp. 67-82.

9. Norman DA. Affordance, conventions and design. Interactions, May – June 1999, 38-42.
10. Cowlagi RV and Saleh JH. Coordinability and consistency in accident causation and prevention: formal system theoretic concepts for safety in multilevel systems. Risk Analysis, 2013; 33: 420–433.
11. Sharit J. A modeling framework for exposing risks in complex systems. Risk Analysis, 2000; 20: 469-482.
12. Pasman HJ, Knegtering B, and Rogers WJ. A holistic approach to control process safety risks: Possible ways forward. Reliability Engineering & System Safety, 2013; 117: 21-29.
13. Goossen LHJ. Risk prevention and policy-making in automatic systems. Risk Analysis, 1991; 11: 217–228.
14. Einarsson S and Rausand M. An approach to vulnerability analysis of complex industrial systems. Risk Analysis, 1998; 18: 535–546.
15. Rastogi A and Gabbar HA. Fuzzy-logic-based safety verification framework for nuclear power plants. Risk Analysis, 2013; 33: 1128–1145.
16. Paté-Cornell M-E and Regan PJ. Dynamic risk management systems: hybrid architecture and offshore platform illustration. Risk Analysis, 1998; 18: 485–496.
17. Podofillini L and Dang VN. Conventional and dynamic safety analysis: comparison on a chemical batch reactor. Reliability Engineering & System Safety, 2012; 106: 146-159.
18. Kobza JE and Jacobson SH. Addressing the dependency problem in access security system architecture design. Risk Analysis, 1996; 16: 801-812.
19. Nan C, Eusgeld I, and Kröger W. Analyzing vulnerabilities between SCADA system and SUC due to interdependencies. Reliability Engineering & System Safety, 2013; 113: 76-93.
20. Sikorski M. Use of digital simulation in reliability analysis for the design of industrial process control systems. Reliability Engineering & System Safety, 1991; 31: 281-295.
21. Garrett CJ and Apostolakis GE. Automated hazard analysis of digital control systems. Reliability Engineering & System Safety, 2002; 77: 1-17.
22. Coulbeck B and Orr CH. Essential considerations in the computer control of water distribution systems. Reliability Engineering & System Safety, 1993; 42: 55-64.

23. Zechman EM. Agent-based modeling to simulate contamination events and evaluate threat management strategies in water distribution systems. Risk Analysis, 2011; 31: 758-772.
24. Yan Z and Haimes YY. Cross-classified hierarchical Bayesian models for risk-based analysis of complex systems under sparse data. Reliability Engineering and Systems Safety, 2010; 95: 764–776.
25. Bobbio A, Ciancamerla E, Franceschinis G, Gaeta R, Minichino M and Portinale L. Sequential application of heterogeneous models for the safety analysis of a control system: a case study. Reliability Engineering & System Safety, 2003; 81: 269-280.
26. Henry MH and Haimes YY. A comprehensive network security risk model for process control networks. Risk Analysis, 2009; 29: 223-248.
27. Ralston PAS, Graham JH and Hieb JL. Cyber security risk assessment for SCADA and DCS networks. ISA Transactions, 2007; 46: 583–594.
28. Kundur D, Feng X, Mashayekh S, Liu S, Zourntos T and Butler-Purry K L. Towards modelling the impact of cyber attacks on a smart grid. International Journal of Security and Networks, 2011; 6: 2–13.
29. Naedele M. Addressing IT security for critical control systems. Proceedings 40th Annual Hawaii International Conference on System Sciences HICSS2007, Jan 3-6th 2007, p.115.
30. Fleury T, Khurana H and Welch V. Towards a taxonomy of attacks against energy control systems. In Para M and Shenoi S (eds.). Critical Infrastructure Protection II. Berlin: Springer, 2008, pp. 71-85.
31. Zhang Y, Wang L, Xiang Y, & Ten C-W. Power system reliability evaluation with SCADA cybersecurity considerations. IEEE Transactions on Smart Grid, 2015; 6: 1707-1721.
32. Ten C-W, Liu C-C and Manimaran G. Vulnerability assessment for cybersecurity for SCADA systems. IEEE Transactions on Power Systems, 2008; 23: 1836-1846.
33. Henry MH, Layer RM, Snow KZ and Zaret DR. Evaluating the risk of cyber attacks on SCADA systems via Petri net analysis with application to hazardous liquid loading operations. IEEE International Conference on Technologies for Homeland Security, 11-12th May 2009 pp. 607-614.

