Xday, XX May 2004.

9.30 am - 11.15am

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) What is Reliability Centred Maintenance?

[3 marks]

b) The European Union recently funded a project into a Reliability Centred Maintenance Approach for the Infrastructure and Logistics of Railway operations (RAIL). This project used risk assessment to guide the monitoring and correction of potential faults across European rail networks. As part of the RAIL project, the team developed a two-stage approach in which the first step was to identify the criticality of network components using the following table:

Factor	Description	Value 1	Value 2	Value 3	Value 4
Technology	Kind of technology on the line or section.	Mechanical	Electro-mechanical	Electric	Electronic
Traffic density	Number of circulations per day.	[1,20]	[21,60]	[61,200]	>200
Revenues	Income from exploitation of network.	Low	Medium	High	Very High
Availability	Number of hours that the line must be available per day.	Low	Medium	High	Very High
Exploitation	Number of passengers or dangerous freight.	Low	Medium	High	Very High
Maintainability	Maintenance process complexity.	Low	Medium	High	Very High
Costs	Costs associated with maintenance.	Low	Medium	High	Very High
Environmental risk	Risk of environmental damage generated by an installation failure.	Low	Medium	High	Very High
Safety risk	Risk of people damage generated by an installation failure.	Low	Medium	High	Very High

i. Briefly explain why this table relies almost entirely on qualitative criteria for criticality assessment.

[4 marks]

ii. Briefly explain why computational failures are only considered indirectly within this table.

[3 marks]

c) The second stage of the RAIL approach focuses on each of the high criticality network components identified using the table in part b of this question. Each of these components is then subjected to a Failure Modes Effect and Criticality Analysis (FMECA). The severity (R) of any failure is defined by the following equation:

$$R = (S+C+D)\frac{P}{MTBF}$$

.

Where S is a numerical measure of safety, C is a measure of costs, D is a measure of punctuality, P is the probability of the failure and MTBF is the mean time between each occurrence of that failure. Recall that there may be many ways, or modes, by which a failure can occur. Briefly comment on whether it is possible to use this formula to consider the severity of a failure where the loss of a programmable system is one of the potential failure modes.

[10 marks]

2.

a) The Boeing 777 is unusual in that the decision was taken early in development to write as much as possible for the software in Ada. Honeywell developed the primary flight controls and purchased DDC-I, Inc.'s Ada Compiler, using it as the front-end for Honeywell's symbolic debugger. The two companies worked together for eighteen months to build the compiler's debugger and the back-end, targeted to an AMD 29050 microprocessor. Briefly explain why so much care was taken to distinguish between the high level language, the compiler and the analysis tools, and the target processor.

[3 marks]

b) One of the other 777 subcontractors, Sundstrand, chose a compiler from Alsys. This generated code for an Intel 80186 microprocessor that relies upon the Certifiable Small Ada Run Time (CSMART) executive. Members of the development team argued that this enables them to reuse code. For example, the Gulfstream V business jet and the Comanche helicopter Sundstrand's library of common generic packages written for the 777. Briefly explain the importance of testing the CSMART executive to support such reuse.

[5 marks]

c) On the 777, three processors provide triply redundant computation for the primary flight control system. Each Primary Flight Computer (PFC) receives data from three control buses. Each PFC only transmits on its associated bus. Each PFC channel contains three dissimilar processor lanes that use different processors and were developed using different Ada compilers. Each lane contains its own power source and has its own terminals to communicate with the buses. Explain how this architecture contributes to the overall safety of the Boeing 777 aircraft.

[12 marks]

3.

a) A recent study by the US Food and Drugs Administration examined 3,140 medical device recalls. This revealed that 242 (7.7%) were attributable to software failures. Of these, 192 (79% of softeware related failures) were caused by software defects that were introduced when changes were made to the software after its initial production and distribution. The majority of these updates stemmed from ‘usability’ problems. The FDA concluded that ‘software validation and other related good software engineering practices … are a principal means of avoiding such defects and resultant recalls’. Briefly explain why regulators now typically rely on certifying the development process or ‘engineering practices’ rather than individual safety-critical systems.

