SCS Exam H 2015-16, Answer 3 Questions

1.

a) Briefly describe the benefits of a cold-redundant back-up system for a safety-critical application.

[4 marks]

[Seen problem]
A cold redundant system is one that does not run alongside the primary application. This has certain disadvantages – in particular it will take time to restore the state of a failed system so that the secondary or redundant application can be started. However, there are also a number of benefits. In particular, cold redundant systems are usually cost effective because the hardware will be used less often than the primary system – this can be important when, for example, a backup redundant processor is then used as the primary system to extend the life of an application. Cold redundant systems are often simpler than warm or hot counterparts – there are less issues associated with locking and contention. There are no issues associated with disagreement between a warm active redundant system and the primary etc.

b) What are the main problems that must be considered when applying cold-redundancy to software rather than hardware applications?

[6 marks]

[Seen/unseen problem]

Hardware obeys the reliability distributions characterized by the bathtub curve – with both burn-in and burn-out time. If both the primary and the back-up are 'burned-in' and a cold redundant system is available then it is very unlikely that the back-up will fail on demand when a primary fails. However, software performance cannot be characterized in the same way. It does not age in the same way that hardware does. Hence a cold redundant backup will fail in exactly the same way as the primary system if they both contain the same bugs or logical errors. The main way to avoid this would be through the use of software diversity and N-version programming – where the cold redundant backup would be written by a different team than the primary.

c) Hot, Cold and Warm redundancy are also used in non-safety applications. For instance, Cisco uses these techniques to ensure that packets can still be forwarded even when hardware components fail. Explain in detail the different amounts of state information that must be held by primary and backup systems when using hot, cold and warm redundancy. (Hint: comment on their relative space and time efficiency).

[10 marks]
[Unseen problem]

In both safety-critical and conventional systems it is important to increase availability through the use of redundancy. Hot redundant systems operate in parallel with a primary application. In most cases, state information is mirrored or duplicated. Arguably this is space inefficient (as the state is duplicated) until the point at which the primary system fails. The results from the backup are not used even tough they are being computed in parallel with the main system
The benefit is that the hand-over is extremely fast – however, this depends upon a swift transition once the secondary has detected that it must take over. In warm redundant

systems, the backup will mirror the state of the primary application but they will not work in step. An example would be the use of checkpoints where a snapshot of the state could be recorded before critical operations. If the primary failed then there would be a slight delay while the secondary was updated using transaction logs from the latest checkpoint. This is a mid-point in terms of both space and time efficiency. The exact balance is determined by the period between checkpoints. Cold redundant systems would have to recreate the state of the failed application, for instance from secondary memory or from backup logs. It would be space efficient but time inefficient as the cold system would have to be brought up.

2.

a) Briefly identify the main components within a safety-management system and explain how they contribute to the safety of complex software.

[6 marks]

The main components within a safety management system are:

a. Risk assessment – this is critical to identify the hazards in the operation of a safety-critical system. These hazards from the focus of subsequent design in order to reduce the likelihood or mitigate the consequences of known hazards;
b. Design and operation – during the course we have covered a range of safety related design and implementation techniques – for instance the use of redundancy described in question 1 – these techniques are applied to mitigate risk. However, it is hard to be sure we have identified all hazards and to ensure that those we have identified have been mitigated so we need to monitor…
c. Incident reporting – is a primary means of monitoring safety-critical systems – accidents are rare but near misses provide valuable information that can be used to refine a risk assessment in an iterative approach to safety management.

[seen/unseen problem]

b) What technique would you use to identify and quantify the risks associated with **human error** within an interactive, safety-critical system? Briefly justify your decision and explain why you did not choose at least one other approach.

[7 marks]

[seen/unseen problem]

There are many answers to this question – with good scope for the more able students. I have covered techniques such as CREAM and Human Reliability Analysis in the course. I have also mentioned specific extensions to HAZOPS and the consideration of human factors in FMECA. The lecture notes and slides include the use of Fault Trees to consider human and system failures. Key concerns would be whether to focus on qualitative or quantitative techniques – and the subsequent concern to validate numerical estimates as well as to link them to underlying psychological theories. It would be important to consider the integration of the approach with more conventional forms of safety analysis. An alternative approach would be to focus on the work of Reason or Rasmussen or of resilience engineering –

where human intervention is perceived to resolve more risks than are created by errors. In this case, we have to ask questions about how to integrate the technique with other forms of more conventional systems engineering – including hardware reliability analysis.

c) What technique would you use to identify and quantify the risks associated with **software failure** within an interactive, safety-critical system? Briefly justify your decision and explain why you did not choose at least one other approach.

[7 marks]

[seen/unseen problem]
As with the previous question, there are many answers. We have looked at standards such as IEC 61508 and ED-153 that focus on the identification of SILs or SWALs and software is assumed to mitigate other hazards in proportion to these risk surrogates. The main emphasis of these approaches is to avoid associating failure probabilities with software – however, we have also covered the work by AT&T with the John Musa formula. Alternative answers focus on the Leveson and Chin Software Fault Trees or on MCDC testing or on formal methods, which are all relevant as approaches that avoid the quantification of software failure probabilities. Higher marks will be awarded to any answers that question the utility of quantification here and also the difficulty of validating those assessments given the problems of test coverage through complex safety-related code.

