
On safety and security requirements in emerging ubiquitous computing models

EMMA OSBORN AND ANDREW SIMPSON

Department of Computer Science, University of Oxford, Oxford OX1 3QD
Email: Emma.Osborn@cybersecurity.ox.ac.uk; Andrew.Simpson@cs.ox.ac.uk

The fields of safety and security are converging due to a number of factors, including the rise in system interconnectivity and an increased dependence on the Internet as part of critical national infrastructures. Partly as a reflection of this, there is a wealth of literature pertaining to the increasing interdependence between safety and security. While much of this research has been concerned with large-scale industrial systems, the rapid emergence of what might be termed *Consumer Cyber Physical Systems* (Consumer CPS) means that it is crucial that such issues are considered in that context also. We evaluate the motivations for implementing Consumer CPS and the novelty of safety and security concerns that such systems give rise to. This evaluation is subsequently used to establish a collection of cyber security requirements for this emerging domain. We also consider how these requirements might impact upon product lifecycles. Our contribution is motivated and illustrated by three representative scenarios.

*Keywords: cyber physical security; safety-critical systems; requirements;
ubiquitous computing; cyber physical systems*

Received 00 January 2009; revised 00 Month 2009

1. INTRODUCTION

Two decades ago, Burns, McDermid and Dobson gave consideration to the relationships that exist between safety and security [1]. Much has changed in the past twenty or so years: software-based systems are now used in contexts that could not have been imagined in the early 1990s; in addition, the security landscape has altered significantly in the intervening period. Going further, the language of safety and security has changed: terms such as ‘on-line safety’ and ‘cyber security’ now form part of our everyday discourse. These factors lead us to reconsider our definitions of security and safety, and the relationships that exist between them in emerging contexts.

While the fields of safety and security have evolved relatively independently of each other, there has been a steady stream of research contributions that have investigated the potential for the ‘cross-fertilisation’ of ideas. Schneier’s presentation of *attack trees* [2] is a manifestation of this cross-fertilisation; Rushby’s introduction of *safety kernels* [3] is another. (Piètre-Cambacédès and Bouissou [4] provide an authoritative survey of the transfusion of techniques.)

We concern ourselves with the emerging context of ubiquitous computing and interconnectivity — in which safety can often only be maintained through cyber

security. In particular, we argue that the traditional views of ‘safety’ and ‘security’ will, inevitably, have to adapt to this new reality. Our focus is *Consumer Cyber Physical Systems* (Consumer CPS): consumer products containing embedded systems that take advantage of the increase in available wireless technologies to connect to the Internet. We evaluate the motivations for industry to implement Consumer CPS and why they may want to include new safety and/or security measures. Our evaluation of the ecosystem driving Consumer CPS development is used to characterise the cyber security requirements for this emerging domain, with the aim of these requirements being to maintain safety levels in Consumer CPS irrespective of new risks.

In Section 2 we give consideration to the domain. We also consider the relationship between safety and security. Section 3 presents our methodology and our motivational scenarios, which are then used to support the narrative of Section 4, where the attributes and market drivers for the development of Consumer CPS are discussed. The scenarios are also used in Section 5, where existing safety legislation is discussed. Section 6 discusses cyber security principles that are less closely linked to safety, and data protection legislation that might prompt their implementation. In Section 7 the issues discussed in the previous sections are summarised

as a set of system properties that can be used to begin building a set of requirements. Section 8 considers interpretation and possible implementation of these requirements. Finally, we draw conclusions and discuss possible areas of future work in Section 9.

2. THE EMERGING CONTEXT

The fields of safety and security are converging, whether intentionally through adaptation of models in one discipline to include aspects of the other (see, for example, [4]) or through the inclusion of aspects of cyber security in safety models to meet on-going safety requirements (see, for example, [5]). This convergence is being propelled by increases in system interconnectivity and automation, as well as by an increased dependence on the Internet to underpin critical national infrastructures.

Research on this convergence has largely concentrated on the fields of nuclear power and Industrial Control Systems, as a result of the changing risks to Cyber Physical Systems (CPS).¹ However, with the emergence of ubiquitous computing models, CPS of a different scale are now entering offices, public spaces and the home [7], with attacks on these systems being reported [8]. Our primary concern is to assess how the safety of what we term *Consumer CPS* can be maintained through the inclusion of cyber security measures, learning from adaptations made to safety and security models in the industrial sector.

2.1. A sectoral view

The use of cyber physical systems has been discussed under a variety of headings, including *ubiquitous* or *pervasive* computing and the *Internet of Things* (IoT). Many attempts have been made to define how ubiquitous computing might develop, including Beecham Research's consideration of the Internet of Things on a sectoral basis.² From this and other sources, we can identify a number of service sectors where ubiquitous computing may have an impact: buildings; energy and water; consumer and home; healthcare and life sciences; industrial and agricultural; transportation; retail; security and public safety; and IT and networks.

Table 1 gives consideration to these sectors, together with where a consumer might come into contact with them in a 'typical' day: home (H), work (W), transport systems (T), retail outlets (R), and social spaces (S). The table shows that the only commonly defined sector for ubiquitous computing that a consumer does not come into direct contact with is the industrial and agricultural sector. This sector, along with the non-public-facing aspects of energy and water production,

Sector	H	W	T	R	S
Buildings	•	•		•	•
Energy & Water	•	•		•	•
Consumer & Home	•	•		•	•
Healthcare & Life Sciences	•	•	•	•	•
Industrial & Agricultural					
Transportation			•		
Retail			•	•	•
Security & Public Safety			•	•	•
IT & Networks	•	•	•	•	•

TABLE 1. Consumer contact with IoT sectors

are those most often discussed when considering the synergy between safety and security [4, 9, 10]. On the other hand, two sectors — healthcare and life sciences, and IT and networks — connect with every activity space. The reason for the former is relatively obvious — the consumer carries their healthcare devices with them in what has been termed the *Body Area Network* (BAN) [11]. For the latter, this is due to the fact that, while the IoT is still an emerging phenomenon, the increasing use of wireless technology means that the availability of connectivity is pervasive — whether or not the technology is in place to exploit it. Modern systems are designed around the expectation of connecting to the Internet, with a lack of network availability becoming the unusual and difficult use case to handle.

As computing becomes ubiquitous, the lines between IoT sectors are likely to blur as new services emerge. An example of this is the overlap of the IT and healthcare sectors, where smartphones carried everywhere become part of the BAN — an integral tool in controlling healthcare products and communicating information. The need to carry a smartphone at all times might be partially due to the healthcare applications it runs, but the introduction of a non-sector-specific tool into the BAN opens up the network for use by other sectors' products.

2.2. Safety and security

Given the increasing interdependence between safety and security concerns, it is timely to reflect upon the wealth of literature in this area. The motivation for the 'merging' of the two fields is discussed in [12]:

"The reason for this is that IT-systems are embedded in ever more influential parts of our living- and working environment and that these embedded IT-systems are networked — be it to enhance their functionality now (or just as an option for future use), be it to ease maintenance." [12]

We provide an overview of existing research pertaining to the relationship between safety and security — considering the language of safety and security, with a view to establishing some terms of reference.

¹We adopt Lee's definition: "Cyber-Physical Systems (CPS) are integrations of computation with physical processes" [6].

²See <http://www.beechamresearch.com/article.aspx?id=4>.

2.2.1. Language

The language of safety and security is, arguably, more prominent than has ever previously been the case. Terms such as *national security*, *cyber security* and *on-line safety* are now part of our everyday discourse — and are examples of how the lines between safety and security can become blurred.

It is worth noting that, even within the traditional safety and security communities, there is a tradition of definitions being relatively ‘fuzzy’. To quote Piètre-Cambacédès and Chaudet [13]:

“Dozens of explicit, but distinct, definitions can be found . . . ranging from slightly different to completely incompatible definitions. In this situation, searching for absolute, universal definitions is bound to fail.” [13]

The situation is not helped by the fact that some languages have a single word for the two terms (German — *sicherheit*, Spanish — *seguridad*, Portuguese — *segurança*, Swedish — *säkerhet*, and Danish — *sikkerhed*).³ Despite the distinctions in English, we still find some blurring of the lines. For example, if we consult the on-line version of the Oxford English Dictionary,⁴ both safety and security have definitions that talk of “freedom from danger”; further, the term *safety factor* is defined as “a margin of security against risks” and the term *security risk* is defined as “a person or thing which poses a possible threat to safety or security.” We see similar overlaps in the definitions given by Merriam-Webster,⁵ where security is defined as “the state of being protected or safe from harm,” and safety is defined as “freedom from harm or danger.”

2.2.2. Technical definitions

Both safety and security have the fundamental premise that they are concerned with preventing something ‘bad’ from happening — effectively, the prevention of threats to safety or security, with the nature of threats differing in either case. Typically (although not exclusively), the former pertains to the prevention of something bad happening accidentally, while the latter pertains to the prevention of someone doing something bad deliberately. Brostoff and Sasse [15] argue as follows:

“It is sometimes argued that a major difference between these domains is that safety failures are frequently accidents, whereas security breaches are often deliberate (and so are likely to happen again and again). This difference is greatly reduced if we assume that the system exists in a dangerous world.

When we focus on the victim/end-user (as a




FIGURE 1. Use of the definitions of *safety* and *security*.

computer security policy must do) instead of the perpetrator/external cracker, we see that safety and security breaches will happen unless the victim takes appropriate steps to avoid them.” [15]

The *SEMA* referential framework of [13] has such distinctions at its heart:

- *System vs. Environment (S-E) distinction*: “Security is concerned with the risks originating from the environment and potentially impacting the system, whereas safety deals with the risks arising from the system and potentially impacting the environment.” [13]
- *Malicious vs. Accidental (M-A) distinction*: “Security typically addresses malicious risks while safety addresses purely accidental risks.” [13]

We base our definitions on this framework. To this end, we characterise safety and security risks thus:

- *Safety* risks originate from the system, accidentally impacting on the environment.
- *Security* risks originate from malicious actors in the environment, intentionally impacting the system.

We characterise the *environment* of these definitions as “the set of other interacting systems whose behaviour and characteristics are generally less known and beyond the control of the system owner” [13]. For the purposes of this paper, this definition is assumed to include the physical environment and underlying communications systems, although Consumer CPS are assumed to be complex systems in their own right.

Figure 1 shows how these definitions are interpreted in the context of interest. The risks discussed encompass aspects of both safety and security, examining use cases where it is possible for the actions of malicious parties to impact on the physical security or safety of the environment. Environments 1 and 2 could be the same environment, or entirely different

³As observed by Burns *et al.* [1], Piètre-Cambacédès and Chaudet [13], and Simpson *et al.* [14] (amongst others).

⁴<http://www.oed.com>

⁵<http://www.merriam-webster.com>

environments due to the increased interconnectivity of systems inherent in ubiquitous computing. Consumer CPS are thus differentiated from typical consumer devices by their capacity to cause physical harm in the local environment (Environment 2) following networked intervention from Environment 1.

Piètre-Cambacédès and Bouissou suggest that, although safety can be legally delegated to manufacturers and designers, malicious risk (including cyber security) often needs to remain the remit of government due to the scope of the problem and the actors existing outside of the system [4]. This approach, whilst effective in the context of critical national infrastructure, is likely to be infeasible when considering Consumer CPS. As such, a different approach will be required.

There are existing legal frameworks in place for dealing with cyber security breaches in the consumer sector. For example, in the UK victims of cyber security breaches are encouraged to report incidents using the same platform provided for reporting fraud. The UK National Cyber Security Strategy⁶ states that this is for “citizens and small businesses to report cyber crime so that action can be taken and law enforcement agencies can establish the extent of cyber crime (including how it affects individuals and the economy).” Consumers typically have a responsibility to take basic measures to protect themselves from both avoidable harm and opportunistic crime, but the existing culture of not reporting cyber incidents makes it difficult to establish to what extent they are mitigating on-line threats.

