Xday, XX May 2006.

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) On 14th August 2003, the company responsible for ensuring the reliability of electricity supply in part of the northern USA suffered a failure in its Real-Time Contingency Analysis (RTCA) software. This program monitored the current state of the power network and then simulated what would happen if failures occurred. The following extract from the official report into the 'blackout' identifies a number of key factors in the failure:

> "about 12:15, the RTCA produced a solution outside the bounds of acceptable error. This was traced to an outage of the Bloomington-Denois Creek 230-kV line, although it was out of service, its status was not updated. Line status information is transmitted by the data network and is intended to be automatically linked to the RTCA software. This requires coordinated data naming as well as instructions that link the data to the tools. The automatic linkage of this line's status had not yet been established. To troubleshoot this problem the analyst had turned off the automatic trigger that runs the RTCA every five minutes. After fixing the problem he forgot to re-enable it. Thinking the system had been successfully restored, the analyst went to lunch" (US-Canadian Presidential Report, 2003).

Explain the importance of data integrity as a cause of this failure.

[5 marks]

b) While the RTCA system was disabled, a second and unrelated software 'failure' affected Emergency Management System (EMS) software. The RTCA program, mentioned above, was used to predict the consequences of a failure in the electricity network. In contrast, the EMS software was used to warn companies when there actually was a failure in the electricity network. The EMS application ran on several servers, any one of which could host all of the EMS functions. However, the normal configuration was to have one server remaining as a "hot-standby". The primary server hosting the EMS failed, due either to "the stalling of the alarm application" or a data overflow on remote terminals that relied on EMS data. As intended, the alarm system application and all other EMS software running on the first server transferred to the back-up server. However, the alarm application "moved intact onto the backup while still stalled and ineffective". The backup server failed 13 minutes later.

Identify the particular strengths and weaknesses of software redundancy that are illustrated by this case study.

[5 marks]

c) The EMS server failures mentioned in part b) slowed the refresh rate on the operators' displays. EMS data was often nested in windows providing more general network information and so the EMS delays affected the operators' ability to refresh more general system information. These problems were compounded because the IT staff did not inform the operators that the EMS servers had failed.

Using the information provided in parts a), b) and c), explain the role that human factors played in both the immediate and longer-term (latent) causes of this safety-critical failure.

[10 marks]

[Cont.]

2.

a) Briefly explain the different relationships that exist between the terms ‘reliability’, ‘safety’ and ‘security’. (Hint: is a secure system necessarily safe?)

[4 marks]

b) Risk based techniques have been advocated for secure applications. For example, the following diagram shows the US government’s Sandia Labs use of fault trees for risk assessment of security threats.

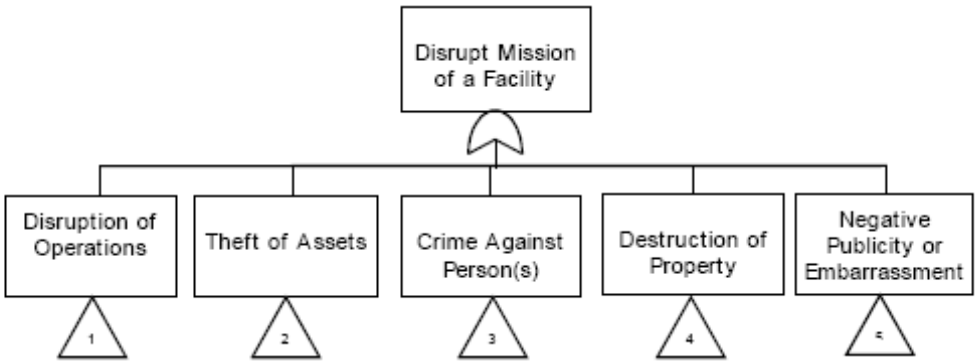


Briefly explain why this risk-based approach to the security of an application might make it more difficult to assess potential safety hazards.

[6 marks]

c) The same US government project that proposed the use of security fault trees has also recently promoted the development of the following risk equation:

$$R = P\_A * (1-P\_E) * C$$

Where R is the risk associated with adversary attack, P_A is the likelihood of the attack, P_E is the likelihood that the security system is effective against the attack, (1 – P_E) is the likelihood that the adversary attack is successful or the likelihood that security system is not effective against the attack and C represents the consequence of the loss from the attack. Using arguments derived from the assessment of safety-critical systems, comment on whether this simple formula is likely to provide important insights into the reliability of complex, security critical software.