34. Henry MH, Layer RM and Zaret DR. Coupled Petri nets for computer network risk analysis. International Journal of Critical Infrastructure Protection, 2010; 3, 67-75.
35. Ingols K, Chu M, Lippman R, Webster S and Boyer S. Modeling modern network attacks and countermeasures using attack graphs. 25th Annual Computer Security Applications Conference ASAC '09, Austin, December 6-10th 2009, pp. 117-126.
36. Long J and Fischhoff B. Setting risk priorities: a formal model. Risk Analysis, 2000; 20: 339-351.
37. Merrick J. and Parnell GS. A comparative analysis of PRA and intelligent adversary methods for counterterrorism risk management. Risk Analysis, 2011; 31: 1488-1510.
38. Rios J. and Rios Insua D. Adversarial risk analysis for counterterrorism modelling. Risk Analysis, 2012; 32: 894-915.
39. McGill WL, Ayyub BM and Kaminskiy M. Risk analysis for critical asset protection. Risk Analysis, 2007; 27: 1265-1281.
40. Rothschild C, McLay L and Guikema S. Adversarial risk analysis with incomplete information: a level-k approach. Risk Analysis, 2012; 32: 1219-1231.
41. Chatterjee S. and Abkowitz MD. A methodology for modeling regional terrorism risk. Risk Analysis, 2011; 31: 1133-1140.
42. Davis G, Garcia A and Zhang W. Empirical analysis of the effects of cyber security incidents. Risk Analysis, 2009; 29, 1304-1316.
43. Rose A, Oladosu G and Liao S-Y. Business interruption impacts of a terrorist attack on the electric power system of Los Angeles: customer resilience to a total blackout. Risk Analysis, 2007; 27: 513-531.
44. Simonoff JS, Restrepo CE and Zimmerman R. Risk-management and risk-analysis-based decision tools for attacks on electric power. Risk Analysis, 2007; 27: 547-570.
45. Ezell BC. Infrastructure vulnerability assessment model (I-VAM). Risk Analysis, 2007; 27: 571-583.
46. Shafieezadeh A, Cha EJ and Ellingwood BR. A decision framework for managing risk to airports from terrorist attack. Risk Analysis, 2015; 35: 292-306.

47. Chatterjee S, Hora SC and Rosoff H. Portfolio analysis of layered security measures. Risk Analysis, 2015; 35: 459-475.
48. Bier V. Choosing what to protect. Risk Analysis, 2007; 27: 607-620.
49. Rios Insua D, Rios J and Banks D. Adversarial risk analysis. Journal of the American Statistical Association, 2009; 104: 841-854.
50. Forrester JW. Counterintuitive behavior of social systems. Theory and Decision, 1971; 2: 109-140.
51. Vaughan D. The Challenger Launch Decision. Chicago: University of Chicago Press, 1996.
52. Reason J. Managing the Risks of Organizational Accidents. Aldershot, UK: Ashgate, 1997.
53. Busby JS, Alcock RE and MacGillivray BM. Types of risk transformation: a case study. Journal of Risk Research, 2012; 15: 67-84.
54. Wilde GJS. The theory of risk homeostasis: implications for safety and health. Risk Analysis, 1982; 2: 209-225.
55. Siponen M and Vance A. Neutralization: new insights into the problem of employee information systems security policy violations. MIS Quarterly, 2010; 34, 487-502.
56. Busby JS and Bennett SA. Loss of defensive capacity in protective operations: the implications of the Überlingen and Linate disasters. Journal of Risk Research, 2007; 10: 3-27.
57. Hynes T and Prasad P. Patterns of 'mock bureaucracy' in mining disasters: An analysis of the Westray coal mine explosion. Journal of Management Studies, 1997; 34: 601–623.
58. Wicks D. Institutional mindsets of invulnerability: Differentiated institutional fields and the antecedents of organizational crisis. Organization Studies, 2001; 22: 659–692.
59. Mascini P. The blameworthiness of health and safety rule violations. Law and Policy, 2005; 27: 472–490.
60. Reason J, Parker D, and Lawton R. Organizational controls and safety: The varieties of rule-related behavior. Journal of Occupational and Organisational Psychology, 1998; 71: 289–304.
61. Kaivanto K. The effect of decentralized behavioral decision making on system-level risk. Risk Analysis, 2014; 34: 2121-2142.

