
CENTRE *for* DOCTORAL TRAINING *in*
CYBER SECURITY

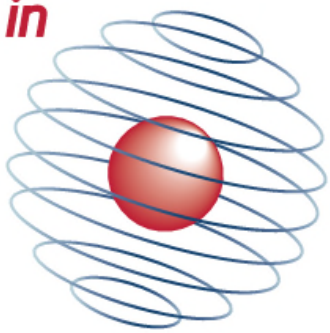

**CDT Technical Paper**

**02/15**

# Safety to security in emerging ubiquitous computing models

**Emma Osborn**

# Safety to security in emerging ubiquitous computing models

**Emma Osborn**
University of Oxford
Centre for Doctoral Training in Cyber Security
emma.osborn@cybersecurity.ox.ac.uk

June 2014

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

*The fields of safety and security have often intersected in the past, and are increasingly converging due to the rise in system interconnectivity, automation and dependence on the internet as part of critical national infrastructure. Research up until this point has mainly focused on safety in the context of large-scale industry, but with the emergence of ubiquitous computing models consumer cyber physical systems (CPS) are entering the home, and attacks on these systems are now being reported. This study evaluates the motivations for industry to implement consumer CPS and why they may want to include new safety and/or security measures. The evaluation of the ecosystem driving consumer CPS development has been used to establish cyber security requirements for this domain, with the aim of these requirements being to maintain safety levels in consumer CPS irrespective of new cyber risks having been introduced. Finally, the interpretation and implementation of the requirements for future work is discussed.*

## Section I - Introduction

The fields of safety and security are converging, whether intentionally through adaptation of models in one discipline to include aspects of the other [1], or through the inclusion of aspects of cyber security in safety models to meet an on-going safety requirement [2]. This convergence is being propelled by increases in system interconnectivity, automation, and remote access as well as our increased dependence on the Internet as part of our critical national infrastructure (CNI).

Research into this cross-fertilisation has largely focussed on the fields of nuclear power and industrial control systems, as a result of the changing risks to Cyber Physical Systems (CPS) and increased government funding post-Stuxnet. However, with the emergence of ubiquitous computing models, CPS of a different scale are now entering offices, public spaces and the home [3], and attacks on these systems are now being reported [4], [5].

The focus of this study is to assess how the safety of consumer CPS can be improved through the inclusion of cyber security measures, learning from adaptations made to safety and security models in the industrial sector.

The study is scenario-based with the scenarios being pragmatic, describing systems which have been adapted to include elements profiting from connectivity rather than new systems designed for purpose. The intention is to reflect probable gradual evolution in system design taking into account markets that are driving the evolution. Throughout the study, an attempt has been made to pull together opinions from diverse sources to provide an accurate account of the properties of these systems and draw up a set of general requirements which a combined safety and security model might have to operate within.

Section II develops the scenarios which are then used as concrete examples in Section III, where the attributes and market drivers for the development of consumer CPS are discussed. The scenarios are also used in Section IV, where existing safety legislation is discussed. In Section V the issues discussed in the previous sections are summarised as a set of system properties which can be used to begin building a set of requirements. Finally in Section VI the interpretation and implementation of those requirements for future work is discussed, before drawing conclusions and discussing future work in Section VII.

The use of cyber physical systems has been discussed under various titles including ubiquitous or pervasive computing and the Internet of Things (IoT). Many attempts have been made to define how ubiquitous computing might develop. An example of this can be seen in Beecham Research's Internet of Things Sector Map [6].

From industry representations of the IoT a number of common dimensions or service sectors where ubiqui-

| IoT Sector | | Consumer Location | | | | |
|---|---|---|---|---|---|---|
| | | *Home* | *Work* | *Travel* | *Retail* | *Social* |
| | *Buildings* | ● | ● | | ● | ● |
| | *Energy & Water* | ● | ● | | ● | ● |
| | *Consumer & Home* | ● | ● | | ● | ● |
| | *Healthcare & Life Science* | ● | ● | ● | ● | ● |
| | *Industrial & Agriculture* | | | | | |
| | *Transportation* | | | ● | | |
| | *Retail* | | | ● | ● | ● |
| | *Security & Public Safety* | | | ● | ● | ● |
| | *IT & Networks* | ● | ● | ● | ● | ● |

**Table 1** Consumer Contact with IoT Sectors

tous computing may have an impact can be collated. These include, among others: buildings, healthcare, consumer products, industry, agriculture, and transport.

Table 1 shows an aggregation of common service sectors discussed in IoT literature with a breakdown of where a consumer might come into contact with these sectors, considering the locations a person might visit in a typical day in the life: home, work, transport systems, retail outlets, and social spaces such as restaurants and bars.

Table 1 demonstrates that the only commonly defined sector for ubiquitous computing that a consumer does not come into direct contact with is the Industrial & Agriculture sector. This sector, along with the non public-facing aspects of energy and water production are those most commonly researched when considering the synergy between safety and security [1], [7], [8].

IT and healthcare sectors connect with every activity space in the scenario. In the case of healthcare the reason for this is obvious — the consumer carries their healthcare devices with them in what has been termed the Body Area Network (BAN) [9]. In the case of IT the reason is because whilst the IoT is still an emerging phenomenon increases in the 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 drawn between IoT sectors is likely to blur as new services themselves provide new opportunities for services. An example of this which can be seen in suggested implementations of the IoT is the overlap of the IT and healthcare sectors, where smart phones 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 the network for use by other sectors' products.

This study discusses topics which are cross-disciplinary by nature, as such there are many conflicting definitions of safety and security under consideration. In the interests of clarity this study uses the definitions described in the SEMA Referential Framework, which are the result of the various definitions being discussed in depth [10]. The definitions focus on the origin and nature of the risk as follows:

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

The environment in these definitions is *"the set of other interacting systems whose behaviour and characteristics are generally less known and beyond the control of the system owner"* [10]. In the context of this study this definition is assumed to include the physical environment and underlying communications systems, although consumer CPSs are assumed to be complex systems in their own right.

Figure 2 shows how these definitions are interpreted in the context of the scenario used in this study. 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 environments due to the increased interconnectivity of systems inherent in ubiquitous computing.

When discussing the cross-fertilisation of safety and security engineering, Pièrre–Cambacédès and Bouissou

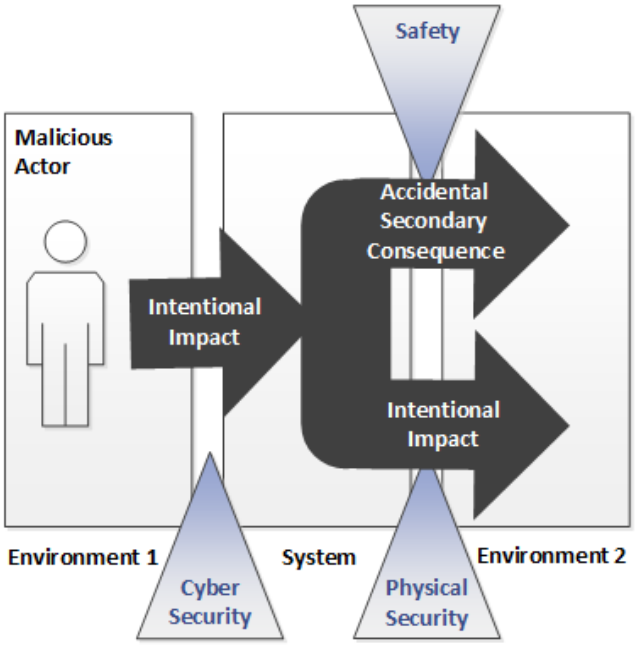


**Figure 1** Use of the Definitions of Safety and Security

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 [1]. This approach, whilst effective in the context of CNI, is likely infeasible when considering consumer CPS due to their number.

Although there are many definitions of Cyber Physical Systems Lee is often cited with his concise definition: *"Cyber-Physical Systems (CPS) are integrations of computation with physical processes."* [11]

The scenario used in this study considers CPS in the context of the consumer market rather than in an industrial setting, and considers system scope to include more than an isolated appliance. For the sake of clarity, this study calls them consumer CPS. The consumer CPS referred to here are consumer products containing embedded systems, which take advantage of the increase in available wireless technologies to connect to the Internet for reasons described in a business model described in Section III.2.