[3 marks]

b) The Da Vinci system is the first fly-by-wire robotic aid to be approved for surgical applications by the Food and Drugs Administration. Briefly explain why conventional forms of black-box testing may not provide sufficient assurance of the safety of such an application.

[6 marks]

c) You have been appointed as a Safety Manager working in a company that produces a programmable ventilator. These devices are used to help anesthetize patients. You have just been sent a report from the Food and Drugs Administration that describes an incident involving one of your devices and the attempts of your colleagues to help the clinician using the device:

THE VENTILATOR ON THE ANESTHESIA FAILED WITH AN ERROR MESSAGE "GAS INLET VALVE FAILURE." PATIENT WAS VENTILATED BY HAND AS PREPARATIONS WERE MADE TO SWAP OUT THE ANESTHESIA MACHINE. the SERVICE REP WAS CONTACTED, HE TELEPHONED INTO OPERATING ROOM. HE WALKED ANESTHESIA ATTENDING THROUGH A SERVICE PROCEDURE TO "BLOW OUT" GAS INLET VALVE. VENTILATOR WORKED AFTER THIS AND FOR REMAINDER OF PROCEDURE. MACHINE WAS TAKEN OUT OF SERVICE AND VALVE IS BEING SENT BACK TO manufacturer. THERE ARE REPORTS OF OTHER RECENT SIMILAR INCIDENTS INVOLVING NEWLY INSTALLED ANESTHESIA MACHINES OF THE SAME MODEL

Using one of the incident analysis techniques introduced in this course, identify any lessons that might be learned from this incident report and explain how you would go about identifying any necessary corrective measures.

[11 marks]

4. To what extent is it acceptable to blame 'systemic failure' rather than operator or managerial 'error' as a cause of accidents involving programmable systems. Illustrate your answer with detailed references to two of the accidents that we have studied in this course.

[20 marks]

[end]

Sample Solutions

1.

a) What is Reliability Centred Maintenance?

[3 marks]

[Bookwork/Unseen problem] Reliability Centred Maintenance as the name suggests is an approach that uses risk assessment to guide maintenance operations. The technique will assess the criticality and frequency of a failure and together with the mean time to repair these estimates will be used to device appropriate maintenance schedules that will minimize the risk both of safety-related failures and of a denial of service provision.

b) The European Union recently funded a project into a Reliability Centred Maintenance Approach for the Infrastructure and Logistics of Railway operations (RAIL). This project used risk assessment to guide the monitoring and correction of potential faults across European rail networks. As part of the RAIL project, the team developed a two-stage approach in which the first step was to identify the criticality of network components using the following table:

Factor	Description	Value 1	Value 2	Value 3	Value 4
Technology	Kind of technology on the line or section.	Mechanical	Electro-mechanical	Electric	Electronic
Traffic density	Number of circulations per day.	[1,20]	[21,60]	[61,200]	>200
Revenues	Income from exploitation of network.	Low	Medium	High	Very High
Availability	Number of hours that the line must be available per day.	Low	Medium	High	Very High
Exploitation	Number of passengers or dangerous freight.	Low	Medium	High	Very High
Maintainability	Maintenance process complexity.	Low	Medium	High	Very High
Costs	Costs associated with maintenance.	Low	Medium	High	Very High
Environmental risk	Risk of environmental damage generated by an installation failure.	Low	Medium	High	Very High
Safety risk	Risk of people damage generated by an installation failure.	Low	Medium	High	Very High

i. Briefly explain why this table relies almost entirely on qualitative criteria for criticality assessment.

[4 marks]

[Unseen/seen problem]. There are a number of answers to this question. On the one hand, some of these concepts are just difficult to quantify. There have been examples where agencies have been very inaccurate in estimating the cost of environmental damage, such as the clean up operation after the Exxon Veldes disaster. Similarly, it seems appropriate to consider the costs of human injury in broad terms. There are further reasons. In an international project, the precise value associated with injury or environmental damage can be culturally determined even with the EC. It can also be argues that by avoiding any reference to quantitative data, the results of this analysis are future-proofed against the inflation that may occur in the values associated with particular adverse events. It might be argued that this is a weakness for the quantitative data in this table, especially the number of operations per day which would be stretched in some areas of the ThamesLink commuting network.

