Xday, XX XXX 2013.

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) Identify three different systems of governance that can be used to ensure the safety of the general public and identify ONE limitation for each of these different approaches.

[6 marks]

[seen/unseen problem]

Market forces – it can be argued that the forces of supply and demand are sufficient to regulate the safety of products available in a market. Consumers will not buy products that are considered to be unsafe. However, markets rarely work in a perfect manner – advertising can create demand for products that are unsafe, consumers often have imperfect knowledge of a products safety record and there are third party effects – such as passive smoking that affect people who do not directly consume a produce (one mark for a description and one for any of the limitations).

Tort and insurance – these approaches provide financial compensation for the victims of an unsafe product. If consumers lack perfect information before purchase, they can use the law to seek redress after an injury has occurred. Tort also enables third party compensation. Insurance enables companies to offset future liabilities that might otherwise act as a disincentive to investment. However, most tort is inefficient in the sense that considerable amounts of money that might otherwise be spent on safety improvements are instead used to maintain the legal system. Tory is retrospective. Other answers might focus on the problems associated with malicious or undue litigation.

State regulation - most governments do not rely only on market forces or on tort to maintain safety. In contrast, they create regulatory agencies who help to ensure that companies employ appropriate standards to guide their work. Fines may also be used even when accidents have not occurred if companies violate their regulatory requirements, hence they avoid the retrospective nature of tort. However, regulatory organizations can be expensive and bureaucratic. They may also be seen to place constraints on innovation, especially if staff lack appropriate technological expertise.

b) The Chair of the UK Health and Safety Executive recently confirmed that her organization has been required to make a minimum of 35% savings; this is the same percentage reduction in costs expected for the Department of Work and Pensions as a whole. What impact might such cuts to regulatory agencies have upon companies that develop safety-critical software across a range of industries?

[5 marks]

[unseen problem]

Regulatory organizations fulfill a number of roles, any of which might be compromised by significant cuts in funding:

- Standard setting. As mentioned in the previous answer, regulatory agencies help to identify the technical standards that embody acceptable means of compliance with national regulations. Lack of funds can starve regulators of appropriate technical input in identifying and assessing potential standards within new and emerging areas or in revising existing standards documents – for example, the HSE played a significant role in shaping IEC61508.

- Guidance and advice giving. In many cases, companies need help in learning how to apply particular tools and techniques, embodied within various standards. Regulatory organizations often publish guidance that helps companies to interpret international requirements within the framework of national regulations. It is costly and time consuming to develop this material – especially outside of the 'core' areas of risk assessment, hazard analysis etc.

- Inspection and enforcement. Regulatory agencies play an important role in determining whether or not a company or public body has met its statutory obligations. Reduced funding can prevent regulators from inspecting organizations or from starting enforcement actions.

- Monitoring, Incident Reporting and Lessons Learned. Incidents and accidents can still occur even if regulators identify appropriate standards and conduct enforcement actions. In such cases, the analysis of previous incidents can help to suggest changes in regulatory regime. This too may be compromised by insufficient funding, especially in technical areas where it can be difficult to retain staff with appropriate skills.

Other answers are possible – include the regulatory role in research (2 marks for each area of concern up to a total of 5 marks). First class answers might mentioned the need for regulatory agencies to monitor accident rates and enforcement actions to ensure that any cuts are not having these potential adverse effects.

c) THE UK Health and Safety Executive have published guidance on the Control of Major Hazards for processes involving programmable systems. One section of the guidance focusses on alarm systems that alert operators to plant conditions, such as deviation from normal operating limits which require timely intervention:

> "Alarm systems are not normally safety related, but do have a role in enabling operators to reduce the demand on the safety related systems, thus improving overall plant safety. However, where a risk reduction of better than 10-1 failures on demand is claimed then the alarm system, including the operator, is a safety related system which requires a suitable safety integrity level (SIL 1 or SIL 2 as defined by BS IEC61508)… For all alarms, regardless of their safety designation, attention is required to ensure that under abnormal condition such as severe disturbance, onset of hazard, or emergency situations, the alarm system remains effective given the limitations of human response. The extent to which the alarm system survives common cause failures, such as a power loss, should also be adequately defined".