3. METHODOLOGY

We employ a meta-study, which brings together the state-of-the-art from different disciplines, fields of research and stakeholder groups. We take a scenario-led approach, where the scenarios are pragmatic — describing systems that have been adapted to include elements profiting from connectivity, rather than new systems designed for purpose. Where possible, they provide a non-proprietary view, taken from different business sectors.

The scenarios identify how overlaps in safety and security models can be exploited; they consider system scope to include more than the individual appliance in isolation. The first two are typically used in motivating IoT. The third has similarities with the first two and consists of a subsystem touching the physical world, together with applications from different stakeholders that interface remotely with the physical system. However, there is no buy-in from the manufacturer of the physical aspects of the system for this type of subsystem to be used as envisaged by the application developers; this raises questions when considering business models and liability.

⁶https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/60961/uk-cyber-security-strategy-final.pdf

The scenarios have been used in conjunction with a literature review to survey the different elements of the Consumer CPS life-cycle. The intention is to reflect probable gradual evolution in system design taking into account markets that are driving the evolution. To this end, we have considered diverse sources to provide an accurate account of the properties of these systems. In such a cross-disciplinary subject it would not be possible to conduct an exhaustive literature review; rather, the question is approached using representative literature from each field to build the bigger picture.

We aim to establish a set of general requirements that a combined safety and security model might have to operate within. These requirements feed into a framework for safety through cyber security in Consumer CPS. The framework uses guidelines from the consumer product safety sector [16], in conjunction with the NIST cyber security principles [17], both of which provide high-level advice for practitioners. The developed requirements are used to describe the intersection between the two engineering processes when considered in the context of Consumer CPS.

3.1. Scenario 1: A ‘smart’ refrigerator

Jen has a state-of-the-art refrigerator. When it was delivered, the technician connected it to Jen’s wireless network. The technician set up a user account on the company’s web-site and told Jen to log in to finish registering and see the services she could use. The site gave her proof of purchase for the warranty, asked her if she wanted the refrigerator to automatically update its systems, and offered a remote-monitoring application plus some third party services. Through the recommended services, Jen could link the refrigerator’s in-built food recognition system to her supermarket shopping application, allowing her shopping list to auto-update when she was running out of her favourite foods. The monitoring application provides information about the food she has in her refrigerator, alerting her when an item is about to go out of date or if the refrigerator isn’t keeping the food cool enough. Jen also uses an application that automatically uploads information from her refrigerator into a diet diary.

3.2. Scenario 2: A home climate-control system

Jen’s building has a sophisticated climate-control system. When she moved into her flat, as well as the keys, the landlord passed on the details for logging in to the manufacturer’s web page. By logging in, Jen was able to change the climate-control settings on a room-by-room basis and download an application to allow her to turn on the heating in her house remotely. The application gave her the option of linking to her car’s satellite navigation system so that when she selected the ‘home’ location, the climate-control

system could automatically detect when Jen was close to home and turn the heating on. As well as linking to applications and the company web-site, the system links to the building's fire alarm system periodically providing health reports to help the alarm system detect faults before they become critical.

3.3. Scenario 3: An in-car media system

Jen has bought an application that is compatible with her car, which, once downloaded, allows her tablet or phone to connect to the in-car system via Wi-Fi. Previous iterations of the product were sold as hardware and wired in to cars. Among other things, the system lets Jen have access to all her music and movies through the dashboard screens. She finds a video online explaining how to install it and, once it is installed, she links her accounts so that her satellite navigation system can automatically look in her contacts to find addresses. The application also provides options for optimising engine control settings to the user's driving style.

4. CONSUMER CPS ATTRIBUTES AND DEVELOPMENT DRIVERS

4.1. Consumer CPS research

Cyber physical systems combine the physical (continuous) world with the digital (discrete) one [18]. Consumer CPS are a type of highly distributed CPS. While the size of the system (at least in terms of number of lines of code and/or budget) is enormously reduced, there is no associated reduction in complexity — any reduction in (say) the number of nodes is offset by the added complexity of many of the system elements being independently produced and with the designer's scope of control being greatly reduced. An example of this is the smart refrigerator scenario, where the product designer can in theory control which supermarkets have access to the refrigerator's API. However, even if the designer feels that a specific supermarket's application introduces too many vulnerabilities for it to be allowed access, the consumer will inevitably expect the service to be available for the largest supermarket chains.

Wolf *et al.* [7] discuss the fact that constraints in traditional CPS for real-time communication and safety are less present in consumer systems. They argue that this opens a market for the widespread sale of data, actuator or computation services. The inference is that, in being less bound by safety regulation, Consumer CPS designers can produce the innovative products that would lead to ubiquitous computing more rapidly than safety-critical systems have been adopted in the past.

Consumer CPS require both dynamic topologies and dynamic reconfiguration, while being potentially tool-neutral with modules not necessarily being combined until run-time [19, 20]. These types of products are also likely to be developed in part by small companies

contracted by a manufacturer. These small companies are competitive as they can keep the overheads of process to a minimum, meaning that the introduction of more rigorous design processes is likely to be resisted.

The number of products with embedded systems that could potentially go on-line are far more numerous than non-embedded systems currently are [21]. This means that the impact of a small safety issue, when multiplied by the potential number of affected devices, could become serious. Despite this, as well as the expectation of designers to have fewer constraints, consumers have a preconceived idea of the acceptable price of an appliance, making the budget to create solutions small [21]. An example of this is with the in-car media scenario. The perceived acceptable price of an application is low, but, as this application is designed to link into the safety-critical system inside of a car, it is, by implication, safety-critical and should be developed accordingly; however, no user would be willing to countenance paying the associated development costs.

Much of the research undertaken on CPS that reach the home has been in the context of *Smart Grids* [9, 18, 22]. Here, the security focus is on protecting electricity company assets in a hostile environment — the consumer's home. Our focus, on the other hand, is consumer safety and security, and, consequently, any existing models where service-provider security is the focus may need to be adapted accordingly.

The movement of CPS into the home potentially poses a high risk to consumers. In an environment where consumers expect any product they buy to be designed safely, few people are likely to undertake any meaningful evaluation as to how cyber threats emanating from the way they choose to use that product may impact their safety or security. For example, in our climate-control scenario, when Jen logs into the manufacturer's web-site, she is unlikely to consider how many past tenants still have access to her system.

4.2. Consumer CPS development drivers

As our scenarios demonstrate, while it is possible to define what a Consumer CPS is, what they interact with and the risks they pose can vary significantly on a case-by-case basis. In order to find the underlying patterns of similarities in these products, it is necessary to give consideration to a higher level of abstraction. To this end, Figure 2 provides an illustrated breakdown of the differences between the typical product manufacture business model and an IT-driven business model.

Traditionally, as depicted in Figure 2 (a), a consumer buys a variety of products and services. The consumer may choose to do some work once they've purchased a product to integrate it into a system, or their system may consist of a variety of items whose only connection is the consumer using them. Kagermann *et al.* [23] suggest that over time this business model is becoming increasingly less profitable, leading to

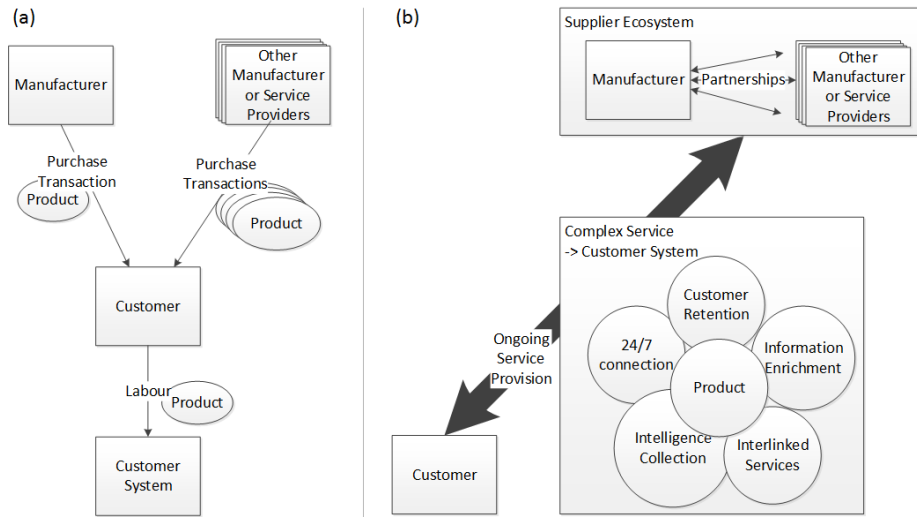


FIGURE 2. (a) A traditional manufacturing business model; and (b) an IT-driven business model.

IT-driven business models: at every opportunity, businesses are attempting to open a continued dialogue with their customers whether through their web-site or via a social media presence. This is highlighted by our scenarios where mobile apps play a key role in introducing interesting ways to interact with traditional appliances. To an extent, these models are shaped by customer value.

This provision of consumer gimmicks is in contrast to the use of an Industrial Control System (ICS) or smart grid scenario, where systems are put on-line in order to reduce monitoring costs [24]. Considered in the context of typically self-contained home appliances it is unlikely that a connection to the Internet of Things would reduce manufacturer operating costs, so a cost-reduction model is limited to the upgrade of home appliances linked to pre-existing services, such as smart meters, owned and operated remotely by third parties.

In the context of IT-driven business models manufacturers make money by reducing the cognitive workload of the consumer as the customer system can be pre-built and configured (Figure 2 (b)). The most mature implementations of this business model are in the IT sector itself where PC or mobile device manufacturers form partnerships with operating system and software development companies. These partnerships allow products to be sold ready to use ‘out-of-the-box’ and let consumers try various software before they buy.

In altering their business model to that of a service provider, manufacturers hope to increase reliability, as design standards are agreed within the supplier ecosystem. They also aim to improve customer retention by purposely building in reasons for continued interaction with the consumer. Being able to download an application and link supermarkets to the refrigerator via the manufacturer’s services means that the customer builds a lasting relationship with the company, and

provides opportunities for them to advertise how they can now link in a new freezer and dishwasher to the same system with no extra effort. All of these services depend on the product going on-line — which might be characterised as *interconnectivity by design*. Unfortunately, this business model draws manufacturers away from their area of expertise. New elements of the product or service may be seen as a fairly modular element of the design, with the possibility of outsourcing the development of applications, web portals, etc. to third parties. While modular design may help manufacturers when developing a complex system, this modularity of both the product and development teams makes it more difficult to judge the risks associated with the use of the system once it is connected.

The evolution of the business model in terms of safety and security requirements can be seen in Figure 3, where it can be seen that, in the past, the security or safety of different stakeholders has been considered separately. In the IT-driven business model, more connectivity leads to broader cyber security requirements and a safety requirement that is shared between the different stakeholder cyber presences, as well as affecting the original product.

This overlap of requirements illustrates why the supplier ecosystem of an IT-driven business model produces a complex system or service: product designers now have to consider the integrity of the services they are offering and the cyber security of the platforms on which they are hosted. As per Figure 1, by connecting a Consumer CPS to the Internet, new risks are introduced in the safety and physical security domains. While the new business model reduces the cognitive workload of the consumer, the broadening scope and complexity of safety and security requirements could over-burden system designers.

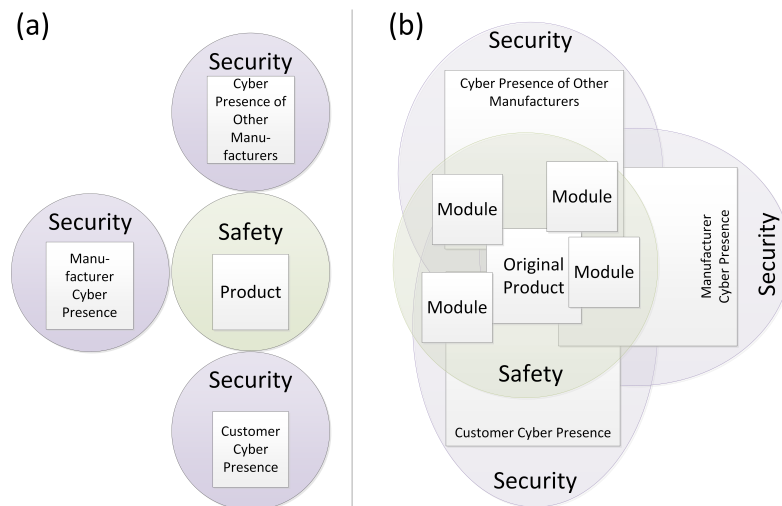


FIGURE 3. Scope of safety and security models in: (a) a traditional manufacturing business model; and (b) an IT-driven business model.