ii. Briefly explain why computational failures are only considered indirectly within this table.
[3 marks]

[Unseen problem]. This table deals with the consequences of a rail network failure. Just as standards such as 61508 focus on the consequences of failures involving Equipment Under Control and not the programmable systems themselves. It is interesting, however, that the RAIL technique does associate a higher consequence with the failure of electronic systems. This may be due to associated knock-on effects because these are likely to be part of wider and more integrated systems.

c) The second stage of the RAIL approach focuses on each of the high criticality network components identified using the table in part b of this question. Each of these components is then subjected to a Failure Modes Effect and Criticality Analysis (FMECA). The severity (R) of any failure is defined by the following equation:

$$R = (S+C+D)\frac{P}{MTBF}.$$

Where S is a numerical measure of safety, C is a measure of costs, D is a measure of punctuality, P is the probability of the failure and MTBF is the mean time between each occurrence of that failure. Recall that there may be many ways, or modes, by which a failure can occur. Briefly comment on whether it is possible to use this formula to consider the severity of a failure where the loss of a programmable system is one of the potential failure modes.

[10 marks]

[Unseen problem] This question centers on whether it is ever possible to derive an accurate estimate for the probability of a failure involving a programmable system, P. Software is not like hardware in that failures are not easily characterize by probability distributions. If it is thought likely that there is a bug then it should be removed, whereas it may not always be possible to remove possible defects in a hardware device. The fact that a piece of code has executed successfully for some period is no reliable prediction of its future operation. Slight changes in the operating environment can expose new sections of code that may cause the system to fail with probability 1 Similar comments can be made about identifying the possible failure modes of systems that rely on safety-critical software. There are so many possible failure modes that it can be difficult to test or even identify them all. Software fault trees and similar inspection methods can be used but may be unreliable over many thousands of lines of code. In any event, it can be difficult to apply FMECA style calculations that are based on the summation of risk equations calculated across a wide range of possible failure modes involving programmable systems.

2.

a) The Boeing 777 is unusual in that the decision was taken early in development to write as much as possible for the software in Ada. Honeywell developed the primary flight controls and purchased DDC-I, Inc.'s Ada Compiler, using it as the front-end for Honeywell's symbolic debugger. The two companies worked together for eighteen months to build the compiler's debugger and the back-end, targeted to an AMD 29050 microprocessor. Briefly explain why so much care was taken to distinguish between the high level language, the compiler and the analysis tools, and the target processor.

[3 marks]

[Unseen/seen problem] In theory, by distinguishing between each of these components it will be possible for suture projects to replace some elements of the development environment with minimal changes to the rest of the architecture. In particular a change in target processor from the AMD29050 should only force revisions in the back-end elements that are targeted for that platform. Other solutions might focus more on the benefits of re-using the existing symbolic debugger that staff at Honeywell will already be familiar with. In any event, it can be argued that each of these components must be carefully considered in the development of a safety-critical programmable system.

b) One of the other 777 subcontractors, Sundstrand, chose a compiler from Alsys. This generated code for an Intel 80186 microprocessor that relies upon the Certifiable Small Ada Run Time (CSMART) executive. Members of the development team argued that this enables them to reuse code. For example, the Gulfstream V business jet and the Comanche helicopter Sundstrand's library of common generic packages written for the 777. Briefly explain the importance of testing the CSMART executive to support such reuse.

[5 marks]

[Unseen problem]
The use of the CSMART executive can be compared to a safety kernel in which it defines a subset of ‘safe’ operations that can be performed on a target processor. By restricting valid language constructs, CSMART also reduces the burden of verification by excluding those features that are difficult to demonstrate are safe in the same way that Praxis and the SPARK Ada subset reduce language constructs to those that are amenable to formal verification. The importance of testing the CSMART executive is that any problems in the implementation of this element would yield potential problems for all of the higher level applications that run upon it. Thus if the CSMART kernel were to be ported to another target processor, it would be possible to increase our assurance in any higher level libraries that were also ported with the executive. The argument here is very similar to that used in part a) by defining appropriate interfaces between software components we can increase modularity. This is an essential commercial consideration for safety-critical systems where even relatively modest changes can incur a heavy burden of verification. One possible limitation with this approach is that the CSMART executive is closely tied to the vendor Alsys.

c) On the 777, three processors provide triply redundant computation for the primary flight control system. Each Primary Flight Computer (PFC) receives data from three control buses. Each PFC only transmits on its associated bus. Each PFC channel contains three dissimilar processor lanes that use different processors and were developed using different Ada compilers. Each lane contains its own power source and has its own terminals to communicate with the buses. Explain how this architecture contributes to the overall safety of the Boeing 777 aircraft.