62. Rietveld E. Situated normativity: the normative aspect of embodied cognition in unreflective action. Mind, 2008; 117: 974-1001.
63. Withagen R, de Poel HJ, Araujo D and Pepping G-J. Affordances can invite behavior: reconsidering the relationship between affordances and agency. New Ideas in Psychology, 2012; 30: 250-258.
64. Stoffregen TA. Breadth and limits of the affordance concept. Ecological Psychology, 2004; 16: 79-85.
65. Şahin E, Çakmak M, Doğar MR, Uğur E and Üçoluk G. To afford or not to afford: a new formalization of affordances toward affordance-based robot control. Adaptive Behavior, 2010; 15: 447-472.
66. Chemero A. An outline of a theory of affordances. Ecological Psychology, 2003; 15: 181-195.
67. Steedman M. Plans, affordances, and combinatory grammar. Linguistics and Philosophy, 2002; 25: 723-753.
68. Reed E S. The intention to use a specific affordance: a framework for psychology. In Wozniak R and Fisscher K (eds.), Development in Context: Acting and Thinking in Specific Environments. Hillsdale, NJ: Lawrence Erlbaum Associates, 1993, pp. 45–75.
69. Cisek P. Cortical mechanisms of action selection: the affordance competition hypothesis. Philosophical Transactions of the Royal Society B, 2007; 362: 1585-1599.
70. Stoffregen TA. Affordances are enough: Reply to Chemero et al (2003). Ecological Psychology, 2003; 15: 29-36.
71. Kanoun W, Cuppens-Boulahia N, Cuppens F, Dubus S and Martin A. Success likelihood of ongoing attacks for intrusion detection and response systems. Proc. International Conference on Computational Science and Engineering (CSE'09), Vancouver, 29-31st August 2009, pp. 83-91.
72. Pietre-Cambacédès L and Bouissou M. Modeling safety and security interdependencies with BDMP (Boolean logic Driven Markov Processes). Proc. IEEE International Conference on Systems Man and Cybernetics (SMC), Istanbul, 10-13th October 2010, pp. 2852 – 2861.

73. Nzoukou W, Wang L, Jajodia S and Singhal A. A unified framework for measuring a network's mean time-to-compromise. Proc. IEEE 32nd International Symposium on Reliable Distributed Systems (SRDS), Braga, 30th September-3rd October 2013, 215-224.
74. McQueen MA, Boyer WF, Flynn MA and Beitel GA. Quantitative risk reduction estimation methodology for a small SCADA control system. Proceedings 39th Annual Hawaii International Conference on System Sciences, HICSS2006, 4-7th January 2006, p.226.
75. Garthwaite PH, Kadane JB and O'Hagan, A. Statistical methods for eliciting probability distributions. Journal of the American Statistical Association, 2005; 100: 680-701.
76. Van Dorp JR. Indirect parameter elicitation procedures for some distributions with bounded support – with applications in program evaluation and review technique (PERT). Structure and Infrastructure Engineering, 2012; 8: 393-401.
77. Gollmann D. Computer Security. 3rd Edition. Chichester UK: Wiley, 2011.
78. Zhang J and Patel VL. Distributed cognition, representation, and affordance. Pragmatics and Cognition, 2006; 14: 333-341.
79. Franchak J and Adolph K. Affordances as probabilistic functions: implications for development, perception, and decisions for action. Ecological Psychology, 2014; 26: 109-124.
80. Chang SE, McDaniels T, Fox J, Dhariwal R and Longstaff H. Toward disaster-resilient cities: characterizing resilience of infrastructure systems with expert judgments. Risk Analysis, 2014; 34: 461-434.
81. Heal G and Kunreuther H. Modeling interdependent risks. Risk Analysis, 2007; 27: 621-634.