[10 marks]

[Cont.]

a) The US Federal Motor Carrier Safety Administration has recently developed guidelines for the design and operation of Lane Departure Warning Systems (LDWS). These monitor the position of a vehicle and warn a driver if the vehicle is about to deviate outside their lane on the road. Currently these LDWS are forward looking, vision-based systems that use software to interpret video images to estimate vehicle state (lateral position, lateral velocity, heading, etc.) and roadway alignment (lane width, road curvature, etc.). The algorithms warn the driver of a lane departure when the vehicle is traveling above a certain speed and the vehicle's indicators are not in use. The LDWS software will also notify the driver when lane markings are inadequate for detection, or if the system malfunctions.

Briefly explain how a truck company might use black box testing to assess the reliability of a number of COTS LDWS applications.

[5 marks]

b) The Federal Motor Carrier Safety Administration describes a generic architecture for LDWS hardware. This includes an electronic control unit that accepts data from a lane boundary sensor through a J1708 or J1939 vehicle network. The control unit monitors the turn signal status and engine power. The output of the system is a status indicator and, when necessary, a warning, which appears on the driver-vehicle interface. Briefly explain how this proposed architecture might aid white box testing of the LDWS software by a system manufacturer.

[5 marks]

c) Celoxica are market leaders in the development of Lane Departure Warning Systems. They have pioneered the use of Electronic System Level (ESL) design. ESL increases the level of abstraction used in circuit design. Rather than building circuits using individual gates or circuit blocks, algorithms are characterized in C-based programming language. These are then implemented in silicon architectures that include software microprocessors, reconfigurable Field Programmable Gate Arrays and heterogeneous System on Chips. Briefly describe the strengths and weaknesses that arise from the use of innovative hardware and software co-design techniques for safety-critical systems.

[10 marks]

4. At present there is little agreement about how to determine whether a software engineer is competent to work on a safety-critical system. Write a brief technical report explaining how you would set up a national scheme so that the public could be confident programmers were competent to work on safety-critical software. Explain the benefits of your proposal by referring to any of the incidents and accidents that we have discussed during the Safety-Critical Systems course.

[20 marks]

**[end]**

**Sample Solutions**

1.

a) On 14th August 2003, the company responsible for ensuring the reliability of electricity supply in part of the northern USA suffered a failure in its Real-Time Contingency Analysis (RTCA) software. This program monitored the current state of the power network and then simulated what would happen if failures occurred. The following extract from the official report into the 'blackout' identifies a number of key factors in the failure:

> "about 12:15, the RTCA produced a solution outside the bounds of acceptable error. This was traced to an outage of the Bloomington-Denois Creek 230-kV line, although it was out of service, its status was not updated. Line status information is transmitted by the data network and is intended to be automatically linked to the RTCA software. This requires coordinated data naming as well as instructions that link the data to the tools. The automatic linkage of this line's status had not yet been established. To troubleshoot this problem the analyst had turned off the automatic trigger that runs the RTCA every five minutes. After fixing the problem he forgot to re-enable it. Thinking the system had been successfully restored, the analyst went to lunch" (US-Canadian Presidential Report, 2003).

Explain the importance of data integrity as a cause of this failure.

[5 marks]

[unseen problem/seen problem]
In the course, we have spent most of the time discussing the processes that are used to develop safety-critical software. The focus of most techniques is on ensuring that software meets its functional and non-functional requirements. There is a preoccupation on the code, the language choice etc (1 mark). However, relatively little attention has been paid top the problems of ensuring data integrity within the design of safety-critical software even though observations such as 'garbage in, garbage out' are cliché's within the subject (1 mark). I have mentioned these issues to them in a lecture and also referred to Neil Storey's (1 mark) work on this issue. We have not however discussed this particular incident, hence there is a key element of problem solving here. For five marks, I would hope that solutions would identify the lack of data from the Bloomington-Denis line as a key weakness in the RTCA software (1 mark). It does not matter how sophisticated the prediction algorithms are, if they do not have real time access to critical data the results will fail to match the observed state of the system over time (2 marks).

b) While the RTCA system was disabled, a second and unrelated software 'failure' affected Emergency Management System (EMS) software. The RTCA program, mentioned above, was used to predict the consequences of a failure in the electricity network. In contrast, the EMS software was used to warn companies when there actually was a failure in the electricity network. The EMS application ran on several servers, any one of which could host all of the EMS functions. However, the normal configuration was to have one server remaining as a "hot-standby". The primary server hosting the EMS failed, due either to "the stalling of the alarm application" or a data overflow on remote terminals that relied on EMS data. As intended, the alarm system application and all other EMS software running on the first server transferred to the back-up server. However, the alarm application "moved intact onto the backup while still stalled and ineffective". The backup server failed 13 minutes later.

