Xday, XX XXX 2008.

9.30 am - 11.15am (check this!)

University of Glasgow

DEGREES OF BEng, BSc, MA, MA (SOCIAL SCIENCES).

COMPUTING SCIENCE - SINGLE AND COMBINED HONOURS
ELECTRONIC AND SOFTWARE ENGINEERING - HONOURS
SOFTWARE ENGINEERING - HONOURS

SAFETY-CRITICAL SYSTEMS DEVELOPMENT

Answer 3 of the 4 questions.

1.

a) The DO-178B definition of COTS is "Commercially available software sold by vendors through public catalogue listings. COTS software is not intended to be customized or enhanced. Contract negotiated software developed for a specific application is not COTS software"

Describe the problems that the use of COTS creates for the verification and validation of safety-critical software.

[5 marks]

[Seen/unseen problem]
COTS systems are typically sold without providing end user access to the code – hence only black box testing is possible in the verification of a safety critical system. The previous quote also suggests that any errors which are found may not be easy to fix given that the developer may not view any particular customer as sufficiently important to warrant large scale changes. This would be the case, for example, if a company were using a mass market COTS products for a safety-critical application.

Validation is slightly more complex. An argument can be made that COTS introduce problems here because it can be difficult to know whether a system will deliver promised functionality until the acquisition process has reached a very advanced stage. Often testimonies about the utility of any software may come from previous customers with slightly different problems or requirements. Hence, it may be difficult to apply comments about the utility of any COTS system to a new context of use. 1 mark for each reasonable argument with plenty of scope for changes to the examples given here.

b) COTS operating systems pose a number of further problems in the certification of safety-critical systems because it can be difficult to make arguments about the safety of a novel application based on the successful use of the operating system with another application.

Explain why arguments about experience in previous applications cannot easily be used to support the certification of COTS operating systems.

[5 marks]

[Unseen problem]
The sample answer to part a) began to address this issue when looking at the problems of validating COTS applications. It can be difficult for a potential user to determine whether or not the characteristics of an operating system will be realized in any other application. The run-time behavior of the system may be very different – for instance in terms of the load on caches (if they are used) as well as in terms of I/O or secondary and primary memory access. These problems are exacerbated in any context that allows multi-tasking – in other words, the precise demands of multiple threads and processes that occur and are supported by an OS in one application are unlikely to be replicated in other applications. Static scheduling can be used to reduce uncertainty but at a significant cost in terms of the associated safety analysis.

In addition to the more specific answers associated with operating system characteristics, good students can call upon the host of arguments that limit the application of 'evidence from use' for software systems in general. Further reference might be made to the repeated attempts to use particular mass market operating systems in safety critical environments even in the face of explicit exclusions from warranties and attempts by the manufacturers to prohibit the use of their systems for this class of applications. 1 mark for each reasonable argument with plenty of scope for changes to the examples given here. An additional mark up to a total of five if they mention the general issues of software certification AND specific issues to do with real time operating systems in this context.

c) Lynuxworks have recently developed a hard real-time operating system that they argue can be certified to DO-178B level A in a number of different applications. It uses both space and time partitioning to support multithreaded and multi-process systems in safety-critical environments and can be run on both Pentium and PowerPC platforms. It exploits many of the concepts familiar in other flavors of the UNIX operating system. In addition, it provides the notion of a Virtual Machine with partitions which are intended to prevent one process from interfering with another.

Describe how you would create an argument to convince regulators that a particular application did not suffer from interference in time or space between multiple concurrent processes and threads.

[10 marks]

[Unseen problem]

In the course, we have introduced a number of analytical techniques that might be used here. These include the development of structured safety cases – they have been told about GSN. We have described the application of formal methods to safety-critical systems and again, they might refer to various types of proof – model checking and theorem proving that might be used here to support the argument. If they mention these techniques then they should argue that informal prose can be difficult to structure and often leads to ambiguities or contradictions in the large-scale, complex and detailed arguments that would be required in this example. We have also introduced the key stages of the DO-178B approach – so the associated documentation and arguments about risk mitigation could be brought into any solution.