Table I. Estimates for the initial affordance graph

Edge	$Node_{start}$	$Node_{end}$	$Time_{min}$	$Time_{mode}$	$Time_{max}$	*Exponent*
0	Start	WiFi	1	10	21	1.8
1	Start	EngWSRoot	5	15	40	1.8
2	Start	OfficeWSRoot	5	20	40	1.8
3	Start	WLGateway	1	4	21	1.8
4	WiFi	EngWSUser	1	10	21	1.8
5	WiFi	EngWSRoot	1	20	30	1.8
6	WiFi	FileSVRRoot	1	16	21	1.8
7	WiFi	OfficeWSRoot	1	16	21	1.8
8	WiFi	FileSVRUser	1	10	21	1.8
9	WiFi	OfficeWSUser	1	10	21	1.8
10	WLGateway	EngWSUser	1	10	21	1.8
11	WLGateway	FileSVRUser	1	12	21	1.8
12	WLGateway	EngWSRoot	1	20	30	1.8
13	WLGateway	OfficeWSUser	1	10	21	1.8
14	WLGateway	FileSVRRoot	1	20	30	1.8
15	WLGateway	OfficeWSRoot	1	20	30	1.8
16	WLGateway	CamSVRRoot	1	4	21	1.8
17	WLGateway	CamSVRUser	1	3	21	1.8
18	EngWSUser	EngWSRoot	1	8	21	1.8
19	FileSVRUser	FileSVRRoot	1	8	21	1.8
20	OfficeWSUser	OfficeWSRoot	1	8	21	1.8
21	CamSVRUser	CamSVRRoot	1	4	21	1.8
22	EngWSRoot	ScadaRoot	1	4	21	1.8
23	FileSVRRoot	ScadaRoot	1	8	21	1.8
24	OfficeWSRoot	ScadaRoot	1	6	21	1.8
25	CamSVRRoot	ScadaRoot	1	4	21	1.8
26	ScadaRoot	Harm	1	2	21	1.8

Table II. Estimates of path durations

Path	$t_{0.90}(p_i)$	$E[t(p_i)]$	*Node 1*	*Node 2*	*Node 3*	*Node 4*	*Node 5*	*Node 6*	$q_{error}(p_i)$
0	28.0	44.9	Start	WiFi	EngWSRoot	ScadaRoot	Harm		0.24
1	25.6	41.1	Start	WiFi	OfficeWSRoot	ScadaRoot	Harm		0.30
2	33.0	49.0	Start	WiFi	EngWSUser	EngWSRoot	ScadaRoot	Harm	0.16
3	27.2	41.7	Start	WiFi	FileSVRRoot	ScadaRoot	Harm		0.33
4	34.0	50.1	Start	WiFi	FileSVRUser	FileSVRRoot	ScadaRoot	Harm	0.15
5	33.0	49.6	Start	WiFi	OfficeWSUser	OfficeWSRoot	ScadaRoot	Harm	0.15
6	22.8	37.8	Start	EngWSRoot	ScadaRoot	Harm			0.44
7	23.5	39.8	Start	OfficeWSRoot	ScadaRoot	Harm			0.38
8	26.3	43.2	Start	WLGateway	EngWSRoot	ScadaRoot	Harm		0.28
9	26.3	43.8	Start	WLGateway	OfficeWSRoot	ScadaRoot	Harm		0.28
10	30.0	47.3	Start	WLGateway	EngWSUser	EngWSRoot	ScadaRoot	Harm	0.17
11	27.1	44.4	Start	WLGateway	FileSVRRoot	ScadaRoot	Harm		0.26
12	32.0	49.0	Start	WLGateway	FileSVRUser	FileSVRRoot	ScadaRoot	Harm	0.16
13	32.0	47.9	Start	WLGateway	OfficeWSUser	OfficeWSRoot	ScadaRoot	Harm	0.23
14	20.8	35.4	Start	WLGateway	CamSVRRoot	ScadaRoot	Harm		N/A
15	28.0	44.1	Start	WLGateway	CamSVRUser	CamSVRRoot	ScadaRoot	Harm	0.26