Write a brief technical report describing the problems that arise during the validation and verification of software in alarm systems for safety-critical applications.

[10 marks]

[seen/unseen problem]

As mentioned, alarm systems are not normally considered to be safety critical because automated functions should be provided to prevent accidents given the relatively high, presumed probability of human error (1 in 10 from the HSE guidance). Validation focuses on the utility of key functions. The validation of alarms is difficult because it is non-trivial to determine all of the abnormal situations that might arise in complex systems, especially when these may be modified over the lifetime of an application. There can also be problems in validation when overload from too many warnings will swamp operators (up to 5 marks).

Conversely, verification ensures the correctness of systems. Verification is hard to demonstrate when abnormal conditions cannot be simulated without risk to the public. This is compounded by the need to recreate operating conditions to accurately determine human intervention. For example, the operators of a production process may respond in different ways at the end of a long shift compared to the beginning, they may also respond differently if they really believe that safety as at risk rather than participating in an exercise or drill (up to 5 marks).

In addition, there are a host of more technical requirements that should be satisfied during the testing of alarm systems under the COMAH requirements:

- "The alarm system should be designed in accordance with IEC 61508 to SIL 1 or 2, with the designated reliability;
- The alarm system should be independent from the process control system and other alarms unless it has also been designated safety related;
- The operator should have a clear written alarm response procedure for each alarm which his simple, obvious and invariant, and in which he is trained;
- The alarms should be presented in an obvious manner, distinguishable from other alarms, have the highest priority, and remain on view at all times when it is active;
- The claimed operator workload and performance should be stated and verified". (2 additional marks for any of these issues, up to a total of 10 marks).

2.

a) The following diagram has been adapted from UK Rail industry guidance on degraded modes of operation. Use the components of the model to analyse any of the accidents or incidents that have been discussed during the course.

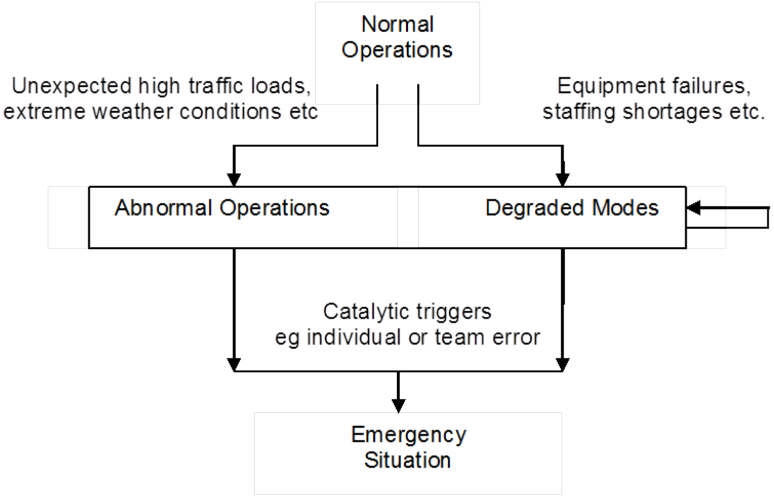
Normal Operations
Unexpected high traffic loads, extreme weather conditions etc
Equipment failures, staffing shortages etc.
Abnormal Operations
Degraded Modes
Catalytic triggers eg individual or team error
Emergency Situation

[5 marks]

[seen/unseen problem]

This model has been presented in the lectures and used to explain a number of accidents. One example might be the Linate runway incursion where an old analogue surface movement radar system gradually became unreliable (degraded modes), sectional lighting failed (degraded modes) and the runway/taxiway markings were inconsistent (degraded modes). On the night of the accident there was heavy fog (abnormal operations), heterogeneous traffic (abnormal operations) and crews that lacked the correct qualifications (abnormal operations). The crew of the Cessna took an incorrect route under low visibility conditions (trigger) and the lack of a digital surface movement system hindered ATC identification of the potential incursion. (2 marks for the description of the accident and 3 marks for the identification of at least one degraded mode, one abnormal aspect of operation and a catalyst/trigger).

b) Redundancy is often cited as a powerful means of maintaining safety under degraded modes of operation. Briefly describe the following forms of redundancy:

- Data redundancy
- Temporal redundancy
- Single and Multi-version redundancy
- Hot and cold redundancy
- Triple modular redundancy

[5 marks]

[unseen problem]

One mark to be awarded for each correct answer.

