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.

**Sample Solutions**

1.

a) The following image shows part of the network schema for the Zurich Air Traffic Control system that was involved in the recent Überlingen mid-air collision.

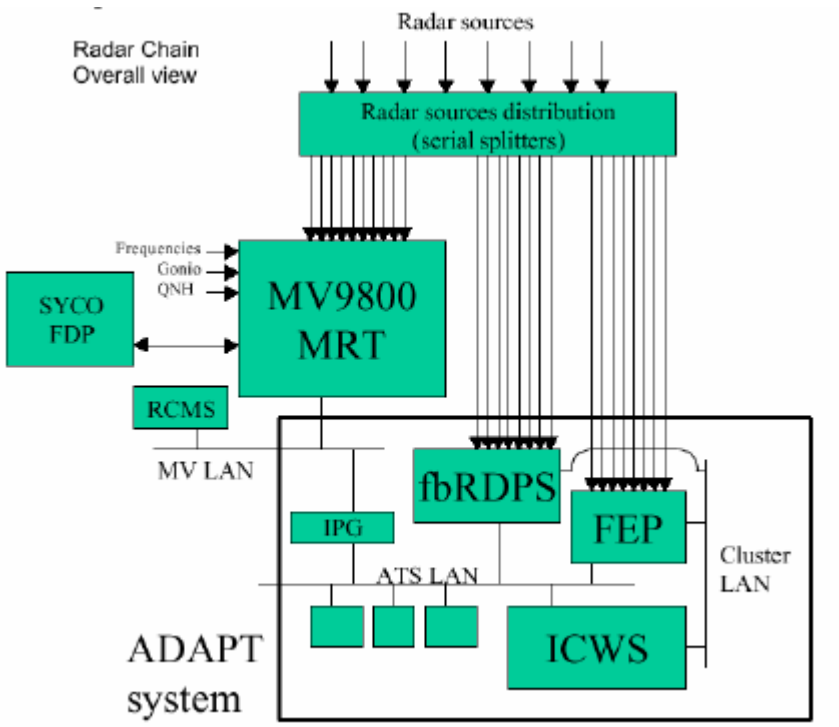


The German BFU investigation into the collision describes the network as follows. "The radar data processing system consists of three main Thomson MV9800 computers. One MV computer generally is in "hot operation", the second is on "hot standby" and the third is used for test purposes and software development. This one is on "cold standby". The MV9800 processes the incoming signals and correlates them automatically with the flight plan data supplied by the flight data processing system (SYCO). The MV9800 is controlled and monitored by the RCMS (Remote Control & Monitoring System). The radar images correlated in the MV computer are sent via the IPG-IPS Interface Gateway to the monitors at the controllers' workstations (ICWS – Integrated-Controller-Workstation). Legal Recording records the radar raw data and the output of the radar data processed by the MV9800".

Briefly describe the way in which hot and cold standbys were used to increase the reliability of the Zurich Air Traffic monitoring system before the Überlingen accident.

[3 marks]

[unseen problem] The primary Mv9800 runs continually in 'hot' operation. The secondary, redundant MV9800 is used as a 'hot' standby. In other words, it is running and ready to take over from the primary computer should it fail. The mechanisms used to trigger this switch are not described. However, the secondary system may use watch-dog processes to monitor the health of the primary system or they may even be some basic manual checks. Similarly, the state or operation replication on the backup are not described and this will determine the lag before the secondary machine can come on-line. In this sort of environment, I would expect tight synchronization. The third machine is a cold bacup. In other words, it may have the same basic configuration as the other two machines but it will not routinely shadow the primary MV9800. As mentioned in the excerpt it will have a support role and only be called on in the (arguably) unlikely event of a double primary and secondary failure.

b) At the time of the accident, maintenance was being conducted on the system. In normal operation, the controller workstation (ICWS) uses flight plan data from the SYCO system and direct radar information so that the controller can see the intended route for each of the aircraft. However, the maintenance activity affected the flight plan system and so each controller had to manually work out the flight plan associated with each radar target on his or her screen. Using any of the techniques that we have looked at on this course, briefly conduct a high-level risk assessment for such maintenance activities and the potential impact on the controller.