## Section II –Scenarios

For the purpose of this study three examples taken from different business sectors have been developed. This should aid in identifying how overlaps in safety and security models can best be exploited. The first two are the type of system frequently used in Internet of Things scenarios; the third is a slightly altered scenario — a consumer system similar to the first two: a subsystem touching the physical world, plus applications from different stakeholders interfacing with the physical system remotely. However, there is no buy–in from the manufacturer of the physical aspects of the system for this type of subsystem to be used in the way the application developers envisage, raising different questions when considering business models and liability.

*Refrigerator*

Jen has a state of the art fridge. When delivered the technician connected it to Jen's home Wi–Fi network. The technician set up an account for her on the company's website and told her she should log in to finish registering and look at the services she could use. The site gave her proof of purchase for the warranty, asked her if she wanted the fridge to automatically update its systems and offered a remote monitoring app plus some third party services. Through the recommended services Jen was able to link the fridge's inbuilt food recognition system to her supermarket shopping app, allowing her shopping list to auto–update when she was running out of her favourite foods. The monitoring app would also provide information about the food she had in her fridge, alerting her when something is about to go out of date or the fridge isn't keeping the food cool enough. Jen also links in an app she finds in the app store that automatically uploads information from her fridge into a diet diary so that she can keep an eye on the calories she consumes.

*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 a dog–eared piece of paper containing the details for logging on to the manufacturer's web page. By logging in Jen was able to change the climate control settings room–by–room to suit her, and download an app that would let her turn on the heating in her house when she was on the way home. The app gave her the option of linking to her car's satellite navigation system so that when she selected the 'home' option her climate control system could automatically detect when Jen was close to home and turn itself on. As well as linking to apps and the company website, the heating system links to the building's fire alarm system periodically providing health reports to help the alarm system detect faults before they become critical.

*In–Car Media*

Jen has just bought an Android app advertised online as being compatible with her car, that once connected

allows her tablet or phone to connect to the in–car system via Wi–Fi. Previous iterations of this product were sold as hardware to wire in to a car, but advances in the types of media car manufacturers expect to be connected mean that this can now be run direct from a mobile phone or tablet. Among other things it lets her have access to all her music through the dashboard screens. She finds a video online explaining how to install it and once installed changes the dashboard display to a prettier theme, linking her accounts so that her sat–nav can automatically look in her contacts to find the addresses she's travelling to, and letting her access music and videos through the in–car screens. The app also provides options for optimising engine control settings to the user's driving style, which Jen leaves for another day.

## Section III –Consumer CPS Attributes and Development Drivers

### III.1 Consumer CPS Research

Cyber physical systems combine the physical (continuous) world with the digital (discrete) one [12]. Until now, development of distributed CPS has mainly occurred with large and difficult to build safety critical systems, however consumer CPS are a type of highly distributed CPS. Whilst the size of the system in terms of number of lines of code or budget are enormously reduced the complexity is not reduced. Any reduction in the number of nodes is offset by the added complexity of many of the system elements being independently produced and with the scope of control a system designer has over the system being greatly reduced. An example of this is the refrigerator scenario, where the product designer can in theory control which supermarkets have access to the API for the fridge. The drawback is that even if the designer feels that a specific supermarket's app introduces too many vulnerabilities for them to be allowed access, the consumer will expect the service to be available for largest supermarket chains. In this case the designer's perspective is likely to be overruled by the company's marketing and strategy teams leading to a system with a weakness the designer can't change.

Wolf et al. discuss the fact that constraints present in traditional CPS for real time communication and safety are less present in consumer systems. In their view this opens the market for the widespread sale of data, actuator or computation services [3]. The inference is that in being less bound by safety regulation, consumer CPS designers can produce the innovative products which would lead to ubiquitous computing far more rapidly than safety critical systems have been adopted in the past.

Consumer CPS require both dynamic topologies and dynamic reconfiguration, whilst being potentially tool neutral with modules not necessarily being combined until runtime [13], [14]. These types of products are also likely to be developed in part by small companies contracted by a manufacturer. These small companies are competitive because they manage to keep the overheads of process to a minimum, meaning that introduction of more rigorous design processes is likely to be resisted.

The number of products such as household appliances with embedded systems which could potentially go online are far more numerous than non–embedded systems currently are [15]. This means that the impact of a small safety issue, when multiplied by the potential number could become extremely 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 extremely small [15]. An example of this is with the in–car media scenario. The acceptable price of an app is extremely low, but the app is designed to link into the safety–critical system inside of a car. In theory, that should mean that the app combined with the car is a new safety–critical system and the app developed accordingly; however, even if this issue was explained to the consumer they would still not be willing to pay the price of the car again to cover the cost of application development.

Much of the research done on CPS that reach the home is in the context of Smart Grid [8], [12], [16], [17]. In that scenario, the security focus is on protecting electricity company assets in a hostile environment— the consumer's home. In this study, the focus is on consumer safety and security, and so 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 safely designed few are likely to evaluate how cyber threats emanating from the way they choose to use that product may impact their safety. In the climate control scenario, when Jen logs into the manufacturer website she is unlikely to ask herself how many past tenants still have access to her heating system.

### III.2 Consumer CPS Development Drivers

As the scenario discussed in Section III demonstrates, whilst it's possible to define what a Consumer CPS is, what they interact with and the risk they pose varies enormously from example to example.

In order to find the underlying patterns of similarities in these products, it is necessary to use a far higher level of abstraction. Figure 3 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 3.(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. suggest that over time this business model is becoming less and less profitable, leading to IT–Driven Business Models [18]. Almost every interaction we now have in shops, restaurants, at work and through advertising includes a link to a company's online presence. At every opportunity businesses are attempting to open a continued dialogue with their customers whether through the official website or via a social media presence.

These models are shaped by customer value. Manufacturers make money by reducing the cognitive workload of the consumer as the customer system can be pre–built and configured (Figure 3.(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 them to sell their products 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. In the case of the refrigerator scenario, being able to download an app and link supermarkets to their fridge 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 hassle.

All of these services depend on the product going online — interconnectivity by design.

The issue with this business model is that it is drawing manufacturers away from their area of expertise. New