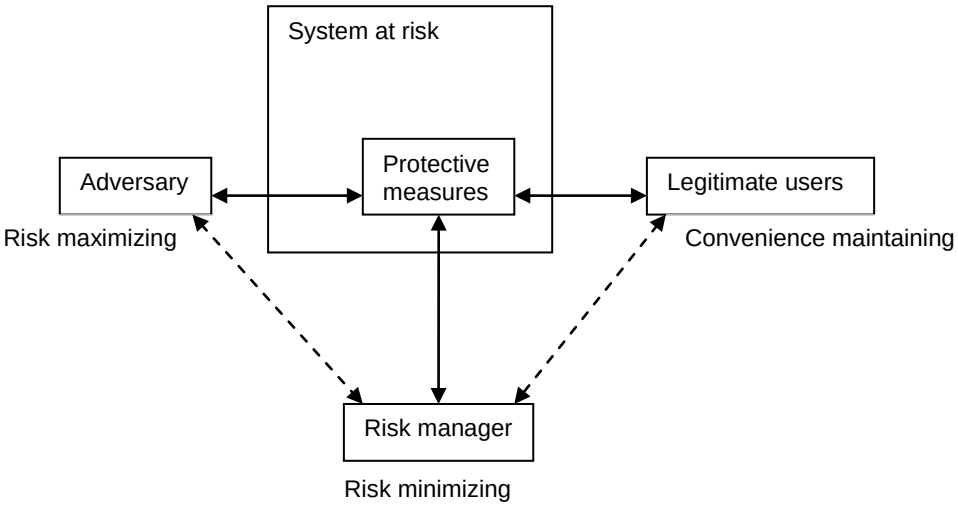

Fig. 1. The essential ecology

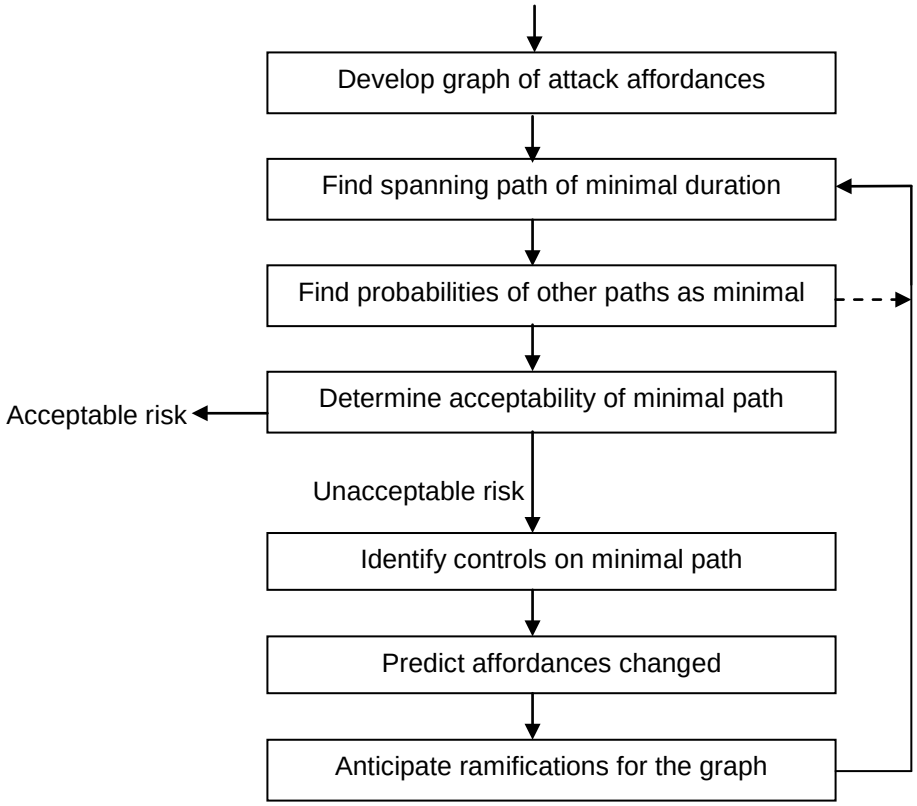

Fig. 2. Outline of the assessment method

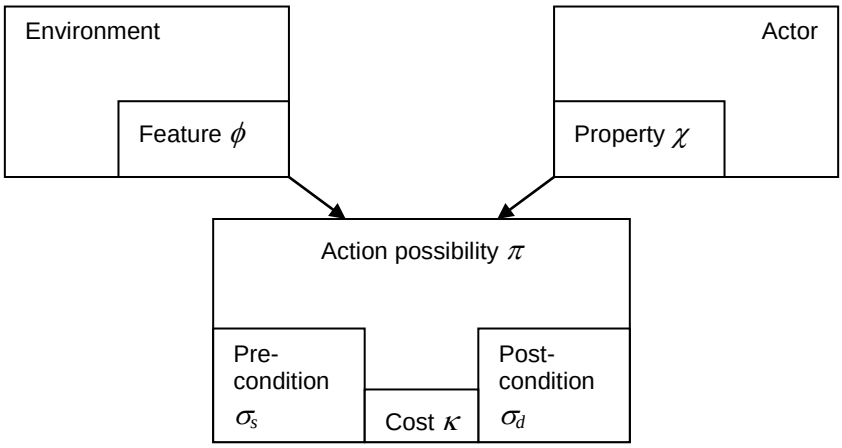