[7 marks]

[Unseen problem] The key issues here are to identify the potential hazards and then to assess the likelihood and consequences of those hazards. We have looked at FMECA, Fault Trees and very briefly at HAZOPs but they could use other techniques from the open assessment. A really good answer might look at the next part of the question and identify the main hazard as being an airborne collision – even if they did not read the full question before starting then they could probably guess at this by the references to the accident earlier on. Anyway, most of the main techniques will identify ways in which a collision might occur. In a fault tree one sub-branch might consider a failure by the controller to ensure adequate separation. This could happen from human factors problems in manually correlating the flight plan with the radar location. HAZOPs and the THEA variant that we looked at in class would take this further by looking at ways in which this might occur using the guidewords such as 'too soon', 'too late' etc. Using one of these approaches should be straightforward for the simple problem decomposition mentioned in the question. More able students should write something about the problems of quantifying risk – the consequence of mid-air collision is fairly obvious at least in crude terms. The likelihood raises many questions. We know that this accident occurred so we could argue the likelihood is '1' – however, can we assess the future probability of a similar failure in the future? In the lectures we looked at Smith and Mosier but it would be a really great answer if anyone used this here.

c) The Überlingen accident occurred at night when two controllers should have been on-duty. Traffic was light so one controller went to rest in another room. The remaining Air Traffic Control Officer had to move between two workstations several meters apart. He had to guide an A320 to land at Friedrichshaven airport and coordinate the movements of a TU-154 and a B757. The maintenance work prevented his computer systems from generating a visual warning (Short Term Conflict Alert) for the TU154 and the B757. The controller apparently also failed to hear the aural warning that was generated. He eventually noticed the problem and asked the TU154 crew to descend. The TU154 crew were told by their on-board computer system (TCAS - Traffic alert Collision and Avoidance System) to climb but they instead followed the controller's instructions. The controller thought they had resolved the problem and moved to the other control station to help the A320. Meanwhile, the B757 descended into the TU154 following instructions from TCAS. (Hint: The TCAS software communicates between each plane to coordinate the instructions to each crew. The TU154 was told to climb by TCAS and the B757 to descend. However, international regulations allow controller's to over-ride TCAS if they think they can avoid a collision).

Briefly explain the role that the controller's situation awareness might have played in the causes of this accident. Comment on whether or not you feel that this was a 'systemic' failure.

[10 marks]

[Unseen problem] In the lectures we covered Flach and Wicken's model of situation awareness. I would not expect students to reproduce this in an exam but they might remember some of the key concepts such as the importance of information from the environment and the impact of anticipation on scoping problem solving. In this accident, the fact that the controller had to perform his colleague's duties, the physical space between the workstations, the dual task of on-route and approach control all combine to stretch his cognitive and perceptual resources. The lack of computational support and the prompts provided by visual STCA warnings clearly would have compromised situation awareness. When he did notice the conflict he would have been rushing to resolve the problem. Hence, the possibly snap decision to order the TU154 decent without checking on the actions of the B757. He may have anticipated that they would hold their altitude or climbed hence did not check and moved back to the competing task. This is arguably a systemic

failure as it stemmed from the interaction of the TCAS software, the crews on each of the aircraft and the controller.

2.

a) The FIA are the governing body of Formula One racing. Their regulations require that "All software must be registered with the FIA, who check all the programmable systems on the cars prior to each event to ensure that the correct software versions are being used". Briefly explain the problems of using 'white box' testing to support the safety regulation of a competitive sport such as Formula One racing.

[5 marks]

[Seen problem/unseen problem] White box testing involves a detailed knowledge of the architecture and function of code. In a competitive environment like Formula One where the software controls most if not all technical subsystems then this code will include critical commercially sensitive material. The regulator wants to check that the code conforms to safety and performance standards and hence must be able to access the code. This opens dangers of unwarranted disclosure. Without access to underlying documentation and additional information, implied in black-box testing, it seems unlikely that FIA officials would be able to conduct thorough tests. Hence there is a clear tension at the heart of software scrutineering in Formula 1.

b) The process of checking that a car's software conforms to the Formula One guidelines causes a number of problems. There are many changes to algorithms during the racing season and the software contains numerous commercial 'secrets'. The FIA offer two choices.