Other answers might look in more detail at the issues surrounding partitioning within an operating system. Traditionally, safety-critical systems restricted the amount of parallelism and multi-tasking that might be allowed by an operating system. This reduced the complexity of guaranteeing safety and liveness properties for complex systems, for instance, in the presence of CPU starvation between competing processes with unpredictable demands from both users and their operating environments. The hard-real time issues referred to in the question also raise concerns by creating additional constraints beyond the eventual progression considered within liveness requirements.

Ideal solutions might go on to consider space partitioning and the types of addressing schemes that would be required to avoid memory contention and over-writing of shared resources. Answers might also mention that these concerns taken together would imply that only a subset or kernel of a full UNIX implementation might be certified for use in safety related applications in order to exclude any interference or side effects from other non-essential processes.

Finally, to obtain full marks I would hope that students consider the dual issues of certification and testing for an operating system and for the applications that run on it. To what extent is it possible or appropriate to partition safety analyses in this manner and how much do application developers need to know about the underlying implementation of the systems that they call upon in safety-critical applications? Hence there is a link between all three elements of this question.

2.

a) The HAZOPS technique typically begins with a form of functional block analysis which breaks complex systems down into their component parts. What problems might arise in applying HAZOPS to analyze the safety of software systems?

[4 marks]

[Seen/unseen problem]

Functional block analysis provides a top down view of a system. In conventional engineering this can be very appropriate. Systems are composed of subsystems and so on. Even in software, hierarchical forms of analysis can be useful, breaking applications down into core functionality. However, these approaches may ignore the dependencies that often at the interfaces between modules and sub-modules. There is also an increasing school of thought which argues for resilience engineering to look more at the demands placed on a system in its context of use at all stages of analysis. Hence in this view it might make little sense to try a functional decomposition of a software application without considering its wider interactions from the start.

b) HAZOPS proceeds by applying guidewords to the functional models of a complex system. The intention is to identify the impact of potential deviations from the intended design. Many applications of the technique are based around a general collection of guidewords such as "NO OR NOT", "REVERSE", "MORE", "OTHER THAN", "LESS", "SOONER THAN", "AS WELL AS", "LATER THAN", "PART OF".

Explain whether or not this list would be sufficient to conduct a HAZOPS for a complex software system and, if the list is insufficient then identify further THREE guidewords that might specifically support HAZOPS for programmable electronic systems.

[6 marks]

[Unseen problem]

A range of solutions are possible here. An argument can be made that this list is appropriate for software failures with 'SOONER THAN', 'LATER THAN' and 'AS WELL AS' all having reasonable interpretations within the context of real time and concurrent software systems. However, these general purpose guidewords were never intended to be sufficient for all domains and so good solutions might identify a number of alternatives. These could focus on traditional software engineering issues, such as 'buffer overflows' or 'type incompatibility'. Some of the Ada predefined exceptions might be mentioned? However, this is a relatively low level approach that runs against the more generic relationships within HAZOPS.

The following table shows how an alternate form of software hazard analysis (SHARD) introduces a number of more specific guidewords for a single language (in this case MASCOT). I would not expect this in the final solution to the above question but it illustrates this important point about the granularity of the keywords that are chosen in any potential solution.

		Failure Categorisation					
Flow		Provision		Timing		Value	
Protocol	Type	*Omission*	*Commission*	*Early*	*Late*	*Subtle*	*Coarse*
Pool	Boolean	No update	Unwanted Update	N/A	Old Data	Stuck at...	N/A
	Value	"	"	"	"	wrong in tolerance	out of tolerance
	Complex	"	"	"	"	Incorrect	Inconsistent
Channel	Boolean	No Data	Extra Data	Early	Late	Stuck at...	N/A
	Value	"	"	"	"	wrong in tolerance	out of tolerance
	Complex	"	"	"	"	incorrect	inconsistent

I would also welcome a more systems oriented approach where the answers considered issues such as side effects or undocumented features that could be considered during software variants of HAZOPS.

The class have been shown Chin and Leveson's Software Failure Trees so this approach might be mentioned in more sustained answers. Each branch in the software fault trees can be thought of as an 'aide memoire' in the same way that the HAZOPS guidewords focus attention on particular hazards.

c) HAZOPS studies can be documented in a tabular form using the following headings:

Deviation	Cause	Consequence	Safeguards	Action

A deviation typically refers to an element of the functional decomposition and one of the guide words. For instance, a sensor reading may be processed by the software 'LATER THAN' a specified deadline. Once a system has been deployed, experience of using software in an operational context can be used to guide subsequent HAZOPS studies. In particular, any deviations that have been observed during operation can be analyzed to identify causes, consequences and so on.