4.3. New sources of risk

With the implementation of new services in the Consumer CPS market, the types of risks associated with products are evolving to include new hazards. We highlight some of these in the following.

4.3.1. Remote operation of devices

As discussed in Section 3, one reason for providing add-on services is to allow 24/7 connectivity to devices and the possibility for remote operation. Operating devices remotely introduces new types of risk due to the operator not being able to assess the state of the environment as part of the decision-making process: operators will not know if there are other people dangerously close to an appliance, or carrying out maintenance, or if there are objects that could be damaged by their actions. In the climate-control example, turning on specific types of electric heater if there are items of clothing draped over them could be enough to start a fire. The operator's lack of presence may also affect the magnitude of an incident: if a malfunction leads to a fire, the fact that there is no one present may mean that it takes longer for the fire to be detected. Another implication of the remote operation of appliances is that those operations will be changing the internal state of the device at a distance. Should the network connection be severed part-way through an operation, or before a critical update, the device could be left unsupervised in an unsafe state.

4.3.2. Data integrity

Non-networked appliances are unlikely to be programmed with cyber security in mind — it is unlikely that signals will be tested for authenticity as those signals will probably originate from a component fixed to the same circuit board; in the case of safety-critical

functions, there may be limited tests for accuracy. Once an appliance goes on-line, the designer can no longer consider the appliance as a closed system. If the appliance has no closed system, the integrity of data received by the controller might be compromised more easily. It might be easy for a neighbour to use the link Jen's refrigerator has with her supermarket account to submit an order for their own food, or for the climate-control system to pass incorrect information to the fire-alarm system. The latter could result in an obvious safety issue; however, both of these examples could also cause financial losses for the consumer, leading back to the question of how cyber crime is handled, as per Section 2. The reporting of cyber incidents and safety faults is discussed in more detail in Section 5.

4.3.3. A third party's motivation to use resources

Proofpoint's incident report of January 2014⁷ was related to the use of various devices as resources to send spam. The motivation of third parties to transform devices into resources introduces a new type of risk to the household appliance sector — a direct consequence of increased interconnectivity. A refrigerator that sends spam is more entertainment than hazard, but, if instead of the refrigerator scenario, we consider the in-car media example, the consequences of an attacker borrowing computational power from one of the car's in-built controllers could be serious.

Data exchanged by the networked devices may also be used as a resource, damaging the privacy and potentially the physical security of consumers. Eavesdropping on conversations between Jen's refrigerator and her applications would give an attacker information about what she ate, but also potentially highlight health is-

⁷<http://www.proofpoint.com/threatinsight/posts/your-refrigerator-is-full-of-spam-part-11-details.php>

sues and bad habits. Jen’s climate-control application and the in-car media application are both linked to the car’s satellite navigation system, so could provide an attacker with her physical location. Privacy and confidentiality issues are discussed in Section 6.

4.3.4. *Malfunction*

The increased complexity of the system, coupled with the potential number of stakeholders, makes it a challenging task to identify all of the potential flaws associated with a Consumer CPS, particularly should an element of the system reach end-of-life and cease to be supported. Malfunctions caused by cyber attacks, as outlined in Section 2, are even more concerning as the ‘accidental’ safety implications of an attack can be replicated either in multiple appliances or multiple times by persistent activity on the same appliance. If ongoing system verification becomes harder, there is increased risk of malfunction leading to injury, damage of possessions or the product self-destructing. Statistics show that malfunctioning electrical appliances are a common cause of fires within the home;⁸ the increased risk of malfunction in a Consumer CPS due to faults (or even malicious actors) might have the potential to increase the risk of fire within the home.

5. THE OBLIGATION TO GUARANTEE SAFETY

5.1. Consumer safety legislation

The motivations for creating better safety models in the consumer sector are not the same as those in the industrial sector. Huge amounts of effort are being put into securing supervisory control and data acquisition (SCADA) systems, as well as other Industrial Control Systems. The motivation for this is clear — there are known attacks proving that these systems are vulnerable. For example, searching the Common Vulnerabilities and Exposures database for SCADA currently (May 2015) produces over 600 results.⁹ These systems tend to be operated by large companies, but often make up part of a country’s critical national infrastructure. The risk in terms of both safety and financial loss at this scale is (rightly) perceived as being unacceptable by its stakeholders.

The consumer market is markedly different: if a product has a cyber security vulnerability, this is unlikely to draw the attention of governments and large industry. Consumers have grown to expect flaws in their computer systems; as such, many do not get reported, even if there is a reporting system in place — this position in respect to cyber crime was illustrated in Section 2. In these circumstances, a manufacturer is unlikely to incur large financial or reputational losses,

and bug fixes may only be planned to be built into the next version of the product. Only in the case of serious malfunction, or where security issues are detected in network traffic, are faults likely to be reported.

If, in the refrigerator scenario, an attacker had managed to reprogram a controller to help mine bitcoins, the refrigerator might then be too occupied with the attacker’s task to turn the refrigeration system on and off at the correct intervals: it freezes the food and the consumer adjusts the temperature; or it stops working altogether and is returned; or it turns itself off more often than it should, reducing the shelf-life of the food and damaging the supermarket’s reputation. None of these are likely to lead to someone beginning a cyber incident response process. The fact that the context makes a cyber attack difficult to envisage may become an incentive for attackers to turn to these systems where they gain access for longer periods before detection, increasing the risks discussed in Section 4.3.

Given these issues (and, in most cases, a lack of financial motivation to improve either software quality or security), the main motivation for improvement in this field would be to ensure that this overlap between security and safety does not negatively impact on a manufacturer’s ability to prove their products meet the requirements of product safety legislation. Unfortunately, as outlined in the following subsections, the three issues of safety, security and software quality are intrinsically linked — forcing the optimistic perspective that, through either design or through-life processes, software quality will have to be addressed to some extent by manufacturers.

5.1.1. *EU product legislation*

An example of legislation intended to protect consumers is provided by the European Union (EU), the most recognisable consequence of which is the CE Marking, which relates to safety, health and environmental requirements. The *Blue Guide* [25] describes the implementation of EU product rules. The legislation covers all new products made available¹⁰ in the EU and second-hand products entering the market from outside the EU. Manufacturers have to conform to the legislation in order to be allowed to use the CE Marking. This typically requires the manufacturer to pass a conformity assessment ensuring that they have met a set of essential requirements, as well as the harmonised standard(s) relevant to that product or equivalent. The assessment process begins at the product design phase and continues throughout the manufacturing lifecycle.

A part of the conformity assessment is a risk assessment; the way that requirements are applied depends on how hazardous a product might be. To

⁸<https://www.gov.uk/government/collections/fire-statistics-great-britain>

⁹<http://cve.mitre.org/cve/index.html>

¹⁰The Blue Guide definition of *made available* is given thus: “made available on the EU market when supplied for distribution, consumption or use on the Union market in the course of a commercial activity, whether in return for payment or free of charge” [25].

quote the aforementioned Blue Guide, “manufacturers have to match a level of protection corresponding to the use they prescribe to the product under the conditions of use which can be reasonably foreseen” [25]. This means that the instructions that come with a product are, in some ways, as important as the product itself. For example, a washing machine having an ethernet port doesn’t necessarily mean the manufacturer has to be concerned about cyber threats to that appliance. Instructions provided by the manufacturer may have to instruct users to plug it in to the Internet for them to be liable for any accidents caused by cyber threats.

There are issues with this approach: instructions for use are obviously important legally, but are sometimes deficient (perhaps because they have been written by an individual who is unfamiliar with the product, or because there are translation issues, or because they are ambiguous or open to interpretation). Some manufacturers try to address the ambiguity of natural language by using images. However, while images can’t be translated incorrectly and take less time to read (and so are more likely to be used), there isn’t, unfortunately, a universal symbol for many of the elements discussed in user manuals — so diagrams and icons might also be open to interpretation.

In the software industry, the bulk of the instructions are provided as a result of risk analyses, aimed at protecting the developer from copyright theft or liability, with limited ‘getting started’ advice also provided. Instructions in the Terms and Conditions or End User Licence Agreements are usually presented in the form of hundreds of lines of impenetrable legalese that are often ignored by the user. This approach, together with the tactile nature of technology, means that users are far more likely to expect to be able to use something safely without reading any instructions, via intuition and learning as they go [26].

There is also the issue of the user not associating a software element of the system with a safety hazard in a physical element of the system. The application may be too removed from the hardware for the connection to be made implicitly meaning that instructions have to be particularly clear.

Liability, as outlined by the Blue Guide [25], is not exclusively held by the manufacturer: component-producers, distributors and sellers also hold a level of responsibility.

It is difficult to define how this legislation handles IT-driven business models. If manufacturers are now advertising goods with services, product add-ons and recommended products, they could be seen as instructing their customers to use their product in that way. However, it is not clear how a vulnerability in a third-party piece of software that provides an attack vector on the original product would be dealt with, leading back to our earlier discussion of culture.

One issue is that products that have been altered to change performance, purpose or type means that those

items can be viewed as new products: if a product add-on makes significant changes to the product, for example re-programming it, would the manufacturer no longer have any responsibility for the product to function safely? Would significant changes in system behaviour without re-programming be sufficient? If the in-car media application makes sufficient changes to the way the car runs, do the developers become entirely responsible for the car’s roadworthiness if the electrical systems fail?

Secondly, there is no liability if, at the time of manufacture, scientific knowledge was not good enough to predict the defect. This raises the question, if an internet-ready product is sold with no cyber security in the design, does the manufacturer avoid all liability? Is the requirement for ‘good engineering practice’ in the legislation sufficient to ensure that security is considered in the design irrespective of known vulnerabilities? For example, assume, in the climate-control system scenario, that the system is built using a component with no known vulnerabilities, but at a time when a competitor has a vulnerability in their system using a similar type of component from a different manufacturer. Would the two systems be viewed as sufficiently technologically different by a court for the manufacturer to say they couldn’t predict the defect?

5.1.2. UK consumer rights

As discussed in Section 4, the line between sale of products and service provision is becoming blurred by the growth of IT-driven business models. Within the UK, consumers also have legal rights covering both goods and services offered, via the Sale of Goods Act 1979¹¹ and the Supply of Goods and Services Act 1982.¹² Together, these state that:

- Goods should be: as described; of satisfactory quality; and fit for their purpose (taking into account their age, price and any claims made in adverts, leaflets or by the seller).
- Services should be carried out: as agreed; with reasonable care and skill; within a reasonable time; and for a reasonable charge — unless a price was agreed beforehand (taking into account the price paid and the way the seller offered the service).

If EU legislation fails to handle issues with software, terms of service may still protect consumers. If a mobile application is, in fact, a service, offering it still has to be carried out with reasonable care and skill. Just as a builder is expected to secure a site at the end of the day, application developers may be responsible for fixing known vulnerabilities in their software for the duration the service is offered for, and for not damaging the systems they connect to. If the refrigerator in our scenario malfunctions due to an attacker gaining

¹¹See <http://www.legislation.gov.uk/ukpga/1979/54>.

¹²See <http://www.legislation.gov.uk/ukpga/1982/29>.

access via a known vulnerability in the application provided by the manufacturer, and this leads to a fire, the manufacturer might be considered not to have put sufficient care into the service they provided. Obviously, as discussed earlier in this section, in order to be held responsible, one of the investigators would have to have considered the hazard of cyber threats in the system.

While companies have obligations to adhere to certain standards and consumers have legal rights, it is also worth making two comments. First, the consumer market for household appliances, etc. is traditionally far less accepting of faults than the IT market, possibly due to existing legislation. Second, the IT-driven business model is one of improving customer retention through continued interaction, and improving reliability through product partnerships. Decreasing the quality of the core product through faults in the value adding side-products would be counter-productive.