Identify the particular strengths and weaknesses of software redundancy that are illustrated by this case study.

[5 marks]

[unseen problem/seen problem]

We have discussed the strengths and weaknesses of redundancy in the course but I have not looked at this particular case study with them. A first class answer would remark on some of the ambiguity in the question (1 mark). It is unclear what is meant by 'stalling of the alarm application'. This ambiguity is deliberate and reflects the publicly available information about this incident. However, there are sufficient details to begin an analysis. The 'hot standby' means that the redundant server was running and ready to take over should the primary server fail (2 marks). This might have been the source of the problem. Such redundancy provides protection against random hardware failures (1 mark). However, if the software is duplicated on both machines then the rapid transition from one server to another without trying to address any potential bugs that caused the primary server to crash will lead to the same failure being repeated (2 marks). In general, software redundancy offers the chance to avoid failures in a primary system but only if some form of diversity is employed, for example using N-version programming (2 marks).

c) The EMS server failures mentioned in part b) slowed the refresh rate on the operators' displays. EMS data was often nested in windows providing more general network information and so the EMS delays affected the operators' ability to refresh more general system information. These problems were compounded because the IT staff did not inform the operators that the EMS servers had failed.

Using the information provided in parts a), b) and c), explain the role that human factors played in both the immediate and longer-term (latent) causes of this safety-critical failure.

[10 marks]

[unseen problem]

There are many different answers here. As far as immediate causes are concerned, solutions could refer to the way in which the operator went off to lunch without checking that the periodic runs for the RTCA software had been enabled. Staff may have lacked critical information to diagnose subsequent problems if the handover from the first operator had not been conducted very carefully (2 marks). In the lectures, we have talked about the critical issues that arise during shift hand-overs (1 mark).

Solutions might also look at the consequences of the failure involving the EMS application. In this case, the failure of the servers removed critical warnings that could further have distanced staff from the state of the underlying application (1 mark). In the lectures, we have spoken about situation awareness and the theoretical material in this area would be relevant (2 marks). First class solutions might look at the degradation of active monitoring that often accompanies the introduction of automated alarms and that any consequent failure of the alarm system becomes very difficult to detect because operators may not notice the absence of alarms when they are not actively monitoring the state of the application itself (2 marks). The final part of the question mentions that IT staff did not tell the operators that the EMS servers had failed (1 mark).

Alternatively, answers might focus on temporal aspects of usability, which are mentioned in the final part of the question (1 mark). Update delays will not only lead to difficulties in planning appropriate control actions but they will also lead to frustration and error (2 marks).

Longer term or 'latent' causes of the failure could focus on many different stages in the development and management of the application. We are told relatively little about these in the extract but for each of the immediate causes mentioned above, it is possible to identify development and monitoring techniques that could have provided some form of barrier against these latent problems. This does not imply that the use of these approached would guarantee the detection and removal of underlying problems. For example, IT staff might have been trained to inform control room staff as soon as they learned of an EMS failure. Such training need not have guaranteed that a warning would be successfully issued and so the automated warnings might have been sent to control room staff as well etc.

2.

a) Briefly explain the different relationships that exist between the terms ‘reliability’, ‘safety’ and ‘security’. (Hint: is a secure system necessarily safe?)

[4 marks]

[seen problem]
In the lectures, the class has been introduced to some of Laprie’s (1 mark) definitions and distinctions. The controversy over this work has also been mentioned. To summarize, in the view of many people, safety and security are attributes to reliability (2 marks). In other words, reliable systems must be both safe and secure. However, others have asked for greater detail in the analysis (1 mark). As the hint suggests, there is a tension between safety and security (1 mark). A totally secure patient information system might deny access to everyone who attempted to log in (1 mark for example). It would be secure because it would never be possible for someone to gain unauthorized access to the patient records. However, such a system could endanger the safety of patients by denying access to medical staff (1 mark for counter argument).

c) Risk based techniques have been advocated for secure applications. For example, the following diagram shows the US government’s Sandia Labs use of fault trees for risk assessment of security threats.