Briefly explain why the post operational application of HAZOPS to software systems requires a very high level management commitment to the monitoring, documentation and audit of a software application.

[10 marks]

[Unseen problem]

After software has been deployed, the company that was originally involved in the development of the software will have an interest in monitoring the operation of their product. However, the opportunities for access may be limited especially once an initial warrantee period has run out. In contrast, there may also be disincentives to some forms of monitoring – if management lacks the resources and commitment to address any potential problems. In other words, logs help identify potential 'Deviations' but do not guarantee the resources to identify or implement 'Safeguards'.

Further issues stem from the complexity and integration of recent software systems – it can be costly and technically difficult to diagnose potential problems from the logs given that the cause of an error report may lie in software produced by several different suppliers. In many cases, the costs of analyzing these error reports are met only to find that the issue is either well understood or relatively benign. In other words, there are significant technical and organizational barriers in moving from the observation of a 'Deviation' to identify the 'Causes' and 'Consequences'.

In addition, software audits can be held at regular intervals to specifically review any previous deviations – this provides managers with means of ensuring that 'Deviations' are not simply ignored and may also stand a greater chance of obtaining necessary specialist expertise from software maintenance teams during the period of each audit.

Irrespective of the approach followed, it is important that the various stages identified in the HAZOPS form are supported by an appropriate safety management system that helps an organization to devote adequate resources of time, money and expertise to support the post operational monitoring of software applications. This is essential if companies are to respond to changes in operating environment and working practices that can undermine the continued operation of safety-critical software.

3.

a) Lloyds Register of Shipping has coordinated studies into the impact that Electromagnetic Compatibility (EMC) can have upon the functional safety of maritime systems. They argue it is essential that designers and operators "account for the effects of EMC in the marine industry" using hazard and risk assessment techniques that are appropriate for particular EMC environments. The main sources of interference above decks are the ship's radio transmitters.

Describe the technical problems that can arise when trying to demonstrate EMC for computational systems in maritime applications.

[5 marks]

[Seen/unseen problem]

In terms of Electromagnetic Compatibility (EMC), it is important to demonstrate that any particular hardware platform is both resilient to external sources of interference and also does not create sources of interference that might affect other systems. These concerns apply to computational systems just as they do to any other on-board applications. However, there are numerous particular issues that arise for programmable applications ranging from the vulnerability of particular low-level circuit configurations through to the susceptibility of high level languages to bit-shift errors induced by interference. The maritime environment provides considerable challenges for the development and operation of computational systems. The power of the radio transmitters is stronger than many other land based sources of electromagnetic radiation given that they may have to transmit over several hundred miles. The source is also asynchronous and unpredictable. Further problems arise because electromagnetic compatibility must also be demonstrated for other sources of radio transmission from ships that can be in close proximity to a vessel in harbor.

b) Lloyds Register is a classification society. These are private companies that are authorized to inspect ships and issue safety certificates for States where ships are registered. A European directive (94/57/EC) describes requirements that classification societies must meet in order for them to be recognized by a Member State under European legislation.

What problems can arise when recruiting well qualified staff to work for certification societies and independent safety assessors?

[5 marks]

[Seen problem but not in this particular context]

It can be difficult for any agency to recruit staff who have the relevant expertise and yet also have the degree of independence necessary for safety assessments. Individuals who are competent in a particular domain will have worked with companies for a number of years and yet may later be called upon to inspect the work of those companies. This can create conflicts of interest, particularly in specialist areas where they may only be a small number of commercial organizations that dominate the market. Lord Cullen identifies similar problems in the aftermath of the Ladbroke Grove rail crash – arguing that the lack of external input can lead to the development of cultural norms that undermine safety within an industry. This raises further competency issues – it can be difficult to identify safety assessors that have domain expertise and also have sufficient understanding of the underlying concepts of risk assessment or safety management.

c) The UK Maritime and Coastguard Agency inspects a proportion of vessels in UK ports. During 2006-7, 4.9% of the foreign ships inspected in the UK were detained. A mean of 4.5 deficiencies were found during each inspection. Some states are named on an international 'black list' which

reflects particular concerns over their safety management. The 8.4% detention rate for ships from these states was over twice the rate for non-targeted flags. A total of 19 ships classed by members of the International Association of Class Societies (IACS) were detained where the Classification Society was found to be held responsible for the detention. The detention rate for IACS ships was 1.2%. Ships classed with Classification Societies which are not members of IACS were over 6 times more likely to be held responsible for a detention than IACS societies.