5.2. Standards

Cyber security issues are not a new phenomenon in the context of safety-critical systems. In hazardous environments, established standards have been updated to reflect new threats; for example, the MoD standard for safety management in defence systems now includes a section on cyber security and data integrity [5]. Contractors supplying products, services and/or systems have to consider cyber security in the context of safety where breaches, or (due to increased dependence on data) data integrity issues, may be a contributory cause of hazards or failure modes. They have to produce and implement suitable mitigations for issues discovered in the analysis. Hazards are usually only reported within the scope of the safety case, but cyber vulnerabilities are an exception — known vulnerabilities impacting only outside of the current safety case still have to be reported.

The BSI standard covering the functional safety of programmable electronic safety-related systems [27] also has a security section, requiring hazard analyses to include the possibility of “malevolent or unauthorised actions.” That means that, if our refrigerator had to meet these standards, the various new risks presented by an internet connection would have had to be evaluated in the design process. Risks identified would have to be mitigated, reducing the likelihood to an acceptable level. Such standards are typically used for very large and/or complex systems. While it is useful to refer to them to see the precedent for considering cyber security threats as a legitimate hazard in CPS, the safety-engineering processes defined for these systems come at an enormous premium in terms of both time and expense. The blanket use of these types of standards in the consumer market is unlikely to be seen as necessary or acceptable by stakeholders.

Security issues have been considered in safety standards for safety-related systems; however, the

overlap from security to safety is relatively inconsistent. ISO 27002 (information security management) [28] has been adapted through British standards to fit various industry sectors. The adaptation covering the energy utility sector, where there are large numbers of safety-critical systems, discusses cyber security through the integrity and availability of safety functions [29]. The standard has also been adapted for the telecommunications industry [30]. In this case, the focus is on the core of the network remaining available in the case of cyber attack. End-users and third parties are mentioned in the context of safety, but only in providing guidelines in case of emergency. There is no mention of the quality of the security measures implemented in home routers provided by the telecommunications industry, or the risks of appliances being connected to a small office or home (SOHO) router with no security.

There are several relevant standards or draft standards that, due to their cross-references and occasionally obscure applications, make it difficult to carry out a comprehensive gap analysis. They range from safety requirements for household appliances, to product interoperability in home networks and building controls. With one exception, the safety standards make no mention of internet connectivity and the IT-related standards talk about security (but not in the context of safety issues). The exception is the building automation and control systems standard, where “life safety messages” are earmarked for network priority by the communication protocol, allowing fire alarm messages, etc. to pass quickly through the building [31].

IT equipment, an example of a consumer product with both safety regulation and inherent cyber security issues, also has its own safety standard [32]. The standard discusses abnormal operation and fault conditions (a description broad enough to cover programming bugs), and states requirements intended to reduce hazards (fire, burns or shocks) and the likelihood of the equipment exceeding temperatures that would degrade components within the expected lifetime of the equipment (far shorter than for household appliances). While safety issues caused by a cyber security breach are not discussed as hazards in their own right, it can be assumed that a programming fault would usually be no more hazardous were the cause malware rather than poor programming skill. The most notable differences between these systems and the ones discussed in the scenarios are that these systems, even when networked, are self-contained in the context of the standard. They are also not being used to control and regulate appliances containing pressurised refrigeration liquids, machinery or other high-risk components.

Another example of a standard in a high-risk environment is AUTOSAR, which exists to standardise the basic software used in cars.¹³ Interestingly, it manages to prove the safety of its core software modules

¹³See <http://www.autosar.org>.

independently of any of the systems that may be plugged in later. By using standard interfaces, etc., the reuse of verified code is made easier, potentially increasing software quality across the whole vehicle.

5.3. Good safety engineering practice

Where standards fail to be prescriptive about how a design issue should be dealt with, the fallback position is that products should be designed following good engineering practice [25]. Good engineering practice in such a cross-disciplinary field is open to interpretation; however, texts such as those by Storey [33] and Leveson [34], which describe the traditional approach to safety in what has now become the field of cyber physical systems, provide excellent starting points.

Sometimes safety and security will be in conflict — an increase in safety might reduce security: for example, the heating control system might prioritise speed in transporting commands to the distributed controllers over maintaining their confidentiality. At other times, one aspect will be a pre-requisite of the other. Using the same example, to ensure that controllers receive commands within a pre-defined safe time, the network used by the system might need to have some security measures to protect against Denial of Service attacks. As we have seen, definitions of one can be given in terms of the other (or, indeed, the two definitions may overlap). (Further, it is worth noting that privacy brings an additional complication. The impact of existing privacy legislation on Consumer CPS design is discussed in Section 6.)

In the traditional approach to safety, products are ideally designed to be intrinsically safe — meaning that they can't produce enough energy to cause harm. An example of this is the safety requirements for IT equipment [32]. Where this isn't possible, a threshold of tolerable risk has to be defined, and measures put in place throughout the design and verification processes to ensure that potential hazards are not likely enough to reach this threshold [33]. This more complex process can be seen in the standards outlined for more safety-critical items, such as those for the automotive or aircraft industries. For software-related hazards, risk mitigation is most often carried out using controllers programmed to ensure a system defaults to a safe state. This means that, in the case of the in-car media application, if there is a problem with the way that the application is communicating with the car, then the car's electronic control units should default to a safe state. This assumes that the application isn't able to make sufficient changes to the on-board systems to override the safety measures put in place by the manufacturer.

Software faults are considered design issues, as algorithms do not degrade over time in the way that electrical components do. This means that software has to be rigorously validated and verified as part

of the design process. Software is often assessed for its reliability as a means of proving safety. Thus, in a traditional approach, if the refrigerator was programmed to be safe at the time of design the manufacturer would not consider potential changes to the software or any cyber security vulnerabilities when reviewing the system safety several years into its lifespan. The issue when considering software in this context is that the software can remain the same throughout the life of the product, but the environment, and the equivalent security level, change over time potentially making the system unsafe.

Leveson [35] shows how, in addition to these traditional approaches to safety, safety practices can be better aligned with systems engineering. Leveson argues that systems are becoming more complex, due to the reduction in physical constraints of electrical components, allowing designers to be more creative [35]. In developing a safety model that links with systems engineering, Leveson notes that reliability has little to do with safety — a piece of software may be reliable but unsafe if the developer has (accidentally) programmed it to be that way. As systems become more complex and contain more software, it becomes more difficult to ensure that the system will act as the designer intends. The event chains used to evaluate hazards are equally difficult to produce, as the more complex a system becomes, the less likely there is to be a single root cause for a given hazard. This is not an issue if considering cyber threats as this provides a clear cause; however, should a developer wish to pre-emptively evaluate the system for vulnerabilities the chain is not as obvious, as the link between collections of vulnerabilities and attack vectors can only be guessed.

The level of safety of a system can evolve over time, and, while the threshold looks at the product as a whole, the components are unlikely to all change at the same rate. If one part deteriorates or evolves while others don't change to accommodate that evolution, the system may become unsafe (due to *asynchronous evolution* [35]). As most systems are designed by multiple people, the interaction between components and their respective safety controls are the areas of a system that are of most concern. An example of this is the refrigerator scenario, with its web-site and multiple application developers. The overlap between areas of responsibility in these areas involves stakeholders from entirely different companies, many of whom are competitors, and where there is a large amount of pressure from the business to make it work.

The validation and verification processes can become difficult in a dynamic system — designers have to manage to either accredit the subsystem without knowing what will be attached, or specify exactly what the system will consist of. This can become complex, not only in terms of design, but also in terms of interaction and collaboration agreements between stakeholders. The SAVI Virtual Integration

Process [20], which is employed in the aircraft industry, specifically defines how designers can work on an architecture together without revealing sensitive intellectual property to other stakeholders.

One of the complexities in determining how to measure risks posed by cyber threats is the expected lifespan of the products being developed when compared to the lifespan of the average computer. Consumers would be unhappy if they were told that the electronic systems in their brand new car would go out of date within two years, making it no longer road-legal. Unfortunately, while legal levels of safety are expected to be maintained for the reasonable lifespan of the product [25], cyber threats evolve. That means that the traditional safety engineering view of software as something that cannot degrade over time no longer holds as soon as cyber security is considered a hazard.

Traditional approaches to dealing with errors, such as voting systems, will not be sufficient to provide safety if a hacker wants to override them using virtualised components. In the case of malicious actors, it may also be necessary to consider how effective non-programmable elements are in a system — does the system provide sufficient complexity for an attacker to find a side channel and avoid those controls altogether?

Hackers often attempt to avoid detection for as long as possible, meaning that changes they make may not be noticeable, or may manifest as an intermittent or transient fault. The complexity of the system would aid in hiding the source of the fault, as engineers employed by a manufacturer to mend faulty climate-control systems or refrigerators are unlikely to be experts in cyber security.

Finally, hazards found during analysis result in design constraints, which, in Leveson's model, require control algorithms [35]. It is difficult to see how the hazard of the controller being re-programmed to suit an attacker's needs might be handled in any of these models.

Safety engineering practice has a through-life approach, with a focus on ongoing operational and management processes of a live system. These elements of the process have largely been ignored in the consumer market, where all that is required are product safety checks at random intervals within a product's reasonable lifespan. While connecting these products may increase system complexity, it also allows the continuous communication between manufacturer, consumer and product, facilitating ongoing maintenance which may be required to successfully mitigate some of the emerging hazards.

Consumers also have an expectation of dependability in appliances. As more of these devices go on-line and safety requirements continue to require high quality software, the safety versus security consumer culture war begins: between the differing expectations for IT devices — which regularly fail — and household appliances — which typically function without fail for perhaps a decade.

6. OTHER OBLIGATIONS TO THE CONSUMER

The overlap between safety and security models tends (due to the weight given to safety requirements designed to reduce harm) to focus on the integrity and availability aspects of the security CIA (confidentiality – integrity – availability) triad, with little emphasis on confidentiality. However, for the sake of completeness of any model, all the facets of security need to be considered.

When moving from selling goods to supplying goods with services there is a second set of legislation that manufacturers are required to adhere to in their system designs — legislation pertaining to privacy when handling personal data.

The requirement to provide privacy to the end-users of Consumer CPS exchanging or storing personal data would result in a similar set of technical requirements to those of confidentiality as part of a small organisation's security policy. In terms of requirements elicitation, while a Consumer CPS design *should* consider the confidentiality of data for security purposes, it *must* consider adequately protecting user privacy in order for the manufacturer to avoid potential fines, making privacy the stronger motivator. It is therefore from existing privacy legislation that the general 'confidentiality' attributes are drawn.

Privacy legislation has been influenced by principles developed in 1980 by the OECD (the Organization for Economic Cooperation and Development) to support a pan-European data protection system:¹⁴

1. Notice: subjects should be given notice when their data is being collected
2. Purpose: data should only be used for the purpose stated and not for any other purposes
3. Consent: data should not be disclosed without the subject's consent
4. Security: collected data should be kept secure from any potential abuses
5. Disclosure: subjects should be informed as to who is collecting their data
6. Access: subjects should be allowed to access their data and make corrections to any inaccurate data
7. Accountability: subjects should have a method available to them to hold data collectors accountable for following the principles

While endorsing the OECD's recommendations, the USA has not implemented them directly; however, all of the seven principles were incorporated into the EU Data Protection Directive,¹⁵ which is an important aspect of EU privacy and human rights law, and regulates the processing of personal data within the European Union. It regulates the processing of personal data, defined as

¹⁴<http://oecdprivacy.org/>

¹⁵http://ec.europa.eu/justice/policies/privacy/docs/95-46-ce/dir1995-46_part1_en.pdf

“any information relating to an identified or identifiable natural person (‘data subject’); an identifiable person is one who can be identified, directly or indirectly, in particular by reference to an identification number or to one or more factors specific to his physical, physiological, mental, economic, cultural or social identity”

regardless of whether such processing is automated.