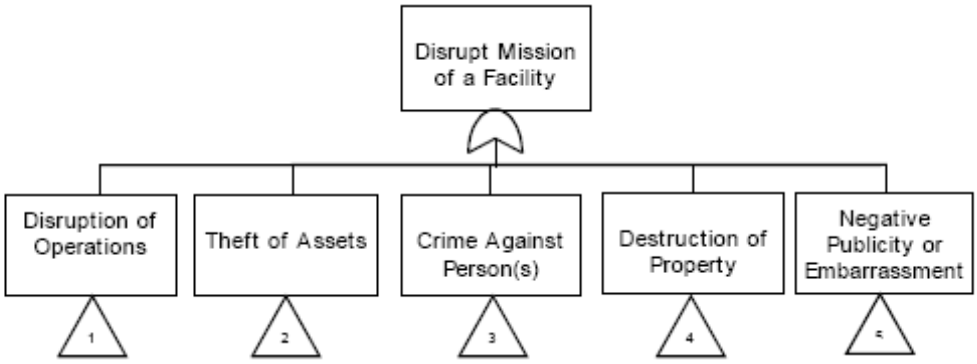


Briefly explain why this risk-based approach to the security of an application might make it more difficult to assess potential safety hazards.

[6 marks]

[unseen problem]
At first sight, this might seem like a complex question but it is really pretty straightforward if they look at the diagram. The key issue is that it becomes impossible to distinguish between those hazards that have safety related consequences and those that are security related (1 mark). For example, ‘theft of assets’ may have no safety related consequences at all, similarly, ‘negative publicity or embarrassment’. In contrast, ‘disruption of operations’ or ‘crime against the person’ may cause physical harm to employees and members of the public (2 marks). In many social and commercial environments there is a concern to treat issues such as ‘negative publicity’ differently from events that might lead to death or injury (1 mark).

First class answers might look more to the pragmatic issues in the application of this form of risk assessment. Although it can be hard to identify the probabilities of basic events in safety-related fault trees, for example, when they involve the likelihood of human error or software failure, it is even more difficult to assess the likelihood of security related incidents (2 marks). For example, many security-related incidents are the result of deliberate violations and targeted attacks (1 mark). Hence some assessment needs to be made of the attractiveness of a target to many diverse groups (1 mark).

c) The same US government project that proposed the use of security fault trees has also recently promoted the development of the following risk equation:

$$R = P\_A * (1-P\_E) * C$$

Where R is the risk associated with adversary attack, P_A is the likelihood of the attack, P_E is the likelihood that the security system is effective against the attack, (1 – P_E) is the likelihood that the adversary attack is successful or the likelihood that security system is not effective against the attack and C represents the consequence of the loss from the attack. Using arguments derived from the assessment of safety-critical systems, comment on whether this simple formula is likely to provide important insights into the reliability of complex, security critical software.

[10 marks]

[seen problem/unseen problem]

We have spent a significant proportion of the course considering whether risk assessment techniques can be used in the development of safety-critical software. We have looked at the structure provided by IEC61508, where risk assessment helps designers match the application of particular safety-critical software development techniques to the hazards that are being mitigated by that software. In other words, greater development resources must be used to the development of software that helps to address high-risk failures (2 marks). We have also looked at the application of various risk assessment techniques directly to software itself, including John Musa's formula for predicting the likelihood of bugs (1 mark) and the use of Leveson and Chin's software fault trees (1 mark, up to 3 for distinguishing different uses of risk assessment in the development of safety-critical systems). In this case, we can see how the previous formula might be 'plugged into' an approach that is similar to 61508. In other words, we can target the analysis of software to those areas that contribute most to the overall security of a system defined in terms of the formula (2 marks). If we try to apply the formula directly to assess the security of software rather than the system as a whole then things become confused, just as they do for software. If we know that there is a potential security problem in a particular piece of code (1-P_E) then we should just remove the bug (1 mark). Just as we should remove a bug if we have concerns about a piece of safety-critical code.