Fig. 3. The affordance principle

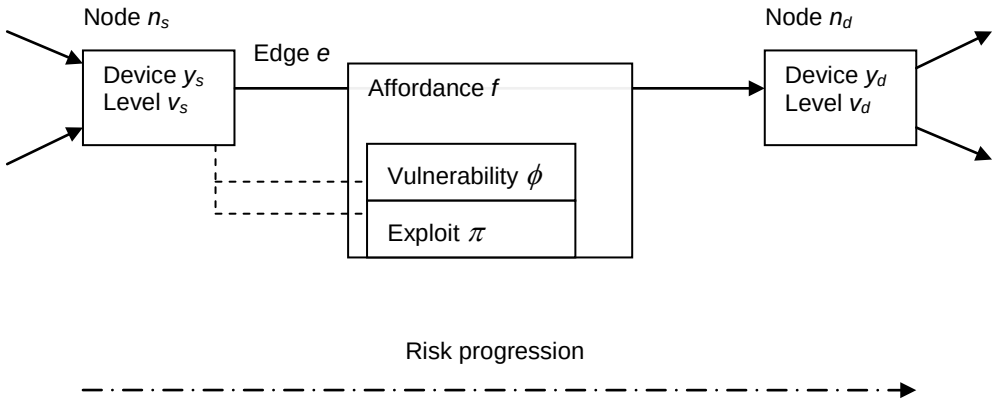

Fig. 4. The principle of the attack affordance graph

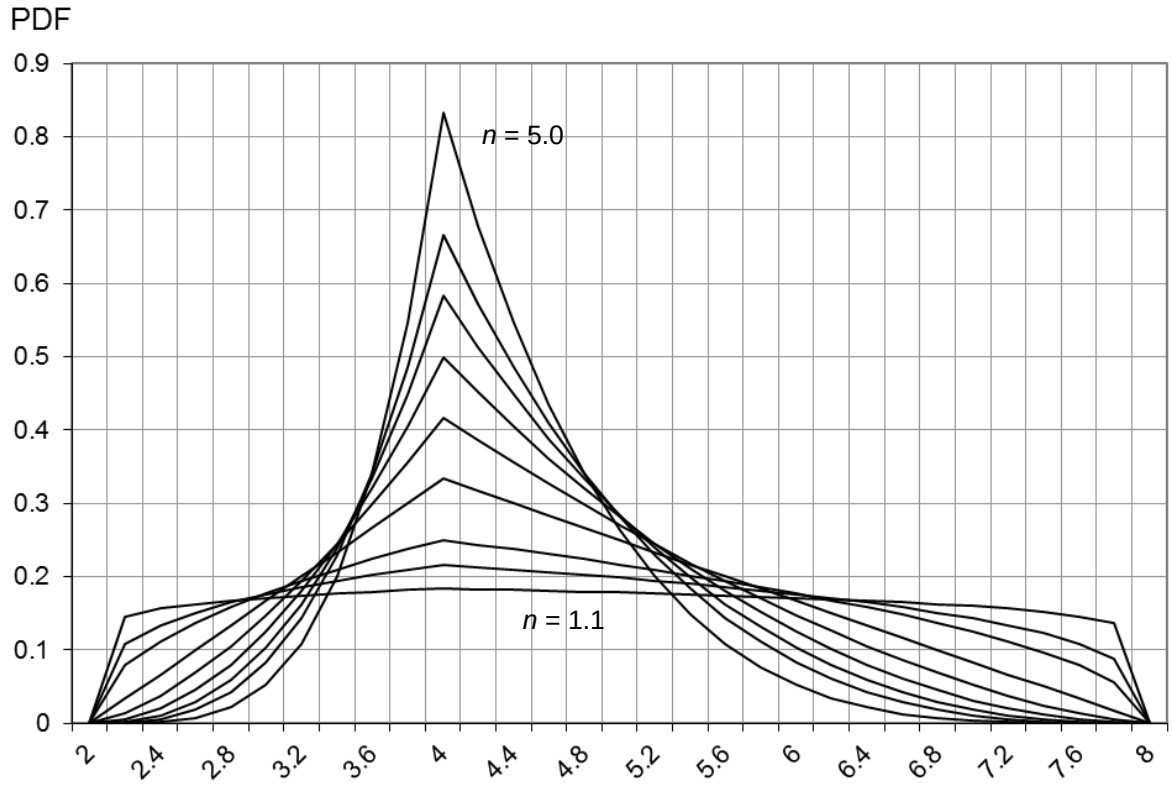