Many of the discussions focused on engineering privacy also use the OECD principles at the core of their discussions. For example, Langheinrich [36] presents six principles for guiding system design in the field of ubiquitous computing, pertaining to: notice; choice and consent; anonymity and pseudonymity; proximity and locality; adequate security; and access and recourse.

With respect to privacy, Roman *et al.* [37] distinguish between protection of personal information and what they characterise as “the existence of entities that profile and track users without their consent” [37].

Finally, users need to be able to obtain the essential household appliances, such as refrigerators, while maintaining the safety of their home and choosing not to share their data. This raises the question of the rights of a user should they ignore appliance instructions and choose to disconnect from the Internet for extended periods — do manufacturers have a responsibility to produce products which remain safe irrespective of the privacy preferences of the user?

7. EMERGING REQUIREMENTS

In this section, we aggregate the attributes discussed in the previous sections to establish a requirements set suitable to meet safety standards set by existing legislation. The intention is to consider the evolution in the way appliances are developed and marketed, as well as the new risks related to cyber security that the consumer faces.

7.1. Summary of Consumer CPS system attributes

7.1.1. Legal obligations

1. Products are expected to be intrinsically safe or have the risks reduced, with some new dimensions of products being viewed as services rather than goods.
2. Ongoing services require ongoing protection and connected products need threat assessment.
3. Emphasis on ‘instructions’ as part of the product, dictating what responsibilities a manufacturer has for consumer safety in various use cases.
4. Protection of personal data held and in transit.
5. Services supplied with goods should be carried out with reasonable care and skill.

7.1.2. Evolution of environment

6. Pervasive but not necessarily reliable connectivity, providing the environment for ubiquitous computing to develop, with huge numbers of devices to potentially go on-line.
7. Cyber threats as direct risks to safety or physical security in the consumer goods market.
 - (a) New risks associated with the business model.
 - (b) Continuously evolving threats.
8. Elements of the system carried with the user at all times.

7.1.3. Evolution of business drivers

9. Moving from selling goods to supplying goods and services in multi-stakeholder IT-driven business models, with complex manufacture/service provision models.
10. Service provision means an opportunity for ongoing interaction with the system.
11. Time to market reduced with the move from appliances to applications.
12. Very constrained budget.

7.1.4. Evolution of development team

13. Increase in the number of separate development teams for the systems.
14. Resistance of designers to adopt process-heavy safety requirements, especially where some producers don’t have the infrastructure to work on long, process-heavy projects.
15. Potential for designers to become over-burdened due to increases in complexity.

7.1.5. Evolution of user expectation

16. Expect products to be usable without prior knowledge or skill, ready ‘out-of-the-box’ and safe even when use is intuitive rather than as instructed.
17. Instructions have to be clear, short and simple, with user-friendly interfaces on the product.
18. Made available without noticeable price increases.
19. Expect products to be safe whether maintained on- or off-line.

7.1.6. Evolution of system

20. Small scale/budget complex systems — systems of subsystems, with associated synchronisation and validation issues.
21. Huge numbers of highly interconnected systems, with small degree of separation from safety-critical systems.
22. Safety- and security-critical systems — software security degrades over time, which in networked CPS constitutes a safety hazard.
23. Differing expected lifespans for subsystems — from months to tens of years — to be balanced.
24. The ability to repeatedly cause harm, accidentally

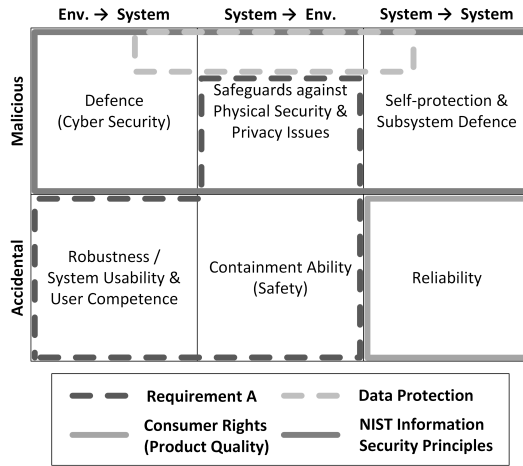


FIGURE 4. Use of the SEMA referential framework in the context of Consumer CPS.

or intentionally, through cyber attacks or subsystem software faults.

25. Evolution in the uses of subsystems, making them multi-purpose.

7.2. Requirements

We have argued throughout that a pragmatic approach is needed in order to encourage cyber security measures in Consumer CPS — that manufacturers need a clear financial motivation which the connection between cyber threats, health and safety, and product liability that the system attributes described in Section 7.1 provide. As we consider safety concerns to be the bigger motivator, it is with a safety definition that we begin our discussion on requirements.

The EU directive on general product safety defines a *safe product* as follows:¹⁶

“safe product’ shall mean any product which, under normal or reasonably foreseeable conditions of use including duration and, where applicable, putting into service, installation and maintenance requirements, does not present any risk or only the minimum risks compatible with the product’s use, considered to be acceptable and consistent with a high level of protection for the safety and health of persons, taking into account the following points in particular:

- i. the characteristics of the product, including its composition, packaging, instructions for assembly and, where applicable, for installation and maintenance;
- ii. the effect on other products, where it is reasonably foreseeable that it will be used

with other products;

- iii. the presentation of the product, the labelling, any warnings and instructions for its use and disposal and any other indication or information regarding the product;
- iv. the categories of consumers at risk when using the product, in particular children and the elderly.

The feasibility of obtaining higher levels of safety or the availability of other products presenting a lesser degree of risk shall not constitute grounds for considering a product to be ‘dangerous’.”

In the context of the system attributes, it is possible to use this definition as the catalyst for the first requirement:

- A. Unskilled users can safely use the system through any recognised interface. (*system attributes 1, 3, 16 and 17 also support this requirement*)

Using the thought processes applied by Piètre-Cambacédès and Chaudet in the SEMA referential framework [13], it is possible to consider this requirement in the context of the sub-notions of safety and security it covers (Figure 4). As demonstrated, the requirement is linked to three notions related to safety, but not to security. The addition of the other legal obligations discussed in Section 5 — Data Protection and Consumer Rights — increases the coverage so that all six notions are partially addressed, however there is a visible lack of coverage of the two main security notions, visually describing the lack of consideration of cyber security in the implementation of safety measures in Consumer CPS.

We have argued that the inclusion of cyber security is key in maintaining safety once the marketplace adapts to the pervasive availability of connectivity discussed in Section 2. In order to address these security concerns, we have selected a holistic framework for traditional information security, whose definition covers the three ‘malicious’ notions of the SEMA Referential framework. This framework makes no mention of safety, but provides a comprehensive set of “system-level security principles to be considered in the design, development and operation of an information system” [13].

The following requirements, based on the system attributes of Section 7.1, are aimed at bridging the gap between engineering practices when considering safety in information systems or cyber security in consumer products, in order to facilitate Requirement A:

Design & implementation:

- B. Cyber threats are recognised as a potential safety hazard in networked consumer cyber physical systems, in order to maintain tolerable risk thresholds. (*attributes 2, 7, 21 and 22*)

¹⁶<http://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX:32001L0095>

- C. New sources of harm introduced by the network or connected subsystems (including, for example, harm to economic well-being and privacy only typically addressed by security measures) are evaluated. (*attributes 4, 5, 8, 24 and 25*)

Through-life:

- D. Modular designs are accepted in safety conformity accreditation so that products are considered as part of a dynamic system including add-on products or services, rather than as an isolated appliance or as a pre-defined system. (*attributes 6, 9, 11, 12, 13, 14, 15, 18 and 19*)
- E. Security degradation over time is minimised to maintain safety. (*attributes 10, 20 and 23*)

8. MOVING FORWARD

The requirements of the previous section have the potential to prompt an overhaul of the manufacturers' product through-life process, as well as to prompt adaptations to the design to include security elements at the same time that a networking capability is added. Certainly, from the safety perspective one might assume that something akin to Requirements B and D will be included in standards as computing becomes pervasive, the ecosystem evolves, and more cyber threats in the Consumer CPS domain are reported and identified. However, the eventual reactive inclusion of elements of cyber security in safety standards is to the detriment of Requirement A, which is intended to reflect the sentiment of consumer product legislation.

In this section we discuss how a set of general cyber security principles may be mapped to guidelines on supplying safe consumer products, using the requirements from Section 7 to form a framework for safety through cyber security. The security principles are those published by NIST [17], which provide a holistic overview of cyber security best practice, although in some cases the wording is for security as a business process rather than as a product component — interpretation in this context is discussed in the following subsections. The safety guidelines are loosely structured by the safety engineering process described in ISO 10377:2013 [16] (to simplify the discussion both sets of information will be referred to as *advice*). In some areas the high level requirements suggested in Section 7 are insufficient to explain how the two engineering processes might interlink — in those areas the requirements have been expanded by sub-requirements to aid discussion.

Figures 7 and 8 illustrate the framework, with the following subsections putting the links between the two sets of advice into context, using the scenarios of Section 3 as motivation.

8.1. Commitment to providing safe consumer products

Both the cyber security and safety advice suggest that their concepts require a commitment from the organisation so that the concepts permeate company culture in addition to including both safety and security of a product during its lifecycle. This relates to the arguments of Sections 4 and 5 about the drivers for introducing new IT-driven capabilities and the breadth of roles that are involved — neither safety nor security can be thought of as an independent function, as both are intrinsic to the design and engineering process. Requirements A and B are used to bridge the gap between general safety and security advice; however, as well as representing the spirit of product safety legislation, Requirement A also covers the topic of instructions, which are intrinsic to safety legislation, but are not addressed in the cyber security principles. The closest mention of cyber security of instructions is the combination of policies and ease of use.

As discussed in Section 5, instructions play an important role in ensuring users know how to use the system, thus limiting a manufacturer's liability if a system causes an accident when used incorrectly.

There are some inherent issues in this approach when considering Consumer CPS — such as the number of stakeholders who own modules in the system — leading to either conflicting instructions or a lack of instructions where there is an overlap in responsibility. There are also some obvious issues in a distributed system with the number of instructions a user would be willing to read. The manufacturer has to be aware of this when designing their instruction set.

Instructions are dealt with differently in fields where the products on sale are appliances, with this difference being due primarily to the different outcomes of risk assessments. On the manufacturing side, instructions are provided so as to alert users to potential hazards and comply with safety regulations. These are seen as vital to the user's safety and so are written in an accessible manner, often using pictograms to ensure that language issues don't impede safety. Whichever format of instruction is provided, it needs to be appropriate for the user interface in question. This makes the question inherently difficult in the context of mobile phone apps, where users expect implicit understanding — as a consequence of intuitive interfaces — combined with trial and error to be appropriate.

Some suggestions for protecting users might be to direct them to the manufacturer's own web-site as a front page for downloading applications, before redirecting them to specific on-line application stores. This has the advantage of advertising all elements of the system on the same page forcing the user to consider them together in terms of safety instructions, as well as the obvious benefit of getting the consumer to continue to make contact with the manufacturer's web-

site. Another option might be for the manufacturer to distribute a safety warning pictogram to their partners with their APIs, requiring that a safety message be displayed to users as they access the system via an application they have downloaded.

In the context of the scenarios, all three could benefit from increasing the scope of their commitment to safety to include security, as well as developing instructions stating the importance of updates, discussing privacy and data use, password security, and most importantly the importance of only using apps recommended or accredited by the manufacturer.

8.2. Design and implementation

8.2.1. Hazard and risk analyses

In describing the inclusion of safety measures in the systems engineering design process, Leveson [35] suggests that hazard analysis should be considered as early as possible in the design process. This approach is facilitated by the IT-driven business model [23] of Section 4, as an element of this business model is the re-use of elements of a previous product along with new value-adding functions. For example, 95% of our refrigerator's parts could be identical to those of a previous model, with the biggest innovation in terms of its physical make-up being the inclusion of a wireless networking card. Because such a large percentage of the product remains the same, it should be possible to work from the existing hazard analysis when evaluating the impact of networking the appliance. As the designer doesn't start with a blank page, it is easier to consider safety from the beginning of the project.