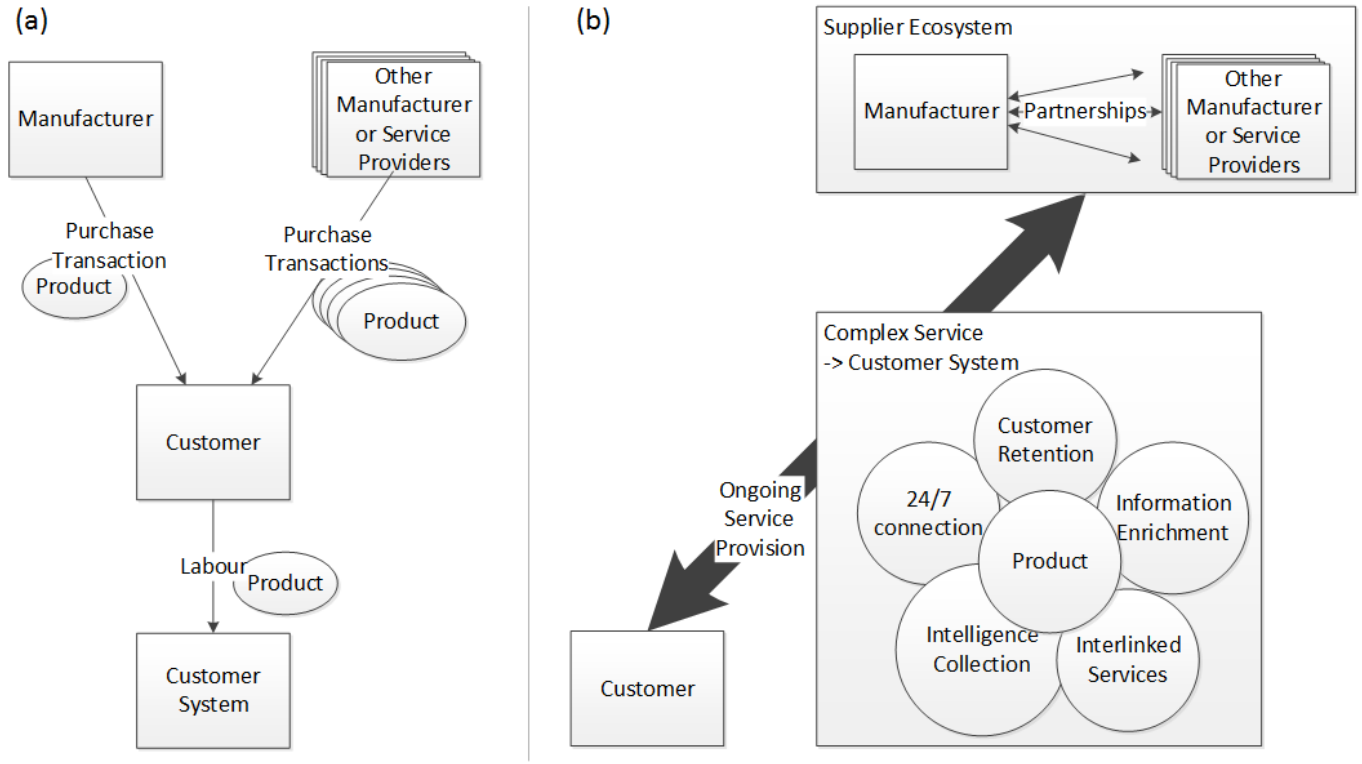


**Figure 2** (a) Traditional Manufacturer Business Model; (b) IT–Driven Business Model

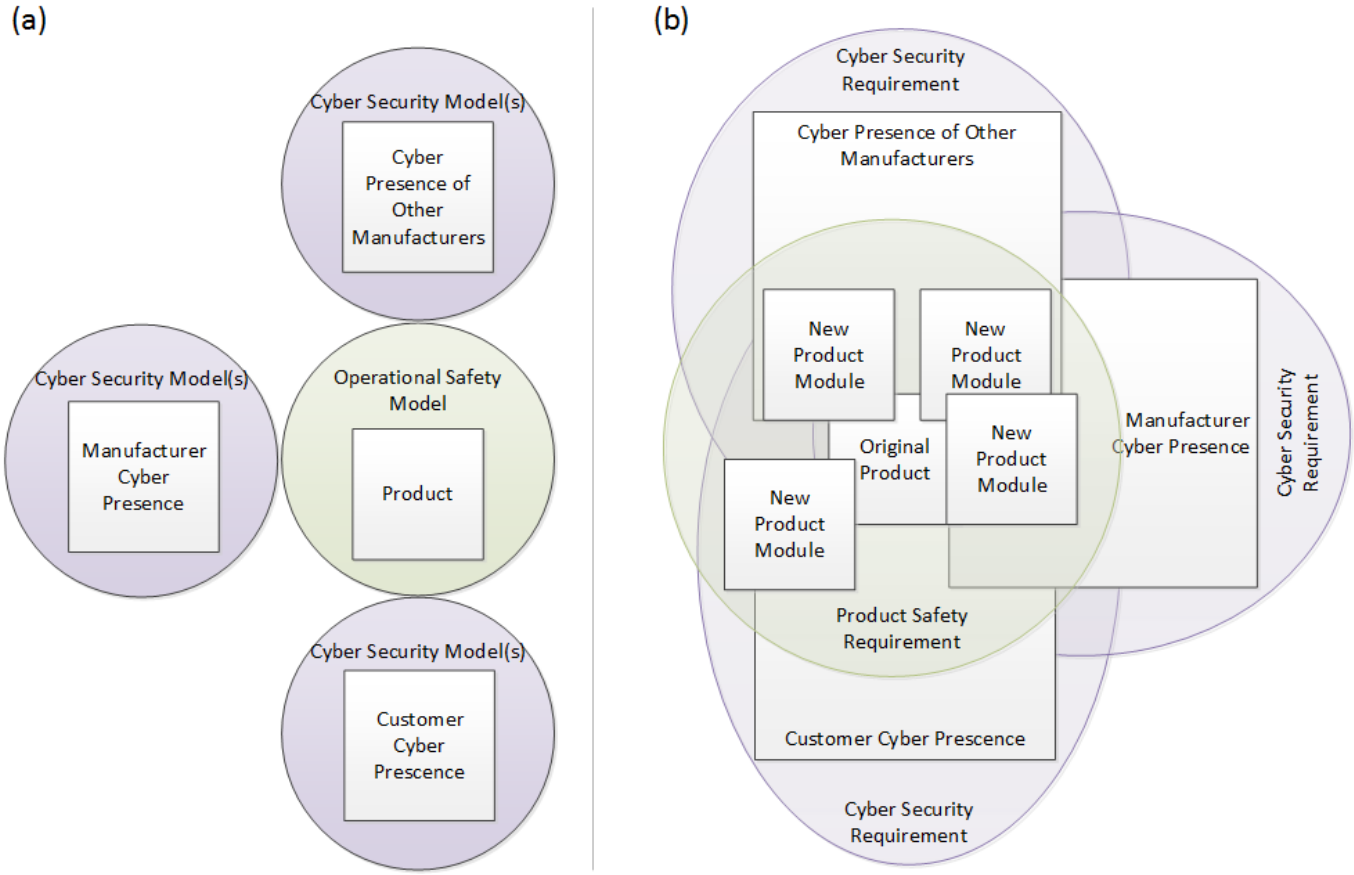


**Figure 3** Scope of Safety and Security Models in (a) Traditional Manufacturer Business Model; (b) IT–Driven Business Model

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 apps, web portals, etc. to third parties. Whilst modular design may help manufacturers when developing a complex system, this modularity of both the product and development teams makes it far more difficult to judge the risks associated with the use of the system once connected.

The evolution of the business model in terms of safety and security requirements can be seen in Figure 4, 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 perfectly 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 new services they are offering and the cyber security of the platforms on which they are hosted. As illustrated by Figure 2, by connecting a Consumer CPS to the Internet, new risks are introduced in the safety and physical security domains. Whilst 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. Developments in safety or security models should aim to address this issue in order to be adopted in the future.

## III.3 New Sources of Risk

With the implementation of new services in the consumer CPS market the types of risks associated with products can be seen to evolve to include new hazards.

*Remote Operation of Appliances*

One of the reasons 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 even if there are now objects in the way which

may 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, for example if malfunction leads to a fire the fact that there is no one present in the building may mean that it takes far longer for the fire to be detected and put out.

*Data Integrity*
Non–networked appliances are unlikely to be programmed with cyber security in mind — they are unlikely to test signals for authenticity as those signals will probably originate from a component fixed to the same circuit board, in the case of safety critical functions they may do limited tests for accuracy. Once an appliance goes online, with control signals coming from mobile phone apps or via a website, the designer can no longer consider the appliance as a closed system. If the appliance doesn't have a closed system the integrity of data received by the controller might be compromised more easily. It might be very easy for a neighbour to use the link Jen's fridge has with her supermarket account to sneak in 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.

*Third party's motivation to use resource*
Proofpoint's incident report of January 2014 [4] was related to the use of various devices as resources to use to send spam. The motivation of third parties to enumerate devices and transform them into resources introduces an entirely new field of risk to the household appliance sector, a direct consequence of increased interconnectivity. A fridge 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 might be quite 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 fridge and her dieting and supermarket apps would give an attacker information about what she ate, but also if there were other people living with her, highlight some health issues and bad habits she might not want people to know about. Her heating app and the in–car media app are both linked in to the car's satellite navigation program, so could provide an attacker with her physical location either putting her in danger or proving to a burglar that she is not on her way home.

*Malfunction*
The increased complexity of the system and the potential number of stakeholders makes it far more challenging to identify all of the flaws. If system verification becomes harder to achieve, there is more chance of malfunction leading to injury, damage of possessions or the product self–destructing.

UK fire statistics show that malfunctioning electrical appliances are a fairly common cause of fires in the home [19]. The increase of risk of malfunction in a CPS due to faults or malicious actors is likely to increase the risk of fire.

# Section IV –The Obligation to Provide Safety

## IV.1 Consumer Safety Legislation

The motivation for creating better safety models in the consumer sector are not the same as those provided in the industrial sector. Huge amounts of effort are being put into securing supervisory control and data acquisition (SCADA) and other industrial control systems. The motivation for this is clear — there are known attacks proving that these systems are extremely vulnerable, for example, searching the Common Vulnerabilities and Exposures database for SCADA produces 488 results [20]. 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 perceived as 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 so many do not get reported, even with a reporting system in place. In this circumstance a manufacturer is unlikely to incur large financial or reputational loss 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 as in the example given in Section I [4], are faults likely to be reported.

In the refrigerator scenario if an attacker had managed to reprogram a controller to help mine Bitcoins the

fridge might then be too occupied with the attacker's task to turn the refrigeration system on and off at the correct intervals — either it freezes the food and the consumer adjusts the temperature, it stops working altogether and is sent back to the manufacturer, 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 those outcomes lead to someone beginning a cyber incident response process. The fact that the context makes 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.