Data redundancy takes a number of different forms – it can be encoding – for example using a parity check or may involve the use of several data sources to cross check the input into a computation. Pre-computed or known values may also be used to cross-check the output from particular processing stages.

Temporal redundancy ensures that more time is available that might otherwise be necessary to perform a computation. If an error or failure is noted then a process can be restarted and completed without missing a time critical deadline – for instance using recovery blocks.

Single and Multi-version redundancy; single version redundancy uses elements within a software module to improve the reliability of that component, using fault detection, containment and recovery mechanisms. Multi-version redundancy employs several different versions of a module to provide alternate means of computing a result if one module fails.

Hot and cold redundancy; hot redundant systems are running in parallel with a primary system while cold stand-by systems have to be started and brought up to an appropriate point at which to resume computation.

Triple modular redundancy allows for three processing elements each running in parallel, we covered multi-layer TMR so they might also refer to three redundant voting elements. I suspect some will produce a diagram but given the small number of marks for each element of the question this would waste time and is not necessary.

c) Identify at least five limitations with the use of software redundancy as a means of ensuring the safety of complex systems.

[10 marks]

[seen/unseen problem]

2 marks per limitation with up to 2 additional marks for the quality of the prose and the analysis. These is a maximum of 10 marks in total.

Software redundancy is hard to achieve because:

- Logical errors. Software redundancy is fundamentally different from hardware redundancy – software fails due to logical errors rather than stochastic fabrication or construction issues. This means that simply having two versions of the same program will not necessarily increase reliability because they will fail in exactly the same way if they contain the same bug.

- Cost. The typical solution to logical errors in software redundancy is to use N-version programming where N independent teams each write code to meet the same requirements and some form of comparison (voting) is used to determine the overall result. However, the costs of software development often determine the overall costs of complex safety critical systems so there may be insufficient budget to use this approach except in very limited areas (there are also issues of complexity that interact with costs – see below).

- Complexity. Some people argue that it is better to spend the software budget on one good piece of code rather the N poor versions of a module. Not only does N version programming stretch the available budget but it also introduces additional complexity through the comparison routines – hence, N-version programming creates a complexity that may itself introduce new failure modes.

- Algorithmic similarity. If redundant diverse software is used then there is no guarantee that they will not contain the same bugs. Many logical errors cluster in the most complex areas of an algorithm – hence it is common to lose a bug that is repeated in the same place in two or more modules because testing makes comparisons between two or more flawed routines.

- Other associated technical issues. Solutions may raise a number of further limitations including physical constrains on memory and processing power – for example in satellite systems. Other issues surround problems in synchronization and timing between redundant software. It can be difficult to eliminate errors in common operating system, communications and network components or COTS products where the source code may not be available and where there may not be any suitable or diverse alternatives.

3.

a) Briefly explain the differences between transient, intermittent, partial and total failures. Which of these different failure modes can be caused by electromagnetic interference with software systems?

[5 marks]

[unseen problem]
Transient faults occur and then may never recur. An example is the electromagnetic interference which affects computation systems from faulty car starter motors. The problem would not recur once the car is driven away. An intermittent fault is one that occurs and recurs at different intervals of time. An example is the EMI that occurs when a local factory starts up a production process; it would not occur all the time but would recur. Partial failures are ones that do not affect all the functionality of a system. This might be the case when EMI causes arbitrary bit flips that are corrected at a higher level within an application program. In contrast, total failures involve the destruction of a system and could be the result of an intense burst of EMI – including some military weapons. From this analysis it is apparent that EMI can result in any of these four failure modes (0.5 marks for each definition, up to 3 additional marks for the analysis of EMI).