In their study on the cross-fertilisation of safety and security models, Piètre-Cambacédès and Bouissou discuss the difference between security risks and safety hazards [4]. Both are present in our scenarios; however, a point made by Piètre-Cambacédès and Bouissou remains valid — that, while cyber threats and any associated risk calculations evolve over time, once identified, hazards remain relatively stable. The physical impact of a system on its environment when it has been adapted to become a Consumer CPS remains the same. If the climate-control system could malfunction, overheat and cause a fire in the old analysis, then it can still cause the same hypothetical fire in the new analysis. The key difference is that, in the new analysis, the trigger for an accident or intentional harm could be system malfunction, or a side-effect of a hacker re-programming the system.

Recognising malicious intent as a part of the system when evaluating potential hazards also means broadening the scope when measuring likelihood from the largely statistical measures used by safety models to that of a cyber security risk analysis, where likelihood considers motivation, resources and the difficulty of carrying out a particular attack [38] — in cyber security the likelihood of absolute harm is low, but the likelihood

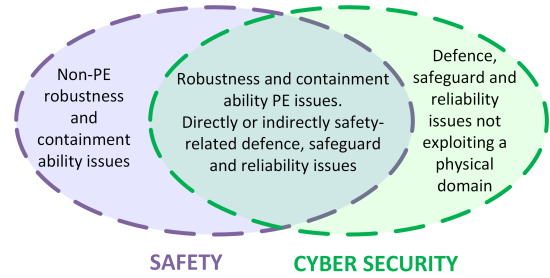


FIGURE 5. Safety hazards, cyber security risks, and their intersection.

of an incident repeating itself once a vulnerability is known increases.

Safety hazards cannot all be linked to cyber security risks: unless the appliance is being 3D-printed by the consumer, no amount of hacking is going to influence the sharpness of its corners once it has left the factory. Hazards that have no relation to programmable elements (PEs) won't need to be re-evaluated in the context of cyber threats, so part of the evaluation process needs to be the compilation of a subset of hazards relevant to cyber security threats, typically dependent on the PEs in a system and, where hazards are associated with PEs, their ability to communicate (in real-time or otherwise) with the network.

The SEMA referential framework, referred to in Section 7.2, can help differentiate between safety hazards, cyber security risks, and their intersection (see Figure 5). The intersection concerns robustness and containment ability issues within PEs, as well as directly or indirectly safety-related defence, safeguard and reliability issues. Indirect issues are those issues that arise as a consequence of data protection issues that can impact upon the physical security of a user.

Our three scenarios could benefit from the following being asked during the risk analysis process: can an existing hazard be produced remotely?; could any new data communicated become a risk to the user if leaked?; could an attacker's use of compromised controllers produce an unsafe level of resource contention within the system?; are there obvious attack vectors that could be deliberately used to cause harm?

8.2.2. Modular design

The framework also highlights modular design as key to linking safety and security advice, as well as being pragmatic in this type of IT-driven business model. Ward *et al.* [20] suggest evaluating systems using a diagram or architectural model, annotated by the various development domains when they discuss the integration of complex systems in aircraft development. While Consumer CPS are unlikely to ever be anywhere near as complex as (say) aircraft systems, they share some key attributes in the development cycle — a modular architecture with multiple stakeholders, some

of whom are competitors, and an expected system lifespan far longer than the average IT system.

Using a diagram to represent the system has another advantage when considering cyber security issues. The modular building blocks defined in the system become clear, because their scope and perimeters are defined. As cyber security is often essentially a process of creating barriers around entities that need protection, and safety issues often arise where subsystems overlap, knowing where the perimeters of the various system elements are is useful. For particularly important elements of the system, it may be necessary to employ multiple barriers at various levels of the system, using approaches such as defence-in-depth to ensure security.

Our three scenarios have different system architectures and pose different levels of risk. In Scenario 1, apps should be considered as modules with strict design criteria and API design. In Scenario 2, apps should be considered as modules with strict design criteria and API design, *and* a manufacturer-owned device or app registration process introduced to reduce likelihood of dangerous unofficial app use. The manufacturer may also wish to introduce more security at physical and logical interfaces to reduce their liability in instances of attack. Scenario 3 is intended to be a ‘rogue app’ example — the vehicle manufacturer may wish to consider the suggestions for Scenario 2 to make this business model less likely to succeed.

8.2.3. Design constraints

From this discussion of the architecture of a CPS, it is clear that there are areas where the different stakeholders have the ability to implement security measures and others where they have little control over the underlying system. Consumer expectation will be continued interoperability between the CPS and their other devices — that they can continue turning on their heating on the way home from work irrespective of their mobile phone’s operating system. Other modules will need to have sufficient security to meet the original hazard threshold despite these limitations, meaning that the main security requirement is pushed on to the appliance portion of the system where the computing constraints are likely to be the tightest.

Limited availability of computational power and communication bandwidth for cyber security is widely discussed in the context of larger scale CPS (see, for example, [39]). However, while the constraints on large-scale safety-critical CPS and their need to put availability and integrity above confidentiality do exist to some extent, in Consumer CPS there are other constraints that are less of an issue. First, while the design of a refrigerator may be used from one product to the next, there is not the same issue of making changes to an operational system — owners will not see enough value or have enough emotional attachment to their legacy appliances to feel the need to update

them with a networking capability. This means that designers can fully evaluate the changes they propose on the bench during the design phase, in a way that the designer of an Industrial Control System would only have the opportunity to do with a brand new multi-million pound system.

Another difference is in the higher level of flexibility that a Consumer CPS has to update and restart. Even if our refrigerator took 10 minutes to restart, there would be no damage to the food it contained and the consumer is unlikely to notice any change as refrigeration systems cycle on and off based on a thermostat. In other scenarios, like the in-car media application, the system is either on or off — when the application is actively connected to a vehicle, it could be considered switched on in the same way a washing machine would be mid-cycle. In such a circumstance, the updates could be planned to coincide with times that the system is not in active use.

In all of these scenarios, the manufacturer aims to make a profit through the sale of large numbers of small systems. The systems’ scale, the extent to which components and architectures are re-used, and the fact that liability is to some extent linked to the instructions for use, should assist manufacturers in being able to test updates against a comparable system before deployment.

Finally, depending on the CPS function, it should be possible to sandbox some new elements, separating them from safety-critical functions. Figure 6 shows possible high-level architectures for our scenarios. The refrigerator scenario has two designs, showing how it might be possible in some instances — in this example, where the system is being monitored but not controlled — to add duplicate hardware (in this case a second thermostat) which could completely isolate the refrigeration system from the consumer IT system. The cost may not initially look appealing, but would probably be cheaper than the through-life cost of maintaining a more hazardous system.

The climate-control system and in-car system are given as counter-examples. With respect to the former, the new system is designed primarily to control a system remotely and so can’t be isolated from the internal control functions. In the latter, the in-car system whose function is being changed is a black box from the point of view of the application developer — those responsible for the system cannot change the architecture of the electrical systems in the car to isolate their application from the safety-critical systems; rather, they have to hope that they don’t override any safety measures that the manufacturer has put in place. In these cases, careful design of APIs (limiting use cases) and perimeter security measures (limiting misuse cases) may be the best options. While in some cases physical elements of the system can’t be isolated, in others this represents a feasible and financially viable option for designers to ensure the safety of their systems.

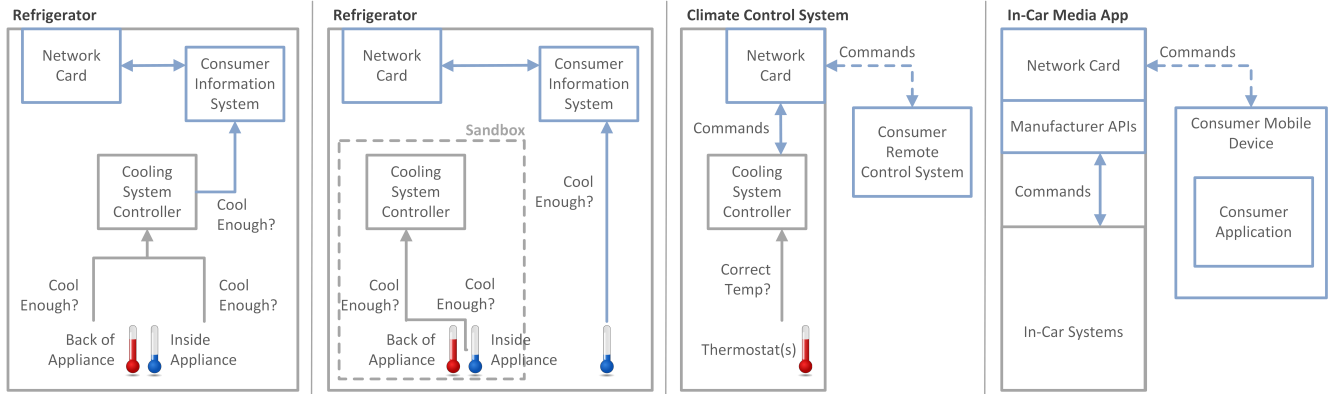


FIGURE 6. Scenario architectures.

8.3. Through-life

The requirements outlined in Section 7 necessitate an evolution to a more involved through-life process involving systematic security updates — either for system software or for the security measures themselves. By considering security requirements in the design phase, decisions can be made to limit the extent to which these are required, thereby reducing the ongoing costs of maintaining the safety of these systems.

It may be possible to push the responsibility for paying for security measures — up-to-date anti-virus software, etc. — on to the user, but, in the case of a serious safety issue requiring customers to make additional purchases to ensure their safety, this may not be sufficient to remove liability. For example, if a manufacturer was selling a refrigerator with sharp edges, it wouldn't be sufficient for users to be told to purchase a strong pair of gloves — there would be a requirement to change the design. However, if the manufacturer was selling an oven, then it is perfectly acceptable to suggest that the user purchase oven gloves. More safety-critical Consumer CPS typically have a far higher price tag than a standard kitchen appliance; as such, it may be the case that the way in which through-life security measures are applied and paid for should vary depending on the level of risk associated with the system. Scenario 1 probably needs occasional free updates, with a suggestion of endpoint security for devices running apps in the instructions. Scenario 2 is a more hazardous system that users have less direct contact with, meaning free automatic updates could be appropriate, with some built-in and maintained (excluding monitoring) security measures. Suggested use of endpoint security for devices running apps and network security inside the building included in the instructions could also be important. In Scenario 3 the car manufacturer might want to facilitate mandatory updates (where a warranty is voided if they are not carried out during a service) should an interface be found to be insecure.

9. CONCLUSIONS

We have drawn together the business drivers for, and the different attributes of, the Consumer CPS ecosystem, with a view to producing a set of requirements for the inclusion of cyber security measures in existing consumer safety models and legislation.

We began by reviewing the relationship between safety and security, before considering their convergence in the context of emerging ubiquitous computing models. These Consumer CPS were illustrated via three scenarios, all based around the same IT-driven business model. Reported attacks and the motivations for manufacturers to update their safety models to include cyber security aided in the collation of a set of Consumer CPS attributes. For the new business models to satisfy the original sentiment of safety legislation and reduce producers' liability, cyber security issues will need to be considered as part of the product lifecycle. While this is needed to maintain a level of safety equivalent to that currently experienced, there is a lack of clear legislation in this area, due to aspects of the new business model being perceived as service provision rather than a product. In other higher risk fields, such as defence, safety standards have already been adapted to recognise cyber threats (Def Stan 00-56 [5], for example), so it is not an unreasonable assumption that these types of measures will be required in other fields.

A set of cyber security requirements — derived from the set of high-level Consumer CPS system attributes — was developed, suggesting how safety might be enhanced through cyber security as consumer goods go on-line. These requirements were then discussed in more detail in the context of system architectures, as well as with respect to existing models for safety engineering in complex systems.