[12 marks]

[Unseen problem/bookwork] Several different approaches can be taken here. I would expect most people to reproduce the diagram of triple modular redundancy from the lecture notes; however, this is not essential for a good mark. The extract is taken from Boeing technical documentation and is potentially ambiguous in the description of the dedicated output bus with the three dissimilar processor lanes on each channel. However, the class has been warned that they may need to make assumptions about such issues and I’ll be generous with the marking. The key issue is to represent the redundancy in processing elements and in the

communications infrastructure. Good solutions should also address the use of diversity in the selection of Ada compilers – N version programming might be mentioned here although we are not told that there were different development teams using each of these different compilers. First class answers might go on to question whether N version programming can be cost effective and whether the use of diverse compilers will contribute much to overall safety when many bugs are introduced through flaws in the common requirements that can be shared between development teams.

3.

a) A recent study by the US Food and Drugs Administration examined 3,140 medical device recalls. This revealed that 242 (7.7%) were attributable to software failures. Of these, 192 (79% of softeware related failures) were caused by software defects that were introduced when changes were made to the software after its initial production and distribution. The majority of these updates stemmed from 'usability' problems. The FDA concluded that 'software validation and other related good software engineering practices … are a principal means of avoiding such defects and resultant recalls'. Briefly explain why regulators now typically rely on certifying the development process or 'engineering practices' rather than individual safety-critical systems.

[3 marks]

[Unseen problem] The FDA statistics are interesting because they indicate that very few are directly attributable to software. This situation may be changing, however, it may also reflect the way in which software failures may be difficult for end-users to diagnose without detailed knowledge of the system architecture. The FDA data also illustrates that most defects were introduced after modification and this is instructive because it suggests that the modification process may be more error prone or less well regulated than the initial development process. The main part of the answer to this question should focus on the difficulty of testing programmable systems. This will only ever demonstrate the presence of bugs but not their absence. If we cannot prove that a product is free from bugs then we can focus on ensuring that the development process finds as many faults as possible.

b) The Da Vinci system is the first fly-by-wire robotic aid to be approved for surgical applications by the Food and Drugs Administration. Briefly explain why conventional forms of black-box testing may not provide sufficient assurance of the safety of such an application.

[6 marks]

[Seen/unseen problem] Black-box testing assumes that the evaluator has little or no knowledge of the implementation details of the system. This can be especially effective when verifying properties of an interface between software modules. This approach is less effective when testing such a complex system as the Da Vinci robot. For instance, it will typically not be possible to test every possible use of the system in all possible working environments or configurations. Having some knowledge of the internal details of the system can help to guide testing towards particularly hazardous operations or states. One area where black-box testing is essential, however, is in the evaluation of a user interface where potential operators must often be assumed not to possess any detailed knowledge of an underlying implementation.

c) You have been appointed as a Safety Manager working in a company that produces a programmable ventilator. These devices are used to help anesthetize patients. You have just been sent a report from the Food and Drugs Administration that describes an incident involving one of your devices and the attempts of your colleagues to help the clinician using the device:

> THE VENTILATOR ON THE ANESTHESIA FAILED WITH AN ERROR MESSAGE "GAS INLET VALVE FAILURE." PATIENT WAS VENTILATED BY HAND AS PREPARATIONS WERE MADE TO SWAP OUT THE ANESTHESIA MACHINE. the SERVICE REP WAS CONTACTED, HE TELEPHONED INTO OPERATING ROOM. HE WALKED ANESTHESIA ATTENDING THROUGH A SERVICE PROCEDURE TO "BLOW OUT" GAS INLET VALVE. VENTILATOR WORKED AFTER THIS AND FOR REMAINDER OF PROCEDURE. MACHINE WAS TAKEN OUT OF SERVICE AND VALVE IS BEING SENT BACK TO manufacturer. THERE ARE REPORTS OF OTHER RECENT SIMILAR INCIDENTS INVOLVING NEWLY INSTALLED ANESTHESIA MACHINES OF THE SAME MODEL

Using one of the incident analysis techniques introduced in this course, identify any lessons that might be learned from this incident report and explain how you would go about identifying any necessary corrective measures.