Imagine you have been asked to act as a consultant to one of the states on the international 'black list', write a technical report describing how you would go about improving the safety record of vessels that operated under the flag of that state. (Hint: you may refer to elements of your answer to part b)

[10 marks]

[Unseen problem]

This is a relatively open ended problem. The hint suggests that the state representatives should consider recruiting the help of independent safety assessors through agencies such as Lloyds Register. Given the statistical differences in the outcome of inspections, it may not be possible to improve 'safety levels' immediately to a position where a nation is removed from the black list, however, it should be possible to develop codes of practice to gradually encourage owners and operators to improve potential deficiencies that are otherwise identified during inspections.

A more extreme approach would be to encourage all owners to register or have their ships classified by a member of the IACS. This enables students to consider the trade-offs that exist between safety and commerce – the costs of classification and of ensuring that any subsequent safety concerns are addressed may be sufficient to dissuade many states from this approach even if it offered the opportunity of beiung removed from the black list.

An important issue in any answer to this question is that the UK Maritime and Coastguard Agency data does not link the outcome of inspections to the likelihood of being involved or implicated in an adverse event. Hence, the increased probability of being detained after an inspection may be viewed as an inconvenience by the owners who choose to register with a 'black list' state. In any event, the report mentioned in the question could propose that direct safety monitoring techniques be used to determine not only that any programme achieved removal from the black list but also that it led to an overall improvement in safety metrics. This is important because the question refers to an improvement in 'safety record' and does not directly ask about being removed from the list itself.

4. Safety can be described as an emergent property that cannot easily be determined before complex applications have been deployed in their eventual context of use. State whether or not you agree with this view. Develop your answer with examples and explain the implications that emergent views of safety have for the resilience engineering of complex systems.

[20 marks]

[Unseen problem/bookwork/Essay]
There are many different approaches to this question. The intention is to give everyone something that they can work on but also to provide the space for more able students to draw links between many different areas of the course. The students have seen sample solutions from the many different previous exam papers where the final question has always been an essay option.

The technical heart of the question refers to a particular aspect of what has recently been termed 'resilience engineering'. This itself is a development of longer term work in socio-technical systems, human factors and control theory. The idea as the question suggests is that workers and managers adapt their behavior after a system has been deployed. For instance, risk homeostasis theory has proposed a model of decision making and behaviour in which individuals have a target level of risk which they are prepared to accept. IIf you introduce a new safety innovation then this reduces the actual level of risk below the target and hence these individuals may alter their behavior – for example by trading additional risk for additional speed so that the safety innovation is counterbalanced and the overall target level is maintained. This behavior has been observed, for instance, in the behavior of cyclists when safety helmets have been made compulsory. The helmets protect the wearer from many serious forms of head injury but it has been noted that the overall frequency of cycling injuries increase. The connection to the question is that the estimated improvements of introducing the helmet legislation were undermined, it has been argued, by the emergent behavior of cyclists who may have believed themselves to be better protected from additional risks by the head gear.

Many people have rejected the arguments behind risk homeostasis or have sought to limit its application. For example, it is often not apparent to individuals that they have the benefits of safety devices. How does a driver know that ABS will protect them from a particular risk? If they cannot understand the potential benefits then how will they know a proportionate response in terms of their target level of risk. Others, including myself, have argued that emergent theories are often used as an excuse for poor engineering and the unexpected ways in which individuals interact with safety critical systems have often been anticipated by designers through, for instance, reporting systems within the design and operations teams.

As mentioned, many different solutions are possible. This is an area of active research and so it is possible to argue both for and against the assertion in the question.