The framework of Section 8 is an initial attempt at linking safety and security processes in the context of Consumer CPS. Many of the new requirements can be

handled by the analysis of cyber threats as hazards and the inclusion of security measures at the design phase, aided by the fact that designers would usually be adapting old designs rather than starting with a blank page. For those that can't — driven by the fact that the security of software diminishes over time as the environment evolves — there is the possibility of the implementation of systematic checks and updates throughout the lifespan of the product. These updates are facilitated by the fact that the business case for the new computing model is one of maintaining contact with the customer over an extended period by providing additional services.

Just as the research underpinning this paper crosses disciplines, so does the future work for taking our framework forward. First, there is a need to build the business case for pre-emptively including cyber security in the development process by gaining a better understanding of how existing safety legislation will be interpreted in the context of cyber security in complex multi-stakeholder systems and the likely costs (in fines and reputational damage) of products lacking security and any perceived benefits of being able to cite security as a product feature. Second, there is a need to explore the inclusion of cyber security as a separate issue in safety standards for consumer products, so that where legislation may already be able to include cyber risk as a hazard it is more likely to be interpreted that way. Third, there is a need for further analysis of the use of existing security measures in the context of consumer CPS, as well as testing new methods for security by design in the different elements of these systems.

Developing these security measures requires future work adapting or amalgamating safety and security models, as well as developing new security tools. Any solutions also need testing against real systems as they are developed. These two sets of technical measures on their own are insufficient to mitigate against the safety hazards that cyber attacks pose. In order to provide comprehensive safety solutions, producers will need to find ways to communicate and collaborate with other members of the supplier ecosystem, ensuring that vulnerabilities don't fall between the zones of responsibility. Another area for potential work is that of user instruction sets. A user also has to be considered as part of the system — while systems become more complex, users become less knowledgeable about the way they work and are less able to evaluate the risks they are taking. Instructions are a key part of consumer safety legislation, but they are presented differently in the safety and IT fields. If manufacturers want users to understand the scope of the systems they are using and be aware of safety instructions, then those instructions have to be presented in an accessible format. Safety instructions in the application subsystems of a Consumer CPS made to look like typical software Terms and Conditions will not gain the attention of a user. Safety instructions become

even more important in the context of ecosystems where interfaces are standardised and producers have no partnership agreement. When the manufacturer loses control of the ways in which a system may be used, the only way they can reduce their liability is through clear instructions on how to use the system safely.

In conclusion, the domain of Consumer CPS is relatively new and has not yet developed to a point where a lack of security has become a serious issue; however, it is clear that substantial future work is required to ensure that safety levels are maintained in this emerging context.

Acknowledgements

The authors would like to thank the reviewers for their helpful comments. Emma Osborn's research is funded by EPSRC via the Centre for Doctoral Training in Cyber Security at the University of Oxford.

REFERENCES

- [1] Burns, A., McDermid, J. A., and Dobson, J. E. (1992) On the meaning of safety and security. *The Computer Journal*, **35**, 3–15.
- [2] Schneier, B. (1999) Attack trees. *Dr. Dobbs's Journal of Software Tools*, **24**, 21–29.
- [3] Rushby, J. (1989) Kernels for safety? In Anderson, T. (ed.), *Safe and Secure Computing Systems*, pp. 210–220. Blackwell Scientific, Oxford.
- [4] Piètre-Cambacédès, L. and Bouissou, M. (2013) Cross-fertilization between safety and security engineering. *Reliability Engineering & System Safety*, **110**, 110–126.
- [5] Def Stan 00-56 (2014) *Interim Defence Standard 00-56 Part 1*. Ministry of Defence. UK.
- [6] Lee, E. A. (2008) Cyber physical systems: Design challenges. *Proceedings of the 11th IEEE International Symposium on Object Oriented Real-Time Distributed Computing (ISORC 2008)*, Orlando, Florida, USA, 5–7 May, pp. 363–369. IEEE.
- [7] Wolf, T., Zink, M., and Nagurney, A. (2013) The cyber-physical marketplace: A framework for large-scale horizontal integration in distributed cyber-physical systems. *Proceedings of the 33rd IEEE International Conference on Distributed Computing Systems Workshops (ICDCSW 2013)*, Philadelphia, Pennsylvania, USA, 8–11 July, pp. 296–302. IEEE.
- [8] Oren, Y. and Keromytis, A. D. (2014) From the aether to the ethernet — attacking the Internet using broadcast digital television. *Proceedings of the 23rd USENIX Security Symposium (USENIX Security 2014)*, San Diego, California, USA, 20–22 August, pp. 353–368. USENIX Association.
- [9] Shafi, Q. (2012) Cyber physical systems security: A brief survey. *Proceedings of the 12th International Conference on Computational Science and Its Applications (ICCSA 2012)*, Salvador de Bahia, Brazil, 18–21 June, pp. 146–150. IEEE.
- [10] Dörnemann, K. and von Gernler, A. (2013) Cyber-gateways for securing critical infrastructures. *Security*

- in *Critical Infrastructures Today: Proceedings of International ETG-Congress 2013*, Berlin, Germany, 5–6 November, pp. 1–6. VDE.
- [11] Banerjee, A., Kandula, S., Mukherjee, T., and Gupta, S. K. S. (2012) BAND-AiDe: A tool for cyber-physical oriented analysis and design of body area networks and devices. *ACM Transactions on Embedded Computing Systems*, **11**, 49:1–49:29.
 - [12] Pfiztmann, A. (2004) Why safety and security should and will merge. In Heisel, M., Liggesmeyer, P., and Wittmann, S. (eds.), *Proceedings of the 23rd International Conference on Computer Safety, Reliability, and Security (SAFECOMP 2004)*, Potsdam, Germany, 21–24 September, Lecture Notes in Computer Science, **3219**, pp. 1–2. Springer.
 - [13] Piètre-Cambacédès, L. and Chaudet, C. (2010) The SEMA referential framework: Avoiding ambiguities when dealing with security and safety issues. *International Journal of Critical Infrastructure Protection*, **3**, 55–66.
 - [14] Simpson, A. C., Woodcock, J. C. P., and Davies, J. W. M. (1998) Safety through security. *Proceedings of the 9th International Workshop on Software Specification and Design (IWSSD 1998)*, Mie, Japan, 16–18 April, pp. 18–24. IEEE.
 - [15] Brostoff, S. and Sasse, M. A. (2001) Safe and sound: A safety-critical approach to security. *Proceedings of the 2001 New Security Paradigms Workshop (NSPW 2001)*, Cloudcroft, New Mexico, USA, 10–13 September, pp. 41–50. ACM.
 - [16] ISO 10377:2013 (2013) *Consumer product safety — Guidelines for suppliers*. International Organization for Standardization. Geneva, Switzerland.
 - [17] Stoneburner, G., Hayden, C., and Feringa, A. (2004) *Engineering Principles for Information Technology Security (A Baseline for Achieving Security)*, Revision A. National Institute of Standards and Technology, Washington, DC, USA.
 - [18] Rajkumar, R., Lee, I., Sha, L., and Stankovic, J. (2010) Cyber-physical systems: The next computing revolution. *Proceedings of the 47th ACM/IEEE Design Automation Conference (DAC 2010)*, New York, New York, USA, 13–18 June, pp. 731–736. ACM.
 - [19] Schneider, D. and Trapp, M. (2011) A safety engineering framework for open adaptive systems. *Proceedings of the 5th IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO 2011)*, Ann Arbor, Michigan, USA, 3–7 October, pp. 89–98. IEEE.
 - [20] Ward, D. T., Redman, D. A., and Lewis, B. A. (2013) An approach to integration of complex systems: The SAVI virtual integration process. *Proceedings of the 2013 ACM SIGAda Annual Conference on High Integrity Language Technology (HILT 2013)*, Pittsburgh, Pennsylvania, USA, 10–13 November, pp. 43–46. ACM.
 - [21] Serpanos, D. N. and Voyiatzis, A. G. (2013) Security challenges in embedded systems. *ACM Transactions on Embedded Computing Systems*, **12**, 66:1–66:10.
 - [22] Mitchell, R. and Chen, I.-R. (2014) A survey of intrusion detection techniques for cyber-physical systems. *ACM Computing Surveys*, **46**, 55:1–55:29.
 - [23] Kagermann, H., Osterle, H., and Jordan, J. M. (2010) *IT-driven Business Models: Global Case Studies in Transformation*. Wiley, Hoboken, New Jersey, USA.
 - [24] Yan, Y., Qian, Y., Sharif, H., and Tipper, D. (2013) A survey on smart grid communication infrastructures: Motivations, requirements and challenges. *IEEE Communications Surveys & Tutorials*, **15**, 5–20.
 - [25] The European Commission (2014). The ‘Blue Guide’ on the implementation of EU product rules 2014. <http://ec.europa.eu/DocsRoom/documents/4942>.
 - [26] Schär, S. G. and Krueger, H. (2000) Using new learning technologies with multimedia. *IEEE Multimedia*, **7**, 40–51.
 - [27] BS EN 61508-1:2010 (2010) *Functional safety of electrical / electronic / programmable electronic safety-related systems. General requirements*. British Standards Institute. London, UK.
 - [28] ISO/IEC 27002:2013 (2013) *Information technology — Security techniques — Code of practice for information security controls*. International Standards Organisation. Geneva, Switzerland.
 - [29] PD ISO/IEC TR 27019:2013 (2013) *Information technology – Security techniques – Information security management guidelines based on ISO/IEC 27002 for process control systems specific to the energy utility industry*. British Standards Institute. London, UK.
 - [30] BS ISO/IEC 27011:2008 (2008) *Information technology – Security techniques – Information security management guidelines for telecommunications organizations based on ISO/IEC 27002*. British Standards Institute. London, UK.
 - [31] BS EN ISO 16484-5:2012 (2012) *Building automation and control systems*. British Standards Institute. London, UK.
 - [32] BS EN 60950-1:2006+A2:2013 (2013) *Information technology equipment – Safety – General requirements*. British Standards Institute. London, UK.
 - [33] Storey, N. (1996) *Safety-Critical Computer Systems*. Addison-Wesley, Harlow, Essex, UK.
 - [34] Leveson, N. G. (1995) *Safeware: System Safety and Computers*. ACM, New York, New York, USA.
 - [35] Leveson, N. G. (2011) *Engineering a Safer World: Systems Thinking Applied to Safety*. MIT Press, Cambridge, Massachusetts, USA.
 - [36] Langheinrich, M. (2001) Privacy by design — principles of privacy-aware ubiquitous systems. In Abowd, G. D., Brumitt, B., and Shafer, S. (eds.), *Proceedings of Ubicomp 2001: Ubiquitous Computing*, Atlanta, Georgia, USA, 30 September 30 – 2 October 2, Lecture Notes in Computer Science, **2201**, pp. 273–291. Springer.
 - [37] Roman, R., Zhou, J., and Lopez, J. (2013) On the features and challenges of security and privacy in distributed internet of things. *Computer Networks*, **57**, 2266–2279.
 - [38] Mayer, N. (2009) Model-based management of Information System security risk. PhD thesis Computer Science, University of Namur.
 - [39] Kuipers, D. and Fabro, M. (2006) Control systems cyber security: Defense in depth strategies. Technical report INL/EXT-06-11478. Idaho National Laboratory, Idaho Falls, Idaho, USA.