[11 marks]

[Unseen problem] As with previous questions, a range of solutions is possible and this sample answer only highlights my initial analysis of this complex mishap. It is also important to reiterate that more information may be required before a safety manager would make any intervention. I would award additional marks to any student who points this out. I have introduced a range of root cause analysis techniques such as MORT, PRISMA, TRIPOD and ECF. These could be used and , if so, they will look beyond the operator and the machine failure to higher levels of management. This analysis could then be linked to a solution to Question 4 on systemic failures (see below).

One of the key observations here is that the service rep was forced to provide backup guidance to clinicians on the telephone in the operating room. The potential hazards associated with such a situation are difficult to underestimate. For instance, problems might have arisen if the service rep had been unavailable or if they had not been able to resolve the problem. In any event, the need to resolve this error message will have consumed valuable resources in terms of clinician's time and attention during the procedure. Either the physician should have been trained on the blow out procedure for the gas inlet valve or the system should have displayed sufficient guidance for them to perform this procedure without recourse to the telephone.

A particular concern here would be that the FDA had received reports of similar incidents. As a safety manager, you would be expected to have noticed such a pattern and ideally to have acted upon them before this piont. If this were the case then it would be important to assess whether any recommendations from previous adverse events had been acted upon. Alternatively, if the safety manager had not received the reports then they must trace back through the reporting chain to establish where this critical information about previous similar events had been lost.

Further lines of analysis could focus on the relationship between the equipment supplier/distributor and the manufacturer of the valve that failed.

4. To what extent is it acceptable to blame 'systemic failure' rather than operator or managerial 'error' as a cause of accidents involving programmable systems. Illustrate your answer with detailed references to two of the accidents that we have studied in this course.

[20 marks]

[Essay]

This is an essay style question that is intended to give plenty of scope for first-class answers. There are many different approaches. One line of argument would be to accept the suggestions put forward by Leveson in her work on STAMP. This has been introduced in the lectures and involves a control-model approach to incident analysis. By focussing on the constraints that hold between many diverse agents, systems and organisations this technique guides analysts towards a systemic view of failure in which it is unlikely that a single individual would be blamed for any failure. In this interpretation, bugs are introduced through inappropriate development practices, poor project management, ineffective regulatory intervention etc. Similarly, the failure to detect and address a bug prior to deployment can be seen as the result of inadequate testing plans, a lack of independent verification, insufficient funding from project management and so on. Other researchers, including Reason, have had a profound influence on UK government thinking in promoting this systemic view. This view is also revealed in many of the comments by inquiries and investigations, such as the Cullen report into the Ladbroke Grove rail crash.

I have introduced a contrary position in this course. By looking too broadly at systemic factors we may overlook an individual's opportunity to intervene and address the causes of failure. This 'opportunity' assessment lies at the heart of proportionate blame. It can be applied at the level of an individual programmer but also at the level of project management. A key idea here is that a number of recent managers have sought to shift responsibility for their involvement in recent mishaps by identifying the 'systemic causes' of the accident. In particular, I have spoken about the causes of the Mars Surveyor mission failures. I have also pointed to some of the arguments that were used in the defence of executives involved in the non-safety related Enron litigation. In these cases, attempts have been made to move responsibility towards the regulators who were responsible for monitoring the conditions in which they operated. In my view, this can stretch the credibility of systemic views of failure beyond the pale.

There is a middle position in which the proponents of systemic views are asked to prioritise their recommendations for avoiding future adverse events. This is a key practical outcome of any investigation. With finite resources, we need to know where to focus our attention. This is less to assign blame than it is to avoid future failures. A systemic view might be acceptable to analyse the causes of accidents but unless it can be used to guide intervention then it may be too general to be of practical use.