Given these issues in incident reporting processes and in most cases a total lack of financial motivation to improve 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.

*EU Product Legislation*

An example of legislation intended to protect consumers is that of the European Commission [21], the most recognisable consequence of which is the CE Marking proving that products conform to this legislation. CE Marking relates to safety, health and environmental requirements.

This legislation is a result of the harmonisation of standards used in the various EU member states to ensure an adequate level of safety in products, for their free circulation in the single market.

EU harmonisation legislation applies to products being made available to an EU market, with the appropriate infrastructure to enable sales to that market (ordering and shipping systems). It covers all new products made available in the EU and second hand products entering the market from outside the EU. The Blue Guide definition of Made Available is as follows: *"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."* [21]

Manufacturers have to conform to the legislation in order to be allowed to use the CE Marking. Typically conforming 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 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: *"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."* [21] This means that the instructions that come with a product are 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 cyber threats cause.

There are of course issues with this approach — instructions for use are obviously important legally, but for those of us who attempt to read them, they are often written by someone who doesn't seem to understand the product themselves, can have some serious errors due to translation, and are usually open to interpretation. The ambiguity of natural language is not a new problem, and in the computer science field has been discussed as early as 1972, but the solution to semantic issues used in computer science is the use of formal languages. Whilst formal languages allow us to define lists of instructions for computers formal languages aren't suitable for communicating with humans [22]. Some manufacturers try and solve this problem by using images, which can't be translated wrongly and take less time to read so are more likely to be used. Unfortunately, there isn't a universal symbol for many elements discussed in user manuals so diagrams and icons might also be open for interpretation.

Liability as outlined by the Blue Guide [21] is not exclusively held by the manufacturer; producers providing components of a product, distributors and sellers also hold a level of responsibility should consumers make complaints about products sold in the EU.

It is difficult to define how this legislation handles IT–driven business models. If manufacturers are now selling 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, providing an attack vector on the original product would be dealt with.

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? If the in–car media app makes sufficient changes to the way the car runs do the app 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 whilst there are no known vulnerabilities for the specific components of the system does the manufacturer avoid all liability? Is the requirement for *"good engineering practice"* in the legislation sufficient to ensure security is considered in the design irrespective of known vulnerabilities? For example in the climate control system scenario, 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?

*UK Consumer Rights*

As discussed in Section III the line between sale of products and service prevision is being blurred by the growth of IT–driven business models. Consumers also have legal rights covering both goods and services offered, via the Sale of Goods Act 1979 and the Supply of Goods and Services Act 1982 [23]. Together these state the following:

*"Goods should be:*

1. *As described*
2. *Of satisfactory quality*
3. *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:*

1. *As agreed.*
2. *With reasonable care and skill.*
3. *Within a reasonable time.*
4. *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 app is in fact a service offering it still has to be carried out with reasonable care and skill. Just as a builder should be expected to secure a site at the end of the day, app developers may be responsible for fixing known vulnerabilities in their software for the duration the service is offered for, and not damaging the systems they connect to. If the refrigerator in the scenario malfunctions due to an attacker gaining access, via a known vulnerability in the app provided by the manufacturer and igniting a fire, the manufacturer might be considered not to have put sufficient care into the service they provided. Obviously, in order to be held responsible one of the investigators would have to have considered the hazard of cyber threats in the system.

Whilst companies have obligations to adhere to certain standards and consumers have legal rights, it is also worth making two comments at this point. The consumer market for household appliances etc. is traditionally far less accepting of faults than the IT market, possibly mainly due to existing legislation. Secondly, 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.

## IV.2 Standards

Cyber security issues are not something new for safety–critical systems. In hazardous environments standards have been updated, for example the MoD standard for safety management in defence systems now includes a section on cyber security and data integrity [2]. 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 also has a security section, this time requiring hazard analysis

to include the possibility of *"malevolent or unauthorised actions"* [24]. That means that if our scenario 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. As these types of standards are used for aircraft, the expected mean time between failure rate might be a little beyond the typical Android app developer.

These standards or similar are typically used for very large and/or complex systems. Whilst it's 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 standard 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, but the overlap from security to safety is inconsistent. ISO 27002 (information security management [25]) has been adapted through British standards to fit various industry sectors. The adaptation coving the energy utility sector, where there are large numbers of safety–critical systems, discusses cyber security through the integrity and availability of safety functions [26]. The standard has also been adapted for the telecommunications industry [27]. 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 standards or draft standards relating to the scenario discussed in this study, which, due to their cross–references and occasionally obscure applications, make it extremely 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 [28]—[40]. 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 [30].

IT equipment, an example of a consumer product with both safety regulation and inherent cyber security issues, also has its own safety standard [41]. The IT equipment standard discusses abnormal operation and fault conditions — a description broad enough to cover programming bugs, and states requirements intended to both 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. Whilst safety issues caused by a cyber security breach are not discussed as a hazard in their own right, it can be assumed that a programming fault would usually not be any more hazardous were the cause malware rather than poor programming skill. The most notable differences between these systems, and the ones discussed in the scenario 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 [42]. This example is interesting as it manages to prove the safety of its core software modules independently to any of the systems which may be plugged in later. By supplying standard interfaces, etc. they also make it easier to re–use code which has already been verified potentially increasing software quality across the whole vehicle.

### IV.3 Good Safety Engineering Practice

Where standards fail to be prescriptive about how a design issue should be dealt with the fall back is that products should be designed following good engineering practice [21]. Good engineering practice in such a cross–disciplinary field is probably fairly open to interpretation, however books about safety critical–systems such as those by Storey [43] and Leveson [44] describe the traditional approach to safety, in what has now become the field of cyber physical systems.

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 [41]. 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 [43]. 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 app, if there's a problem with the way the app is communicating with the car, the car's electronic control units should default to a safe state. This assumes that the app wasn'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 software does not degrade over time in the way electrical components do. This means that software has to be rigorously validated or verified as part of the design process. Software is often assessed for its reliability as a means of proving safety [43]. That means that 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.

In addition to these traditional approaches to safety Leveson has written a second book adapting safety practices to align better with systems engineering. She feels that systems are becoming more complex, due to the reduction in physical constraints of electrical components, allowing designers to be more creative [45].

In developing a safety model that links with systems engineering, she notes that reliability really has nothing 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.

The safety of a system can change over time, and whilst 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 whilst others don't change to accommodate that evolution, the system may become unsafe (asynchronous evolution [45]). As most systems are designed by multiple people, the interaction between components and their respective safety controls are the areas of a system of most concern. A good example of this is the refrigerator scenario, with its website and multiple app 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 become extremely difficult in a dynamic system, designers have to manage to either accredit the sub–system 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 in terms of interaction and collaboration agreements between stakeholders. In the aircraft industry, one of the approaches that has been developed is the SAVI Virtual Integration Process [14], which specifically defines how designers can work on an architecture together without revealing sensitive intellectual property to other stakeholders.

Whilst the books mentioned above give a good indication of how safety is treated with regards to computer systems, there is no detailed information on how to measure hazards posed by cyber threat. Part of the issue 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 in two years, making it no longer road legal. Unfortunately whilst legal levels of safety are expected to be maintained for the reasonable lifespan of the product [21], cyber threats continuously 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. The software can remain the same throughout the life of the product but the environment, and equivalent security level, change over time potentially making the system unsafe.

Traditional approaches to avoiding errors, for example voting systems, may no longer 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 to a system

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 a faulty climate control systems or fridges are unlikely to be experts in cyber security.

Finally, hazards found during analysis result in design constraints, which in Leveson's safety model require control algorithms [45]. It is difficult to see how the hazard of the controller itself 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 [45]. These elements of the process have been largely ignored in the consumer market, where all that is required are product safety checks at random intervals within a product's reasonable lifespan. Whilst 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 reliability and safety in appliances. As more of these devices go online and safety requirements continue to require high quality software, consumers may force software standards to improve in general.

# Section V –New Requirements

## V.1 Summary of Consumer CPS System Attributes

The attributes discussed in Sections I–IV have been aggregated in the following list in order to act as a springboard to building 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 cyber security related risks that the consumer faces.

*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, 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.

*Evolution of Environment*

4. Pervasive connectivity, providing the environment for ubiquitous computing to develop, with huge numbers of devices to potentially go online.
5. 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.