> **Option 1:** The software can be checked, line-by-line at the start of the season. A copy of the machine code is then kept by the FIA so that it can be compared at any time during a Grand Prix.
>
> **Option 2:** The software is not inspected at the start of the season. The FIA can demand an upload of all software at any time for a 'full check' if deemed necessary.

Each team must choose option 1 or option 2 at the start of the season. Most have chosen Option 2. Compare and contrast the strengths and weaknesses of each approach as a means of ensuring driver safety in motor sport.

[7 marks]

[Unseen problem] Option 1 has considerable advantages for the scrutineers. They can invest considerable resources at the start of the season to examine and analyse the code used by each team. If there are safety issues then these can be identified and corrected not just within a team but between teams before the season begins. However, this approach effectively freezes the code that may be used in a car. This creates safety related problems because changes in the tracks or the performance of the vehicle can create dangers that rely on software to be corrected. It also prevents teams from innovating as they learn lessons about the handling of a car during the season. For example, if code is inspected and tested before the season begins then it will only have been used in simulated conditions and not during a real race. Hence, any bugs that may be identified or dangers that emerge during a race cannot easily be corrected. All of these reasons explain the popularity of option 2 with the current teams. However, this creates problems for the scrutineers who have to ensure that the code used during each race is both safe and legal. It is impossible for them to test all of the software used by every team at each race. Hence they must determine the optimal period between inspections to maximise their inspection resources and deter teams from introducing code that might otherwise not be passed as safe or within FIA guidelines. Better answers may question whether it is technically possible for these inspections to really look at the safety of the code involved. Greater reliance can arguably be placed on the team's motivation to look after their drivers and the cars rather than the deterrent effect of FIA inspections that focus more narrowly on the performance regulations.

c) A recent summary of this scrutineering process argued that for any car control systems, the software is inspected to ensure that "ALL the code is accessed". It went on to argue that "System software is inspected to ensure that uploading, de-compiling and comparison software do not eliminate or modify the control software in any way. Especially important are checks to ensure that it is not possible to download an illegal program into RAM which is overwritten and lost when the system is switched off and then on again" (Peter Wright – Formula One Technical Guide, Scrutineering). Briefly explain the technical difficulties of achieving these checks for complex, safety-critical systems.

[8 marks]

[Unseen problem] This question builds on the closing comments of the sample solution to the previous question. As mentioned, the FIA is charged with ensuring the safety of competitors and of ensuring that the performance regulations are followed to insure the long-term health of the sport. However, it is particularly difficult to meet these two objectives when software is involved. For example, it is unclear what is meant by a check to ensure that 'ALL code is accessed'. If this implies that every path is used within a complex control structure then this is technically infeasible for anything but the simplest of systems. If it means that there are no unused subroutines then this poses similar challenges depending on the nature of the control structures that are used. Identifying redundant code is easier but can also be complex depending on the language involved (eg more difficult in C with pointers to functions than in Ada). The subsequent comments about code alteration is interesting because it is a variation on some of the testing issues that have been mentioned in the course. Usually, you are concerned that if you test the source code then the machine code preserves the intended semantics of the high level language after compilation so that the test condition will continue to hold. The concern is that there are no bugs in the compiler etc. Here the quotation seems to imply testing ensures the translation or interpretation of code does not introduce any performance enhancing instructions that would not otherwise be visible to the scrutineers.

3.

a) Two devices can be defined to be electromagnetically compatible (EMC) when the electromagnetic disturbances produced by one device are not powerful enough to affect the operation of another, or when that equipment is protected enough from EM disturbances to continue working. Briefly explain why it can be difficult to use this definition to ensure that any particular device is safe from the hazards of electromagnetic radiation.