Well-developed answers could also mention the difficulty of estimating the likelihood of attack (P_A), mentioned in the sample solution to the previous part of this question (1 mark), or of the potential consequences, C, just as we have questioned the assessment of outcomes in safety-related incidents (1 mark). It is also possible to question whether the formula would provide a conservative assessment of risk in the face of contingent failures. In other words, attacks are unlikely to be entirely independent and hence we cannot use R to calculate the probability and consequences of subsequent attacks given that a previous attempt has been made to compromise a system just as we cannot easily use the standard risk formulae to account for situations in which (for example) a poor programmer introduces numerous bugs into the same piece of code (3 marks).

3.

a) The US Federal Motor Carrier Safety Administration has recently developed guidelines for the design and operation of Lane Departure Warning Systems (LDWS). These monitor the position of a vehicle and warn a driver if the vehicle is about to deviate outside their lane on the road. Currently these LDWS are forward looking, vision-based systems that use software to interpret video images to estimate vehicle state (lateral position, lateral velocity, heading, etc.) and roadway alignment (lane width, road curvature, etc.). The algorithms warn the driver of a lane departure when the vehicle is traveling above a certain speed and the vehicle's indicators are not in use. The LDWS software will also notify the driver when lane markings are inadequate for detection, or if the system malfunctions.

Briefly explain how a truck company might use black box testing to assess the reliability of a number of Commercial Off The Shelf (COTS) LDWS applications.

[5 marks]

[seen problem / unseen problem]
We have looked at the issues of back box testing in the course but we have not discussed the case study before. The commercial development of LDWS is still relatively new. Black box testing assumes that the testing team does not have access to the internal, implementation details of the software (1 mark). In this case, they might use a prototype version of the LDWS and install it in a test vehicle (1 mark). This could then be driven in a controlled environment to determine whether it met the high-level requirements mentioned in the question (1 mark). These would include various lateral positions, velocities etc. It would also include different environmental conditions to test the vision system, eg rain and fog (2 marks). Subsequent tests might focus on the COTS issues by extending the study to different motor vehicles (1 mark). However, if a test was not passed then it would be difficult for the group doing the tests to diagnose the reasons for the failure given that they do not necessarily have access to the internal structure of the applications (2 marks).

b) The Federal Motor Carrier Safety Administration describes a generic architecture for LDWS hardware. This includes an electronic control unit that accepts data from a lane boundary sensor through a J1708 or J1939 vehicle network. The control unit monitors the turn signal status and engine power. The output of the system is a status indicator and, when necessary, a warning, which appears on the driver-vehicle interface. Briefly explain how this proposed architecture might aid white box testing of the LDWS software by a system manufacturer.

[5 marks]

[seen problem/ unseen problem]
The FMCSA architecture breaks down the LDWS into a number of recognized components. This has considerable advantages for white box testing (1 mark). In theory, it would be possible to identify interfaces between these different components and then develop test cases based on assumptions about the performance of each unit (1 mark). This would support incremental test strategies during development rather than the summative approach of black box testing mentioned in the previous answer (1 mark). For example, the testing of the control unit might be done separately from the J1708 or J1939 architecture assuming that the network data link to the control unit could be simulated, for example to mirror the information from turn signals and engine power (1 mark). Of course, this white box approach would make it far easier to diagnose the cause of particular problems (1 mark). Equally, it is not sufficient for most safety critical systems where the interactions between components can lead to hazards that are seldom apparent in unit testing. For example, there may be interactions between the LDWS and other control systems sharing the vehicle network under pathological conditions that should be anticipated even though they are rare if they have adverse consequences (2 marks).

c) Celoxica are market leaders in the development of Lane Departure Warning Systems. They have pioneered the use of Electronic System Level (ESL) design. ESL increases the level of abstraction used in circuit design. Rather than building circuits using individual gates or circuit blocks, algorithms are characterized in C-based programming language. These are then

implemented in silicon architectures that include software microprocessors, reconfigurable Field Programmable Gate Arrays and heterogeneous System on Chips. Briefly describe the strengths and weaknesses that arise from the use of innovative hardware and software co-design techniques for safety-critical systems.