b) The European Electromagnetic Compatibility (EMC) 2004/108/EC came into force on 20 July 2007; products must not generate electromagnetic pollution. Equipment must also be resilient to interference. 2004/108/EC does not state the maximum levels of emissions or resilience. In practice, however, companies must develop tests across five different areas:

- Radiated emissions - Checks to ensure that the product does not emit unwanted radio signals;

- Conducted emissions - Checks to ensure the product does not send out unwanted signals along its supply connections and connections to any other apparatus;

- Radiated susceptibility - Checks that the product can withstand a typical level of radiated electromagnetic pollution;

- Conducted susceptibility - Checks that the product can withstand a typical level of noise on the power and other connections.

- Electrostatic discharge - Checks that the product is immune to a reasonable amount of static electricity.

Explain how safety cases can be used to record the relationship between evidence and the arguments that might be used to convince a national regulatory authority that a particular product meets the requirements of 2004/108/EC.

[5 marks]

[unseen problem]
In the course, we have discussed the use of Goal Structuring Notation (GSN) as a means of developing safety cases. In GSN, a network is created to describe the links between a safety claim and the evidence that is used to support that claim (1 mark). In this case, the top level assertion would be that the product is acceptably safe to operate; a contextual node might be used to refer to the requirements of 2004/108/EC (2 marks). Individual lines of argument could refer to each of the five tests listed above (1 mark). The results derived from those tests would be linked to these arguments in the manner illustrated by the following figure (1 mark – a diagram is not essential).

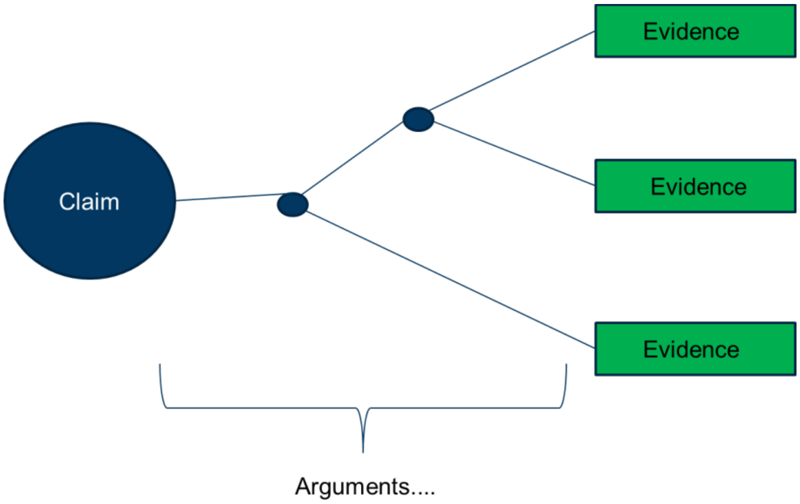

c) There are a host of other directives that deal with EMC. These include 93/42/EC on Medical Devices and 95/54/EC on automotive applications. Under European Law, the areas addressed by these more specific directives are excluded from the provisions of 2004/108/EC.

You have been hired by a company developing a new family of safety-rated processors. Write a brief technical report for company management explaining the technical challenges that these different directives create when trying to sell hardware across different European industries. Your report should also identify ways to address the problems that different industry directives create.

[10 marks]

[unseen problem]

We have discussed different approaches to software development across industries with different standards and regulatory bodies in the lectures. This question extends the scope of the discussion to consider hardware requirements. 2004/108/EC is a general requirement that does not focus exclusively on safety-critical applications. Hence, the additional standards extend the requirements of the five different forms of test mentioned in part b. I do not expect answers to tell me what those additional requirements are but a first class answer will refer to the modular nature of safety cases and GSN structures. For example, new evidence of more stringent tests can be linked into the nodes shown in the previous diagram. Additional tests would appear as new lines of argument. (2 marks for identifying the generic nature of 2004/108/EC, 2 marks for mentioning the more stringent nature of the alternate standards, 2 marks for mentioned the modular use of GSN through changes in evidence nodes and 2 marks for identifying the need for additional tests and lines of argument in the GSN). A first class answer might also mention the use of mappings between a company testing programme used across all industries and the requirements for particular standards. In this case, the processor might be internally tested against a range of criteria represented by a generic GSN – if the processor was to be sold to a new market that the generic GSN would be adapted and the mapping retained for any subsequent products to minimize future work.