[3 marks]

[seen/unseen problem] EM radiation can lead to transient failure modes created by one-off interactions between devices. Car starter motors are the example that we used in the lectures. It is difficult if not impossible to show that your device is resilient against all possible external interference from all other devices. All that you can do is conduct a form of risk assessment and shield if necessary. Operational rules can be used to keep potential radiation sources away from your device (eg in aviation and medicine) but these rules have been violated and cannot be relied on.

b) The European EMC directive requires computer and light industrial equipment to function in electromagnetic fields up to 3V/m. Computers and equipment in heavier industry must operate in EM fields up to 10V/m. Environmental conditions also affect performance. For example, systems that integrate many different items of equipment are more susceptible to electromagnetic problems than might be predicted by examining each of their components. Briefly explain how you would demonstrate that a proposed system design would be safe against the hazards of electromagnetic radiation.

[7 marks]

[seen/unseen problem] It would depend on who you were acting for. If you are the customer of the equipment then the 'E' mark should be sufficient to ensure that a single device met the EC requirements. If you were a supplier then there is a legal obligation for you to get specialist help and demonstrate that your device meets the required standards. If you were using the equipment as part of a larger system or if you were using an 'E' marked device in a safety related environment then you should additionally conduct an EMC risk assessment. If there was a likely, high-consequence hazard caused by EM radiation then shielding should be considered or the acquisition of higher-rated specialist equipment. The theoretical performance of these devices should also be check against specific bench-marking tests in the eventual environment because reflective metal surfaces and other fixtures can amplify the effects of EM interference.

c) A recent report by the Institution of Electrical Engineers argued that it is increasingly difficult to keep potential sources of radiation, including mobile phones and Wireless LANs, away from sensitive safety-critical devices. The summary at the end of the report identified redundancy as a potential solution. Briefly explain the problems of using redundancy and diversity as protection against electromagnetic radiation.

[10 marks]

[Unseen problem] There are many answers to this. For example, redundant wiring looms can be used to carry signals that can then be compared at their destination to see if there has been any EM interference. However, this creates additional expense and design complexity, for instance in voting and check algorithms that can also introduce failure modes. Alternatively, data redundancy can be used at a higher level with checksums etc but again this comes with the costs of complexity. Multiple processors can be used but they add expense, weight, power consumption and additional hazards from the synchronization and coordination algorithms. For extreme environments, such as space, it is possible to buy specialist EM 'hardened' processors but the costs are often prohibitive and they lack many of the associated development tools that are essential in many domains. We described the architectures of several specialist processors for space and aviation applications during the course. Above all, EM radiation can act as a common failure mode for multiple redundant systems that use the same design and implementation mechanisms. Hence, higher levels of protection may require the additional costs of different development teams etc.

4. Explain the strengths and weaknesses of the ALARP ('as low as reasonably practicable') approach to design, as embodied within the Health and Safety legislation of various countries.

[20 marks]

[Unseen problem – with some discussion in class]

There are many different potential answers here. The key issues focus on what is meant by 'as low as reasonably practicable'. The assessment of whether a risk is as low as is reasonable is something that depends upon a range of values and attitudes within society. Different groups in society also have different attitudes towards particular risks. Hence, the courts often only reconcile these different views, too often in the aftermath of an accident.

It is also important to understand that attitudes towards risk vary greatly within society depending on the social benefits and costs that are associated with the activity. For example, the risks associated with road

and rail transport or with the use of fossil fuels versus nuclear energy are affected by a host of technical and non-technical factors that are difficult for engineers to assess.

Further problems arise when attempts are made to associate metrics with ‘reasonable’ safety related investments. For example, the attitudes towards ATP and other forms of rail signaling and protection systems have varied over the last decade. If we take crude measures like the cost per life saved by such equipment then a succession of accidents can force an organization to the point where an investment is suddenly justified because a large number of people have been killed in accidents that might have been avoided if the investment had been made earlier.

As mentioned, this is an essay style question where I am looking for some reasoned arguments and also some specific examples drawn in from the systems and industries that we have examined in the course. I think this is a deceptively hard question so I’m prepared to be quite generous.