*Evolution of Business Drivers*

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

*Evolution of Development Team*

10. Increase in the number of separate development teams for the systems.
11. 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.
12. Potential for designers to become over–burdened due to increases in complexity.

*Evolution of User Expectation*

13. 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.
14. Instructions have to be clear, short and simple, with user–friendly interfaces on the product.
15. Made available without noticeable price increases.

*Evolution of System*

16. Small scale/budget complex systems — systems of sub–systems, with associated synchronisation issues.

17. Huge numbers of highly interconnected systems, with small degree of separation from safety critical systems.

18. Software security degrades over time, networked devices making this a safety hazard.

19. Differing expected lifespans for sub–systems, from months to tens of years, to be balanced.

## V.2 Requirements

Working from the Consumer CPS system attributes defined in subsection V.1, it is possible to begin drawing together a list of system requirements to provide requisite safety and security.

*Standards or Legislation*

There are certain attributes, such as the extremely low budget manufacturers would have available for system alteration, which make collating a set of feasible requirements particularly difficult. Whilst this attribute does not spawn safety or security requirements in its own right, it may alter the way requirements are handled. In order to continue meeting existing safety standards certain requirements may need to be recognised by standards bodies. This should also prompt the adoption of other requirements as good engineering practice. There are two overarching requirements at a standards level:

**i. The recognition of cyber threat as a potential safety hazard in networked consumer cyber physical systems.** Prompted by attributes 1, 2, 5, 17 and 18.

**ii. The acceptance of a modular design in safety conformity accreditation so that products are considered as part of a dynamic system including add–on products or services, rather than an isolated appliance [28]—[40] or predefined system [2], [24].** Prompted by attributes 10 and 19.

*Design*

Designers need to consider the change in requirements of the original product and a change in scope to include a wider system of applications and services.

*Original Product*

**iii. Identify hazards which could have an increased level of risk due to the product being networked.** *Is it possible for the climate control system to be reprogrammed remotely so that it malfunctions and overheats?* Prompted by attributes 1, 2, 3, 5b, 14, 18 and 19.

**iv. Identify new hazards introduced by networked product.** *Does the in–car media app share personal data in a way that could be hazardous for the consumer?* Prompted by attributes 1, 2, 3, 5a, 13 and 14.

**v. Identify new hazards introduced by networking capability.** *How does the refrigerator connect to the internet? Is it automatic? Can it be easily disconnected if a dangerous cyber threat is uncovered? Can the group of fridge owners or users who are not internet users remain safe?* Prompted by attributes 1, 2, 3, 5, 13 and 14.

*System Level*

**vi. Identify connections to safety–critical systems and faults which might be propagated.** *The climate control system connects directly to the fire alarm system; could a virus propagate to this system?* Prompted by attributes 1, 2, 3, 5, 14 and 17.

**vii. Ensure developers of product add–ons are aware of system hazards, and adhere to pre–agreed programming standards which can reduce these risks.** *Can the developers of the supermarket apps be provided with a specified standard to follow or communication protocol to use which limits the threat of an attacker using their app as an attack vector? Can product APIs be designed so that safe/secure usage is the only option?* Prompted by attributes 1, 2, 3, 5, 6, 10, 12 and 14.

*Through–Life*

A networked service provision approach provides an opportunity for the manufacturer to retain a communication link and some control over the products they sell. New risks introduced by networking encourage the use of through–life service provision.

**viii. Ensure any new product add–ons adhere to the manufacturer's standards, don't introduce new hazards, and that add–on standards remain up to date.** *When a new supermarket asks to make an app to communicate with the scenario refrigerator the manufacturer informs the supermarket of the standards they require, also checking that the new supermarket app isn't directly connected to a safety critical system. The programming standards have been regularly reviewed and updated to ensure sufficient levels of security are maintained.* Prompted by attributes 1, 2, 6 and 18.

**ix. Add automatic fault reporting mechanism to the product to improve speed of attack detection.** *The climate control system begins to malfunction, only working in bursts of 15 minutes. Jen tells her landlord about the problem and he calls a heating engineer who can't replicate the fault. The climate control system sends an automatic report to the manufacturer who can send out an update to fix the system, and remove the vulnerability.* Prompted by attributes 1, 2, 7, 17 and 18.

**x. Require suppliers to notify if any new vulnerabilities are found in product components.** *A microcontroller used in the refrigerator is found to have a security vulnerability. Its manufacturer notifies its customers who can then send out new firmware in the next update.* Prompted by attributes 1, 2 and 5b.

**xi. Provide ongoing software update services to combat evolving cyber threats.** *The manufacturer of the in–car media app sends out automatic updates whenever they find a security issue that could make the app unsafe to use.* Prompted by attributes 1, 2, 5b, 18 and 19.

These requirements will be developed in the following section and compared to existing relevant models or solutions for CPS, safety or security. This gap analysis type approach should establish where new or adapted models are required for this field and how well–organised developments in this field have been up until now.

## Section VI — Looking Ahead

The new requirements listed in Section V prompt an overhaul of the manufacturer's product though–life process, as well as adaptations to the design to include security elements, at the same time that a networking capability is added. It is assumed that requirements i and ii will be included in standards as computing becomes pervasive, the ecosystem evolves and more cyber threats in the consumer CPS domain are reported.

### VI.1 Design phase

In describing the inclusion of safety measures in the systems engineering design process Leveson [45] suggests that hazard analysis should begin to be considered as early as possible in the design process. This approach is facilitated by the IT–driven business model [18] described in Section III, 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 the refrigerator described in our scenario could have more than 95% identical parts to a previous model. The biggest innovation in terms of its physical makeup is 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 Pièvre–Cambacédès and Bouissou discuss the difference between security risks and safety hazards [1]. In the scenarios used in this study both are present, however a point they make remains valid — whilst cyber threat 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 difference is that in the new analysis the trigger for an accident could be system malfunction, or a side effect of a hacker re–programming the system. The fire may also be caused intentionally by a hacker, which becomes a physical security risk rather than simply a safety hazard. The fire hazard remains exactly the same in each of the three examples, but the likelihood, and so the risk, varies enormously.

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 [46].

This analysis could be approached using existing hazard analysis tools such as HAZOP, which have been adapted to identify security threats [47]. However, using HAZOP in the scenarios described in this study would make system design and implementation extremely process–heavy as it is designed for large safety critical systems. For any type of security analysis to be adopted in this field, a more lightweight option for evaluating the hazards of potential cyber security triggers is required.