4. NASA Software Safety Standard (NASA-STD-8719.13B) states that:

"Software shall be classified as safety-critical if it meets at least one of the following criteria:

a. Resides in a safety-critical system (as determined by a hazard analysis) AND at least one of the following apply:
   i. Causes or contributes to a hazard.
   ii. Provides control or mitigation for hazards.
   iii. Controls safety-critical functions.
   iv. Processes safety-critical commands or data (see note 4-1 below).
   v. Detects and reports, or takes corrective action, if the system reaches a specific hazardous state.
   vi. Mitigates damage if a hazard occurs.
   vii. Resides on the same system (processor) as safety-critical software (see note 4-2 below).
b. Processes data or analyzes trends that lead directly to safety decisions (e.g., determining when to turn power off to a wind tunnel to prevent system destruction).
c. Provides full or partial verification or validation of safety-critical systems, including hardware or software subsystems.

a) Use the previous paragraph to contrast the approach to safety-critical software engineering embedded within NASA's STD-8719.13B with key concepts in IEC61508.

[10 marks]

b) NASA-STD-8719.13B is being revised – extend your answer to part a) to suggest ways in which the concept of safety integrity levels in IEC61508 might be integrated into future revisions of the NASA Software Safety Standard.

[10 marks]

[structured essay/unseen problem]

This is an open-ended essay with two parts, as illustrated above. I will tell the class that there is a question on NASA software safety standards but not that it is about their relationship with IEC61508.

It is clear from the paragraph cited above that NASA-STD-8719.13B focuses on the direct role that software can play in creating hazards (2 marks) – for example, bugs that might affect safety-critical processes and data. This creates technical problems because conventional forms of risk assessment focus on the likelihood and consequence of failures (2 marks). It is difficult to talk about the probability of bugs – for the reasons mentioned in the sample solution to question 2c). Code does not suffer the same stochastic failure rates that can be identified for hardware (2 marks). Techniques such as fault injection can be used to assess whether testing has uncovered particular known faults. Other approaches, including Leveson's Software Fault Trees have not been widely applied in industry although they support testing rather than risk assessment. In the course, we have covered the Musa formula for software reliability and pointed to the difficulty in validating many of the terms that it contains – it is not necessary to repeat the formula in this solution (2 marks).

In contrast, IEC61508 focuses more narrowly on the hazards associated with Equipment Under Control (2 marks) – this avoids the need to quantify failure rates for software components. Instead, a Safety Integrity Level is calculated for the systems that are employed to mitigate the risk from Equipment Under Control (2 marks). If the probability and consequence of failure for EUC is high, then the protective systems will inherit a high SIL. These mitigation measures include software components. The higher the SIL then the more exhaustive are the techniques to be used to ensure the correctness of any mitigation (2 marks). 61508 focuses on EUC because software will cause very limited damage if it is not used to influence the behavior of Equipment Under Control.

The second part of the question asks for ways in which the concept of safety integrity levels in IEC61508 might be integrated into future revisions of the NASA Software Safety Standard. This is an open exercise and I will look for the feasibility and originality of the proposals – as appropriate for a level H module. One technique might be to develop a two-stage programme where an initial analysis depends upon the use of SILs following 61508 (2 marks). The second stage might then focus on the consequences of any failure to mitigating software using the approach in 8719.13B (2 marks). A first class answer would point to the costs associated with this approach and to the technical confusion that might arise in any hybrid technique (2 marks). A similar solution would be to use the template in 61508 and then associated the recommended techniques for analysis and testing in 8719.13B with the recommended methods for particular SILs (2 marks). 8719.13B talks briefly about assurance levels but does not formalize them in the way that has been described for 61508.