3.

a) Briefly explain why we tend to use process rather than product based approaches to the validation and verification of safety-critical software.
[4 marks]

[seen problem]
This has a connection to question 2 – with complex, safety-critical software there are ethical and practical concerns associated with the validation of the product –t he code itself. We cannot test all possible execution paths – in a realistic environment especially when there are risks associated with the processes being controlled. This also creates barriers to regulators who cannot be expected to understand and assess every aspect of many millions of lines of code. IPR issues may also restrict access to source code for external validation/verification. In contrast, the focus is on assuring that companies follow approved development processes. These approaches do not rely on access to source code, you can audit development documentation and look at the techniques used during different stages of the development process.

b) Explain the problems that can arise when attempting to validate and verify any maintenance tasks that have been conducted on legacy software components.

[6 marks]

[unseen problem]

With legacy code you do not always know the verification tests that were conducted nor can you always recreate the validation criteria that were used by previous generations of software engineers. Even if the documentation is very good, in millions of lines of code a huge investment is required to reach the level of understanding achieved by the original developers. The changes made to code might invalidate these previous validation/verification criteria – a related issue in DO-178C is that of traceability – in other words can you easily map from validation and verification criteria down to particular lines of code to know whether or not you have to consider those criteria from particular changes. As with previous questions this is relatively open and a number of alternate approaches could be used here.

c) A Nordtest review of PLC validation techniques argued that; "The complexity of such a validation task is further increased by the fact that a proper validation of the system must address not only the application software, but that as much as 5 different aspects must be considered:

- The PLC context, i.e. all external I/O interfaces, sensors, actuators and power supply.
- The PLC input/output hardware providing the interface between the context and the PLC processor.
- The PLC processing HW (e.g. CPU, memory, timers, interrupt, watchdog)
- The PLC operating system (OS) often called the kernel that provides the environment for the application soft- ware, as well as a certain amount of fault handling.
- The PLC application software, implementing the user specified functionality for:
  - Primary safety functionality
  - Derived functionality to handle context fault conditions".

Explain why each of these different aspects must be considered within the validation of a PLC within a safety-critical application.

[10 marks]

[unseen problem].

The first concern focuses in "The PLC context, i.e. all external I/O interfaces, sensors, actuators and power supply". This is important because failures in any of these external interfaces must be dealt with by the PLC or by external assurance techniques – no matter how good the code on the PLC if the environment fails then the PLC may not function correctly.

The second concern focused on "The PLC input/output hardware providing the interface between the context and the PLC processor." In high reliability applications, redundancy may be used on the buses and other interface components because these represent a single point of failure. Even if we have redundant sensors and redundant PLCs then the data comms between these devices will bring the system down if they are not also duplicated eg consider Byzantine failures – however, redundancy at this level increase, cost, complexity, power etc.

The third area addresses "The PLC processing HW (e.g. CPU, memory, timers, interrupt, watchdog)" this is relatively self-evident. Aspects of the previous questions could be brought in here. Some students might oick up on Program Counter monitoring and other techniques for the implementation of watchdogs.

The fourth area focuses on "The PLC operating system (OS) often called the kernel that provides the environment for the application soft- ware, as well as a certain amount of fault handling". This raises a lot of concern – typically the PLC kernel is not open source so at some level we have to either conduct black box testing or trust the suppliers. We might use an approved vendor approach and also insist on appropriate safety management techniques mentioned in previous questions.

The final area considers "The PLC application software, implementing the user specified functionality for: Primary safety functionality and Derived functionality to handle context fault conditions". In some respects this is the easiest area to consider and good solutions wil draw on the previous sections of this question – you can in most cases use White box approaches but within the constraints imposed by legacy systems and by the concerns over traceability raised in part b).

4. A number of researchers have identified the problems associated with responding to 'weak signals' in the design and operation of safety-critical systems. Briefly enumerate some of these problems and identify techniques that might be used to address these concerns.

[20 marks]

[Essay – unseen/seen problem]
The number of accidents in most safety-critical systems across the Western world and most other areas of the world is relatively low. In consequence, it can be hard to identify what the potential risks might be. If we focus on previous rare accidents then we have very little data to go on. In consequence, we might choose to look at near miss incidents – there will be many more of these. However, there is often a big difference between an incident and an accident. Multiple levels of defence might intervene to prevent someone from being killed so it can be hard to argue about how near a miss might actually have been. Also, the large number of safety concerns means that the signals that incidents provide about future accidents will be relatively weak. Often accident reports talk about 'a failure of imagination' – to go from an incident report to imagine the way that multiple fatalities could occur. This is important because in the course we will have discussed Linate and Ueberlingen where similar incidents were logged in the days before these major accidents. However, in many cases there are so many other incident reports that did not lead to loss of i.e. that the signal provided is very, very weak. We cannot respond to every safety concern with limited finances and the need to maintain levels of service so the question is often about the prioritization of resources to these weak signals.

Another approach would be to look at the ideas of resilience engineering – where instead of focusing on a small number of very rare accidents we look to reinforce the behaviors on a 'good day' that made sure things went right. There are a set of issues here – for instance, identifying those aspects of a safety-critical system that contributed to the absence of an accident or indeed identifying how close we were to a failure even though there was no loss of life.

I would expect some of the more able students to talk about the use of counter factuals as a key technique for analyzing incidents and accidents. Also, the use of data mining approaches to consider patterns even when the signals are relatively weak.