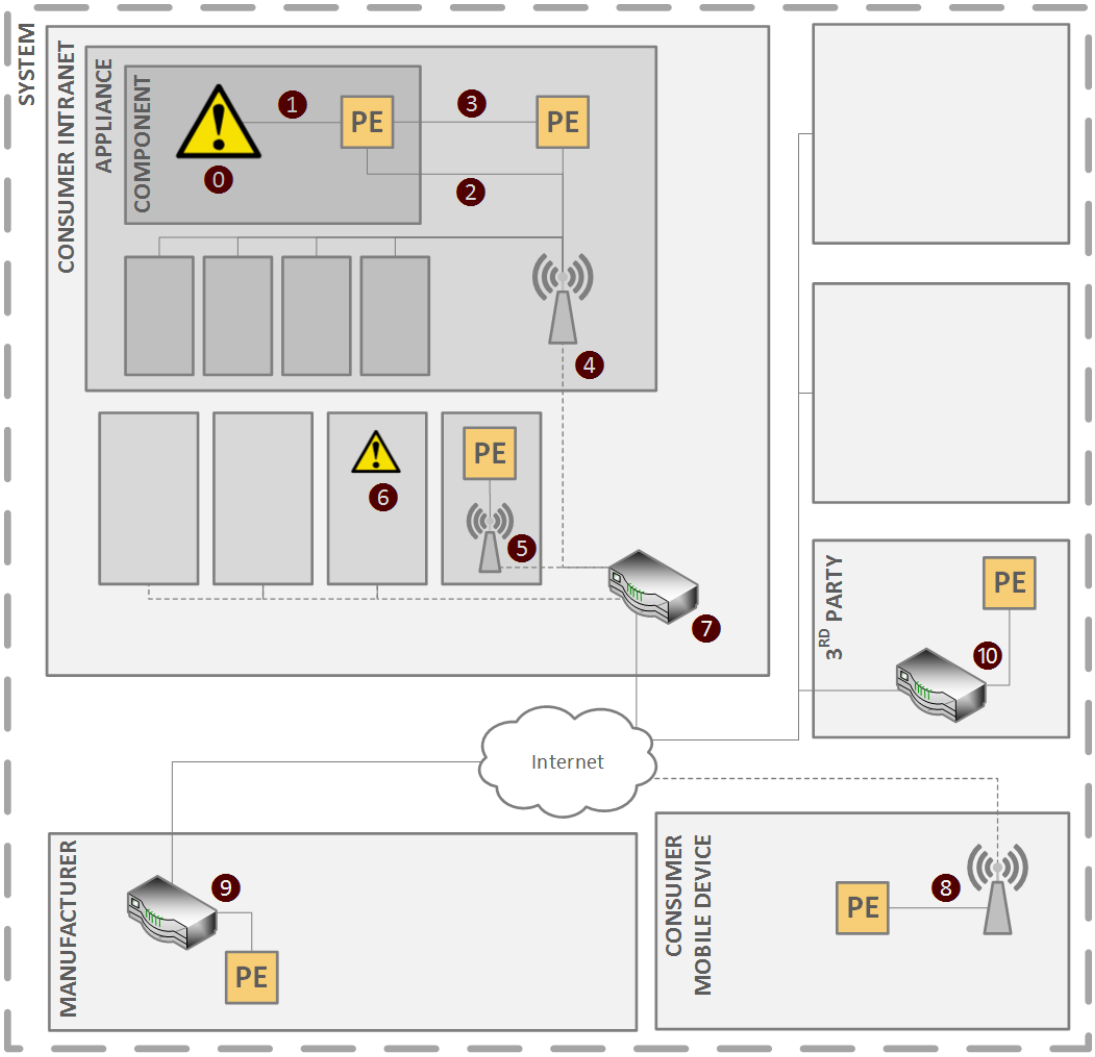


**Figure 4** Generic Consumer CPS Architecture

Figure 5 shows a generic consumer CPS architecture, with a system of distributed subsystems containing programmable elements (PE). The subsystems all have defined perimeters, although systems can interact with each other via electrical or radio signals. The consumer intranet could be a home, in a vehicle, or a BAN, with a wireless router as the connection point to the internet. Inside of the intranet are a number of networked appliances, including the one which is the element of the consumer CPS that interfaces with the physical world. This element of the consumer CPS has a number of components with programmable elements connected to the network card.

Both Leveson and Ward et al. advocate considering systems as sets of subsystems [14], [45]. Ward et al. suggest evaluating the system using a diagram or architectural model, annotated by the various development domains when they discuss the integration of complex systems in aircraft development. Whilst consumer CPS are unlikely to ever be anywhere near as complex or safety–critical as 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 very clear, because their scope and perimeters are defined. As cyber security is often essentially a process of creating barriers around things which need protection, and safety issues often arise where subsystems overlap, knowing where the perimeters of each system element fall is extremely useful. For particularly important elements of the system it may be necessary to employ multiple barriers at various levels of the system, using the defence–in–depth model to ensure security [48].

*Relevant Hazards*

Once the system elements have been defined, the architecture should highlight the existing hazards (shown as

hazard triangles in Figure 5). It should also be possible to see which hazardous elements are controlled by a programmable element (PE) — for example via the connection labelled 1. Unless the appliance is being 3D printed by the consumer no amount of hacking is going to influence the sharpness of the corners of an appliance once it has left the factory. Hazards which have no relation to programmable elements aren't going to need to be re–evaluated in the context of cyber threat.

The next question is whether the PEs identified are re–programmable? Where hazards are associated with PEs their ability to communicate, in real time or not, with the network card either directly (Label 2), or via a higher level controller in the system (Label 3) is vital to their being a cyber security risk, as it is this connection that would allow a malicious actor to alter system function. If the PE is not networked then there is no change to the existing hazard analysis.

This should provide a sub–set of hazards in the appliance which could be effected by cyber attack, to be considered when evaluating the rest of the architecture, in line with requirement iii.

*Architecture Evaluation*

The final element to consider inside of the appliance is the networking card and the level of security it provides to the perimeter of the appliance (label 4) , in line with requirement v. Various security measures have been suggested in the CPS field to provide security at this level [7], [17], [49], although their level of development and their adaptability to the computational constraints of a household appliance are unclear. Consideration should also be given at this point to what information is readily shared by the appliance as part if it's new capability — is the consumer's privacy protected to a safe level fulfilling requirement iv?

The next issue is the presence of other networked CPS in the same consumer network (labels 5 & 6). There's a potential for interconnections and/or interference from other networked appliances, as well as other systems with hazardous elements which could be tampered with via any foothold an attacker gains in the original appliance. These elements could significantly increase the complexity of a consumer CPS, and their safety implications need to be considered to fulfil requirement vi. The ISO/IEC 18012 series provides product interoperability standards [50], whilst a draft series of standards (ISO/IEC 30100 [40]) extends the interoperability to network resource management, including remote management.

Label 7 is the consumer's Small Office or Home (SOHO) router. Vulnerability researchers have been focusing on these products recently, as Internet Service Providers (ISPs) often provide them as part of their service provision poorly configured, with few or no security measures [51]. This represents, for example, the perimeter between the consumer's home WiFi network and the Internet. Although this would be an excellent place to add additional security to the system it is also a perfect example of why these distributed systems are difficult to secure. The ISPs don't add much security to their devices, and don't allow consumers or third parties to make significant changes. Even if the manufacturer happened to be an ISP and had the freedom to change the router, the consumer's Wi–Fi network is used by many different types of system — if one of those was a laptop that the consumer used solely for the purpose of downloading malware the consumer would still have the right to expect their CPS to remain safe.

The consumer is likely to be using an application either on a mobile phone or via a web browser so their device may or may not be inside the consumer's home Wi–Fi network (Label 8). Irrespective of the location, the manufacturer can impose some security through controlling the quality of the software they develop as per requirements vii and viii; however, they are still placing an application on an untrusted platform and allowing it to remotely control a physical element of a CPS. Some suggestions have been made by researchers on how to make mobile apps secure enough for use in safety critical–systems, for example using virtualisation techniques to isolate safety–critical applications from the other applications [52].

Labels 9 and 10 concern the links from members of the supplier ecosystem to elements of the CPS — communication relating to fault reporting, updates, request and retrieval of data to populate applications etc. The manufacturer and their partners are probably used to securing their own networks and protecting customer data, but need to consider the cyber security they are offering the consumer as part of the new service they offer, ensuring that it would be difficult for malicious actors to gain access to client systems via a breach in their own security. This should satisfy requirements ix—xi.

Having considered the architecture in terms of cyber security, which could influence the level of risk associated with hazards in the CPS, Leveson suggests considering the controls which could be put in place to reduce the hazards identified to an acceptable level [45].

*Design Constraints and Hazard Controls*

From this discussion of the architecture of a CPS it becomes very clear that there are some areas where the different stakeholders have the ability to implement security measures and others where they have little to no control over the underlying system. In those cases, 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 the operating system their mobile phone is running. Other modules of the system need to have good enough security to meet the original hazard threshold despite these limitations, meaning that the main security requirement is pushed onto the appliance portion of the system where the computing constraints are likely to be the tightest.

Severely limited availability of computational power and communication bandwidth for cyber security is widely discussed in the context of larger scale CPS, for example in ICS and SCADA systems [53]. However, whilst 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 which are less of an issue. Firstly, whilst the design of a refrigerator or heating system may be used from one product to the next there is not the same issue of making changes to an operational system — owners are not going to 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 an ICS designer would only have to opportunity to do with a brand new multi–million pound system.