Fig. 5. The two-sided power distribution with $t_{min} = 2$, $t_{mode} = 4$, $t_{max} = 8$, for n from 1.1 to 5

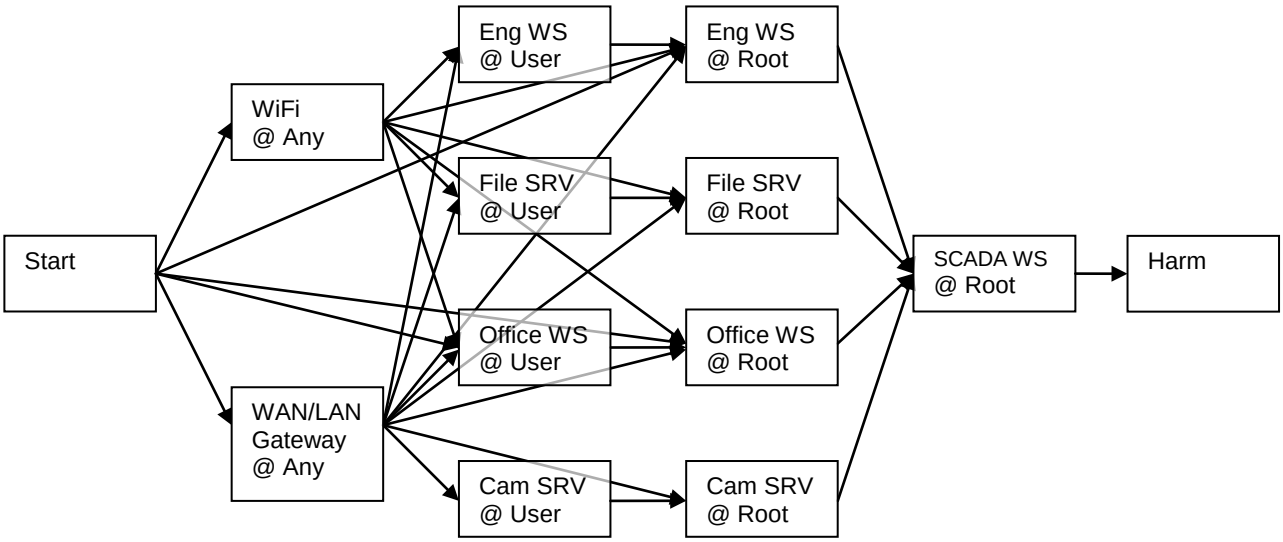

Fig. 6. The initial affordance graph