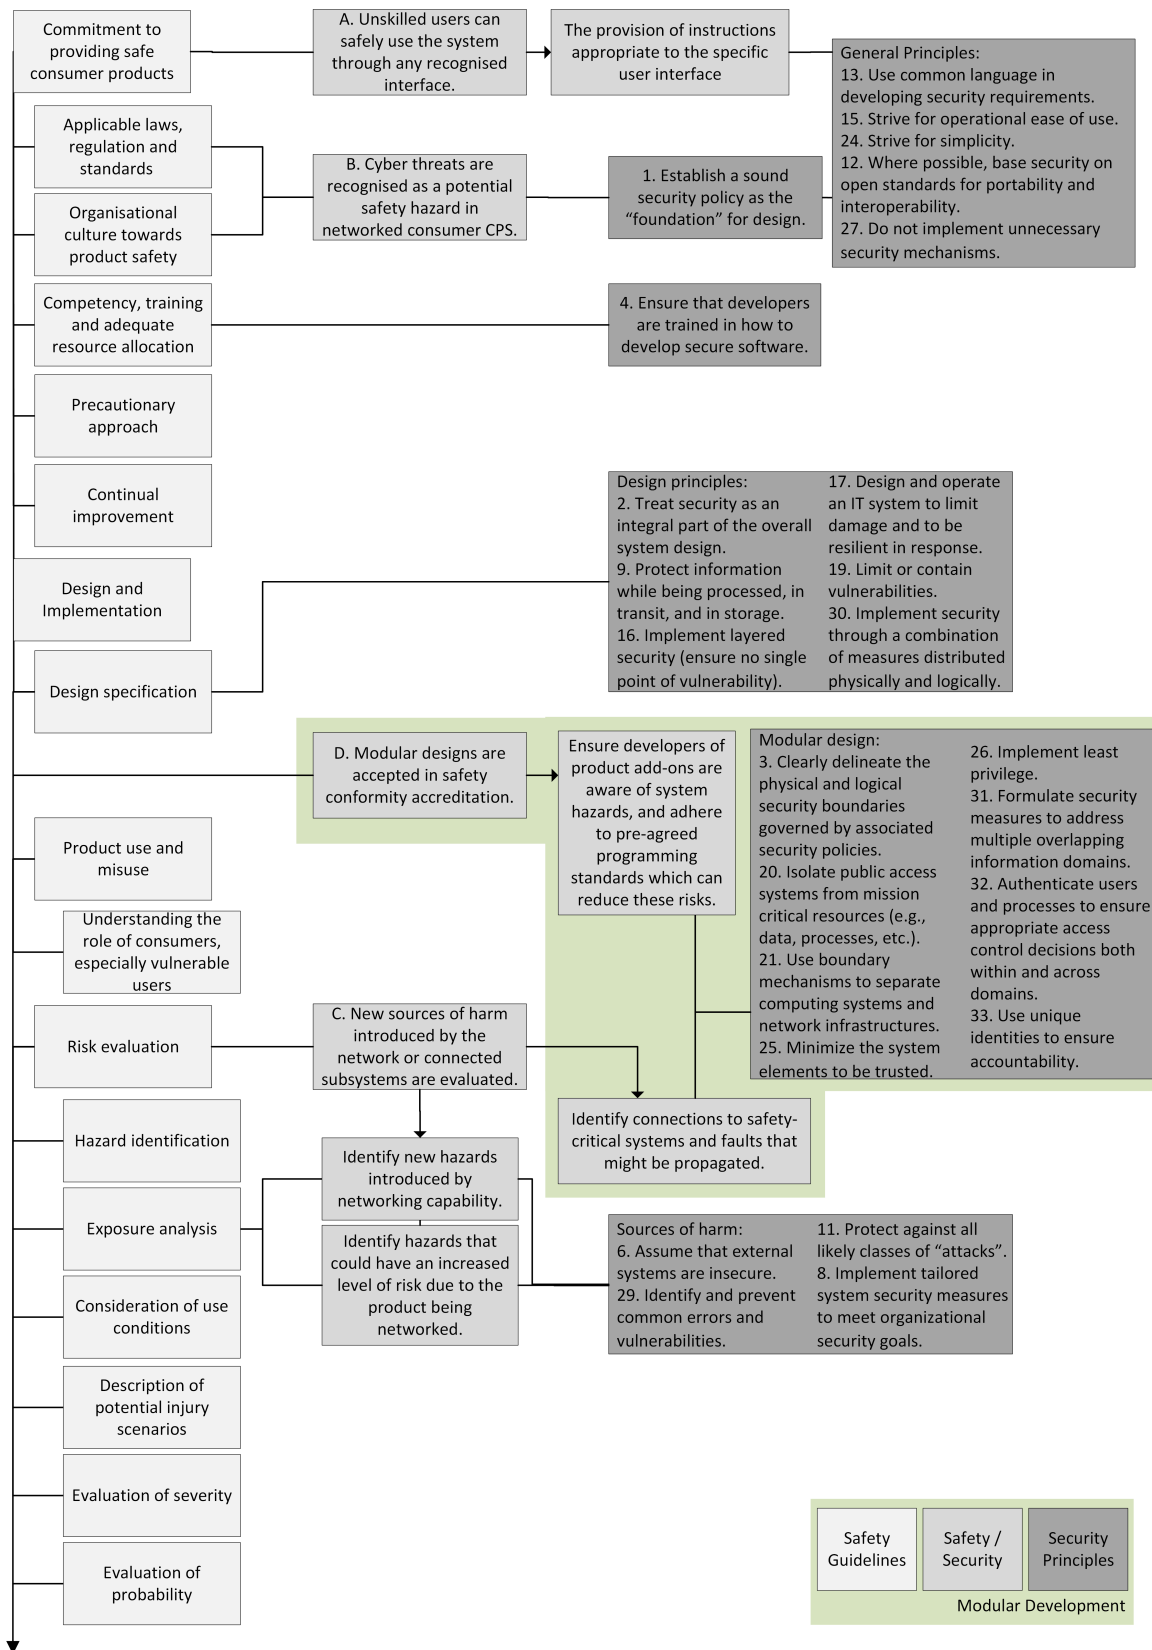


FIGURE 7. Safety through cyber security framework (part 1).

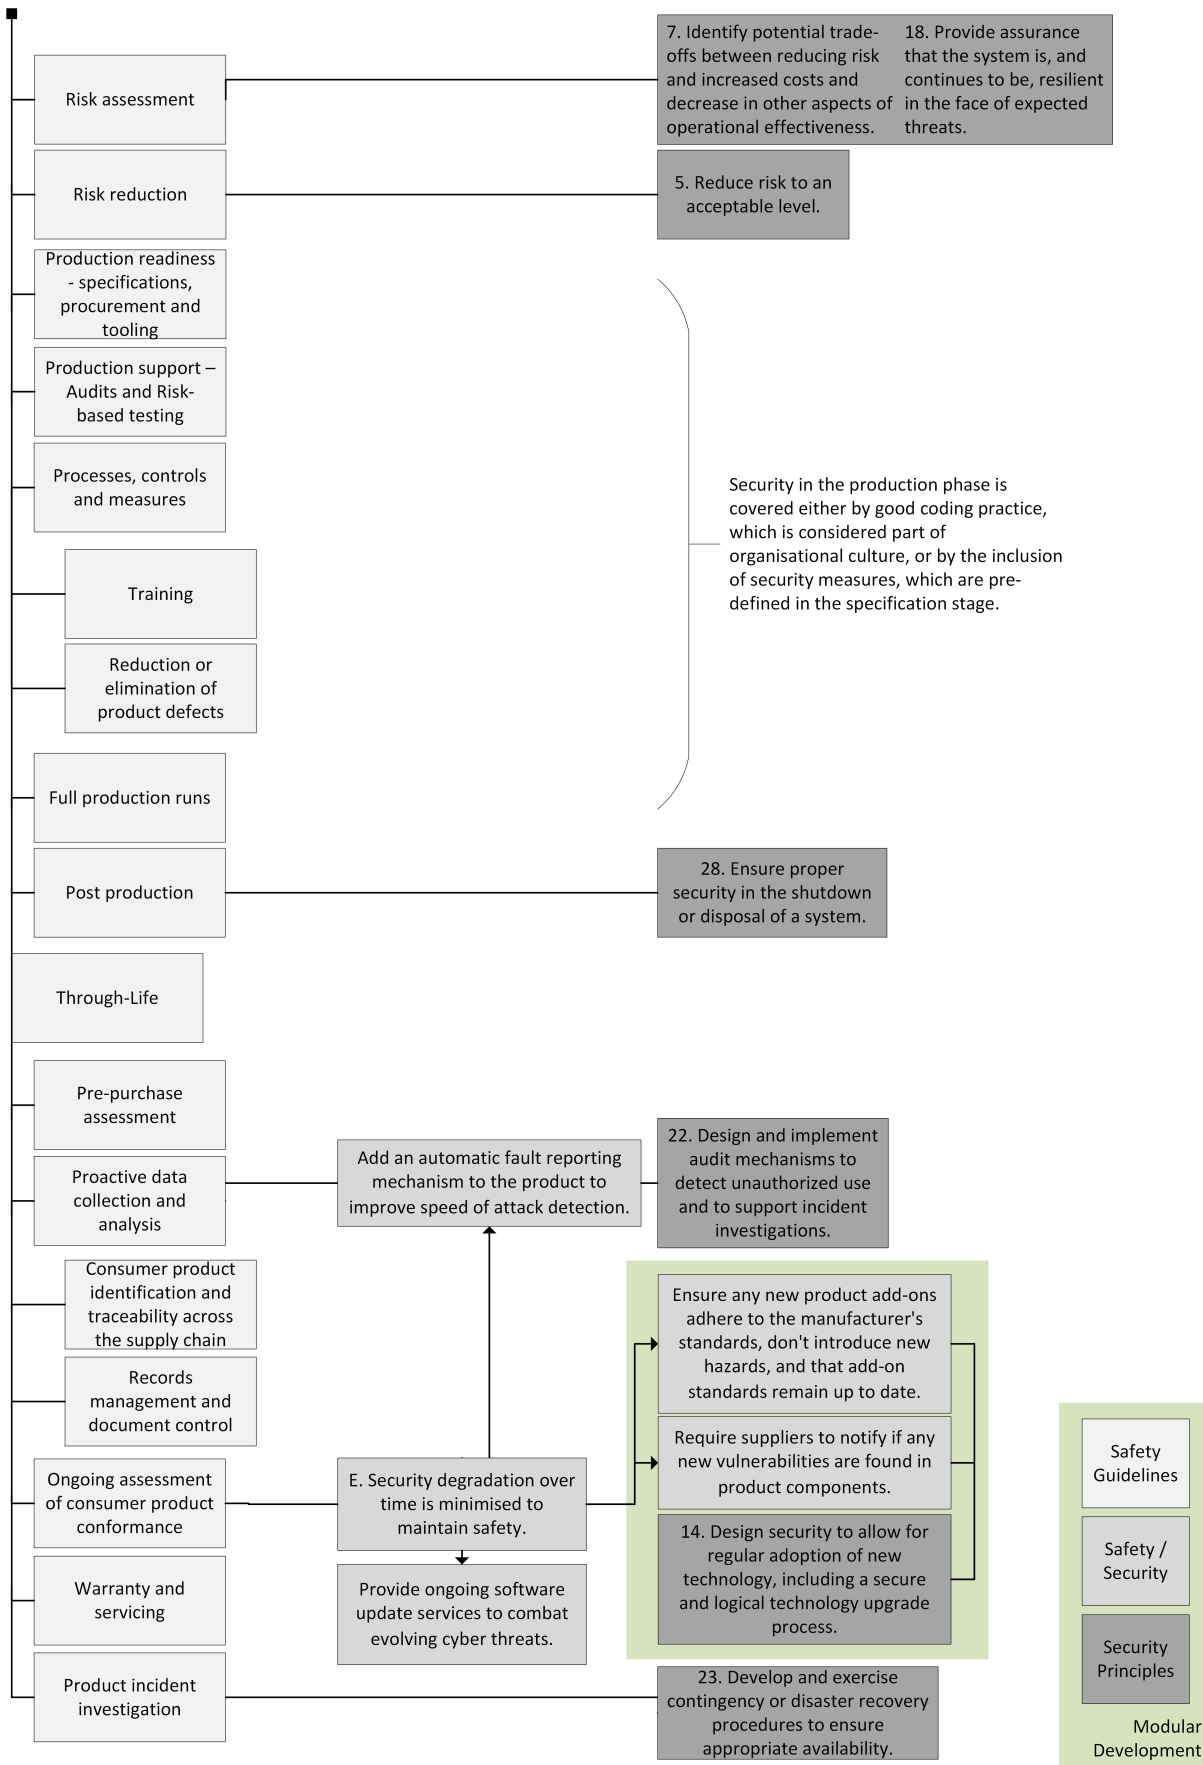


FIGURE 8. Safety through cyber security framework (part 2).