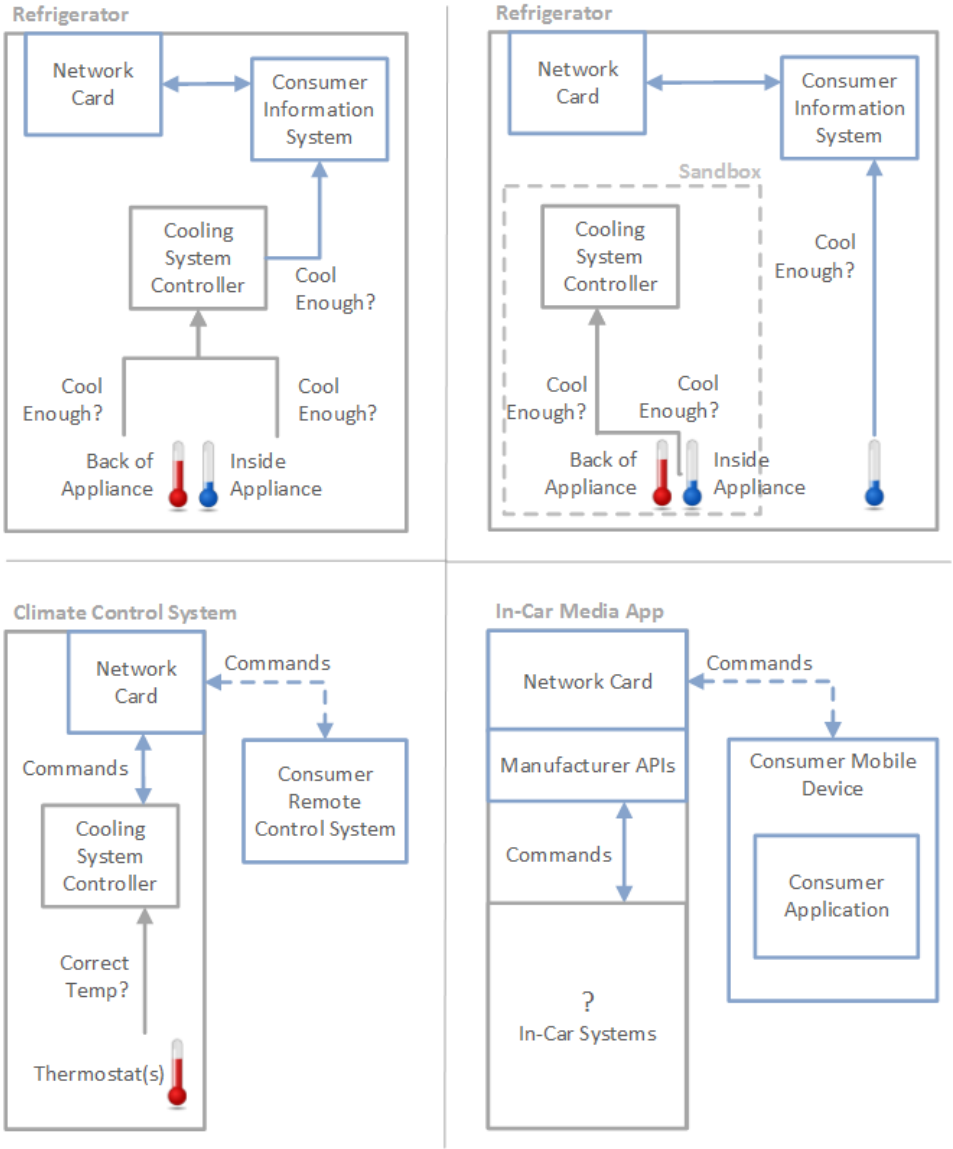


**Figure 5** Scenario Architectures

Another difference is in the higher level of flexibility a consumer CPS has to update and restart. Traditional safety–critical systems tend to control processes, where downtime is unacceptable or where unscheduled maintenance or under–tested updates could cause an accident. In the fridge scenario, even if the system 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 app the system is either on or off — when the app is actively connected to a vehicle it could be considered switched on in the same way a washing machine would be mid–cycle. In that circumstance the updates could be planned to coincide with times that the system is not in active use. In any of these scenarios the manufacturer is aiming to make a profit not through highly complex bespoke systems, but through the sale of large numbers of small systems. The scale of these systems, 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 function of a CPS it should be possible to sandbox some new elements, separating them from safety–critical functions. Figure 6 shows some possible high level architectures for the scenarios discussed in this study. The refrigerator scenario has two designs, showing how it might be possible in some instances 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 to the manufacturer, but is probably 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 the former being included because 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 — they cannot change the architecture of the electrical systems in the car to isolate their app from the safety–critical systems, they just have to hope that they don't override any safety measures that the manufacturer has put in place. In these cases cyber security solutions more commonly associated with ICT systems, such as firewalls and IDS, may be the only answer. Whilst 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.

*Instructions*

As mentioned in Section IV instructions play an extremely 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 very differently in fields where the products on sale are appliances, as opposed to software. This is mainly due to the completely different outcomes of risk assessments. On the manufacturing side instructions are provided to alert users to potential hazards and comply with safety regulation. 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 on safety.

In the software industry the bulk of the instructions are provided as a result of a risk analysis, aimed at protecting the developer from copyright theft or liability, with some limited "getting started" advice also provided. Instructions in the Terms and Conditions are usually presented in the form of hundreds of lines of impenetrable legalese which no one expects the user to read before they agree to them. Where advice on the use of a piece of software or device is provided it is usually displayed at the time of use in single line instructions, for example in the Android operating system. This approach, along with the tactile nature of technology and the age we now start using it, 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 [54].

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.

Some suggestions for protecting users might be to direct them to the manufacturer's own website as a front page

for downloading applications, before redirecting them to specific online 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 website. 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.

### VI.2 Through–Life

The requirements outlined in Section V necessitate an evolution from the random safety checks required by the safety legislation to a more involved through–life process involving systematic security updates — either for system software, or 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, reducing the ongoing cost of maintaining the safety of these systems.

It may be possible to push the responsibility for paying for security measures — up to date antivirus etc. onto the user, but in the case of a serious safety issue requiring the customer to make an additional purchase to ensure their safety may not be sufficient to remove liability. If a manufacturer was selling a refrigerator with razor sharp edges, it wouldn't be sufficient for them to tell users to purchase a strong pair of gloves — they would have to change the design. However, if the manufacturer is selling an oven, it's perfectly acceptable to suggest that the user purchase oven gloves. Many of the more safety–critical systems in the consumer CPS field also have a far higher price tag than a standard kitchen appliance, so it may be the case that the way through–life security measures are applied and paid for varies depending on the level of risk associated with the system.

## Section VII –Conclusions

In conclusion, this study has drawn together different attributes and business drivers of the consumer CPS ecosystem, in order to produce a set of requirements for the inclusion of cyber security measures in existing consumer safety models and legislation.

The study began by discussing the convergence of safety and security fields in the context of emerging ubiquitous computing models — consumer CPS. These consumer CPS were described through 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 discussed in the scenario to satisfy the original sentiment of safety legislation and reduce producer's liability, cyber security issues will need to be considered as part of the product lifecycle. Whilst this is needed to maintain a level of safety equivalent to that currently experienced there is a lack of clear legislation, 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 threat [2], so it is not considered 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 online. These requirements were discussed in more detail in Section VI in the context of the system architectures of the study's scenarios as well as existing models for safety engineering in complex systems.

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 which 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. 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 cyber attacks pose. In order to provide comprehensive safety solutions, producers need to find ways to communicate and collaborate with other members of the supplier ecosys-

tem, ensuring that vulnerabilities don't fall between the zones of responsibility.

Another area for future work is the user instruction sets. A user also has to be considered as part of the system — whilst systems become more complex users become less knowledgeable about the way they work and less able to evaluate the risks they take. 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 they need to be presented in an accessible format. For example if safety instructions in the application subsystems of a consumer CPS are made to look like software terms and conditions that almost guarantees that they will be ignored.

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 way a system may be used the only way they can reduce their liability is through clear instructions on how to use the system safely. This is highlighted in discussions on the third scenario, that of the in-car media app, throughout this paper.

The domain of consumer CPS 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 safety levels are maintained in a world of ubiquitous computing.

### Acknowledgements

*With thanks to EPSRC for funding the project and Andrew Simpson for all his brilliant advice and supervision.*

## References

[1] L. Pièrre–Cambacédès and M. Bouissou, "Cross-fertilization between safety and security engineering," Reliab. Eng. Syst. Saf., vol. 110, pp. 110—126, Feb. 2013.

[2] Ministry of Defence, "Interim Defence Standard 00-56 Part 1." Dstan, 21-Feb-2014.

[3] T. Wolf, M. Zink, and A. Nagurney, "The Cyber-Physical Marketplace: A Framework for Large-Scale Horizontal Integration in Distributed Cyber-Physical Systems," in 2013 IEEE 33rd International Conference on Distributed Computing Systems Workshops (ICDCSW), 2013, pp. 296—302.

[4] Proofpoint, "Your Fridge is Full of Spam, Part II: Details," ThreatInsight, 21-Jan-2014. .

[5] Y. Oren and A. D. Keromytis, "From the Aether to the Ethernet—Attacking the Internet using Broadcast Digital Television," To appear in the Proceedings of 23rd USENIX Security Symposium (USENIX Security 14), http://iss.oy.ne.ro/Aether, San Diego, CA, 2014.