[10 marks]

[unseen problem]
This is a more difficult question but it creates opportunities for the more able students to show their ability to think analytically about the safety implications of techniques that they may have met in other areas of the course or in their general reading about computing related topics. Hardware and software co-design create considerable analytical challenges in safety-related systems. On the one hand they create the opportunity to develop safety-related systems at a higher level of abstraction than would otherwise be the case (1 mark). Arguably, it will be cheaper and less error prone to use a higher level programming language than lower level circuit design (2 marks). I have not mentioned VHDL because Celoxica integrate VHDL into their approach but if a student advocated this as an alternate to the co-design mentioned explicitly in the question then that would be great (2 marks).

Co-design approach creates a host of further problems. The use of FPGAs and reconfigurable hardware implies that potential problems may depend not simply on the code but on the way in which the code is realized in hardware (1 mark). Hence it may involve the physical as well as the logical analysis of the intervening layers of abstraction (1 mark). It is unclear how to deal with the exponential branching of possible behaviors given the increased blurring of hardware and software (1 mark). Formal methods are widely regarded to provide some solutions here but there is little agreement on how they can be integrated with the co-design of safety-critical applications (2 marks). At present these architecture are not widely used in safety-critical systems. Hence they lack the evidence provided by larger user populations that have supported the safety-critical application of specialist safety microprocessors (eg AAMP5 etc) (2 marks).

These are initial thoughts and alternate answers are possible for maximum marks.

4. At present there is little agreement about how to determine whether a software engineer is competent to work on a safety-critical system. Write a brief technical report explaining how you would set up a national scheme so that the public could be confident programmers were competent to work on safety-critical software. Explain the benefits of your proposal by referring to any of the incidents and accidents that we have discussed during the Safety-Critical Systems course.

[20 marks]

[Essay]
This is an open-ended question, the class have been told to expect an essay and I will provide them with some broad hints to consider the issues of competency before the exam. The UK Health and Safety Executive have recently published some guidelines on this topic.

The class have been introduced to the ethical guidelines published by both the ACM and the BCS as part of the course (1 mark). I would expect everyone to mention this at some point. These guidelines make it clear that programmers must report any potential bugs or problems that they identify in both their own code and that of their colleagues (1 mark). The guidelines also make it clear that programmers have a duty of care to the public that over-rides all other considerations in the development of safety-related systems (2 marks). However, I have also taken them through the events leading to the Bristol Infirmary Inquiry and to the difficult position that 'whistle blowers' often find themselves in (2 marks).

I'm hoping that this question will offer some of the more able students an opportunity to reflect about their experience on the course (1 mark). Hopefully, they will recognise that there is a core of generic material that makes safety-critical programming different from other areas of software engineering. For example, we have explored the dangers of code reuse and the inclusion of redundant code (e.g., Ariane V) (2 marks). We have also looked at the issues of resource management and risk assessment (e.g., London Ambulance

Servce Computer Aided Dispatch System) (2 marks). We have looked at the impact of the choice or programming language and of programming techniques, such as caching and garbage collection, that can threaten safety if not handled carefully (1 mark). In other words, there are key ideas that must be understood before any general software engineer is 'competent' to work on safety-critical systems (2 marks).

Another issue to consider is the difference between programming expertise and domain expertise (1 mark). It is unclear whether competent programmers must also be domain experts (1 mark). Is it possible for someone who specialised in the software engineering of safety-critical applications to move between different domains on a contract by contract basis or do programmers have to specialise in the development of safety-critical avionics systems? I would expect a balanced approach, referring to the development of specialist teams within companies and the ability of small numbers of programmers to move between teams on a periodic basis providing there is core continuity of specialist skills (2 marks). Even in such an approach, it is impossible for programmers to become experts in all areas of domain expertise and so a key competency must be the ability to communicate with these experts. Answers might refer to the Lockhead programmers involved in the Mars Climate Orbiter Star Camera who resolved a data format error but did not check that the data referred to valid coordinates with the navigation team (2 marks).

The mention of team work raises a series of issues to do with more personnel qualities that make for competent software engineers. Most of these qualities would be useful in any application domain, such as the ability to think analytically, the ability to communicate complicated technical concepts to non-technical audiences, a form of paranoia when interpreting test results and so on (2 marks). However, first class answers should point to the more serious consequences should software engineers fail to exhibit these more general qualities. I have given numerous examples of communication failures between medical physics teams, device manufacturers and clinicians that could be used to illustrate this argument (2 marks).