[6] Beecham Research, "M2M/IoT Sector Map". http://www.beechamresearch.com/article.aspx?id=4.

[7] K. Doeornemann and A. von Gernler, "Cybergateways for Securing Critical Infrastructures," presented at the Security in Critical Infrastructures Today, Proceedings of International ETG-Congress 2013; Symposium 1:, 2013, pp. 1—6.

[8] Q. Shafi, "Cyber Physical Systems Security: A Brief Survey," in 2012 12th International Conference on Computational Science and Its Applications (ICCSA), 2012, pp. 146—150.

[9] A. Banerjee, S. Kandula, T. Mukherjee, and S. K. S. Gupta, "BAND-AiDe: A Tool for Cyber-Physical Oriented Analysis and Design of Body Area Networks and Devices," ACM Trans Embed Comput Syst, vol. 11, no. S2, pp. 49:1—49:29, Aug. 2012.

[10] L. Pièrre–Cambacédès and C. Chaudet, "The SEMA referential framework: Avoiding ambiguities in the terms security–and safety,–" Int. J. Crit. Infrastruct. Prot., vol. 3, no. 2, pp. 55—66, Jul. 2010.

[11] E. A. Lee, "Cyber Physical Systems: Design Challenges," in 2008 11th IEEE International Symposium on Object Oriented Real-Time Distributed Computing (ISORC), 2008, pp. 363—369.

[12] R. Rajkumar, I. Lee, L. Sha, and J. Stankovic, "Cyber-physical Systems: The Next Computing Revolution," in Proceedings of the 47th Design Automation Conference, New York, NY, USA, 2010, pp. 731—736.

[13] D. Schneider and M. Trapp, "A Safety Engineering Framework for Open Adaptive Systems," in 2011 Fifth IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO), 2011, pp. 89—98.

[14] D. T. Ward, D. A. Redman, and B. A. Lewis, "An Approach to Integration of Complex Systems: The SAVI Virtual Integration Process," in Proceedings of the 2013 ACM SIGAda Annual Conference on High Integrity Language Technology, New York, NY, USA, 2013, pp. 43—46.

[15] D. N. Serpanos and A. G. Voyiatzis, "Security Challenges in Embedded Systems," ACM Trans Embed Comput Syst, vol. 12, no. 1s, pp. 66:1—66:10, Mar. 2013.

[16] S. Karnouskos, "Smart houses in the smart grid and the search for value-added services in the cloud of things era," IEEE International Conference on Industrial Technology (ICIT), 2013, pp. 2016—2021.

[17] R. Mitchell and I.-R. Chen, "A Survey of Intrusion Detection Techniques for Cyber-physical Systems," ACM Comput Surv, vol. 46, no. 4, pp. 55:1—55:29, Mar. 2014.

[18] H. Kagermann, H. Osterle, and J. M. Jordan, "IT-

Driven Business Models: Global Case Studies in Transformation". Wiley, 2010.

[19] "Fire statistics Great Britain - GOV.UK." https://www.gov.uk/government/collections/fire-statistics-great-britain.

[20] Mitre, "Common Vulnerabilities and Exposures," CVE List Main Page. http://cve.mitre.org/cve/index.html.

[21] ENTR.C.DIR, "The Blue Guide–on the implementation of EU product rules 2014." European Commission, 04-Feb-2014.

[22] I. D. Hill, "Wouldn–t it be nice if we could write computer programs in ordinary English - or would it," Comput. Bull., vol. 16, no. 6, p. 306, 1972.

[23] Trading Standards, "A Trader–s Guide to Goods and Services." Norfolk County Council.

[24] BSI, "BS EN 61508-1:2010." British Standards Institute, http://www.bsigroup.co.uk/, 2010.

[25] ISO, "ISO/IEC 27002:2013." International Standards Organisation, http://www.iso.org, 2013.

[26] BSI, "PD ISO/IEC TR 27019:2013." British Standards Institute, http://www.bsigroup.co.uk/, 2013.

[27] BSI, "BS ISO/IEC 27011:2008." British Standards Institute, http://www.bsigroup.co.uk/.

[28] BSI, "BS EN 60335-2-30:2009+A11:2012." British Standards Institute, http://www.bsigroup.co.uk/, 2012.

[29] BSI, "BS EN 60335-2-73:2003 +A2:2009." British Standards Institute, http://www.bsigroup.co.uk/, 30-Apr-2010.

[30] BSI, "BS EN ISO 16484-5:2012." British Standards Institute, http://www.bsigroup.co.uk/, 2012.

[31] BSI, "BS EN 60730-1:1995." British Standards Institute, http://www.bsigroup.co.uk/, 05-Apr-2004.

[32] BSI, "BS EN 60730-2-14:1997 +A2:2008." British Standards Institute, http://www.bsigroup.co.uk/, 30-Apr-2010.

[33] BSI, "BS EN 14908-1:2014." British Standards Institute, http://www.bsigroup.co.uk/, 2014.

[34] BSI, "BS EN 12098-3:2013." British Standards Institute, http://www.bsigroup.co.uk/, 2013.

[35] BSI, "BS 5760-13.5: 1996 IEC 605-3-5: 1996." British Standards Institute, http://www.bsigroup.co.uk/, 1996.

[36] BSI, "Draft BS EN 60335-2-24:2010/Fp." British Standards Institute, http://www.bsigroup.co.uk/, 08-Apr-2014.

[37] BSI, "Draft BS EN 60730-1." British Standards Institute, http://www.bsigroup.co.uk/, 04-Nov-2011.

[38] BSI, "Draft BS EN 60730-2-6." British Standards Institute, http://www.bsigroup.co.uk/, 06-May-2014.

[39] BSI, "Draft BS EN 60730-2-9." British Standards Institute, http://www.bsigroup.co.uk/.

[40] BSI, "Draft BS ISO/IEC 30100-1." British Standards Institute, http://www.bsigroup.co.uk/, 15-Jul-2013.

[41] BSI, "BS EN 60950-1:2006 +A2:2013 Incorporating corrigenda June 2006, August 2006, September 2010, August 2011, October 2011 and August 2012." British Standards Institute, http://www.bsigroup.co.uk/, 31-Oct-2013.

[42] AUTOSAR, "AUTOSAR - AUTomotive Open System ARchitecture." http://www.autosar.org/.

[43] N. Storey, "Safety-critical computer systems". Harlow: Addison-Wesley, 1996.

[44] N. G. Leveson, "Safeware: System Safety and Computers". Reading, Mass: Addison Wesley, 1995.

[45] N. Leveson, "Engineering a safer world: systems thinking applied to safety". Cambridge, Mass; London: MIT Press, 2011.

[46] N. Mayer, "Modelbased Management of Information System Security Risk," University of Namur, 2009.

[47] R. Winther, O.-A. Johnsen, and B. A. Gran, "Security Assessments of Safety Critical Systems Using HAZOPs," in Computer Safety, Reliability and Security, U. Voges, Ed. Springer Berlin Heidelberg, 2001, pp. 14—24.

[48] Directorate for Command, Control, Communications, and Computer Systems, "Information Assurance through Defense-in-Depth". U.S. Department of Defense Joint Staff, 2000.

[49] K. K. Venkatasubramanian, T. Mukherjee, and S. K. S. Gupta, "CAAC — An Adaptive and Proactive Access Control Approach for Emergencies in Smart Infrastructures," ACM Trans Auton Adapt Syst, vol. 8, no. 4, pp. 20:1—20:18, Jan. 2014.

[50] ISO, "ISO/IEC 18012-1:2004." International Standards Organisation, http://www.iso.org, 2004.

[51] C. Heffner and D. Yap, "Security Vulnerabilities in SOHO Routers." Sourcesec, http://www.sourcesec.com/Lab/soho_router_report.pdf, 2010.

[52] M. Geier, M. Becker, D. Yunge, B. Dietrich, R. Schneider, D. Goswami, and S. Chakraborty, "Let–s Put the Car in Your Phone!," in Proceedings of the 50th Annual Design Automation Conference, New York, NY, USA, 2013, pp. 143:1—143:2.

[53] D. Kuipers and M. Fabro, "Control systems cyber security: Defense in depth strategies." United States. Department of Energy, 2006.

[54] S. G. Schar and H. Krueger, "Using New Learning Technologies with Multimedia," IEEE Multimed., vol. 7, no. 3, pp. 40—51, 2000.