ENGR 401 handout Week 4 -Scan into ecs…

Student ID:

1. Therac-25 Computer Controlled Radiation Therapy Incident (Risk Management Review):

Note – This is the same material provided in the week 3 workshop handout. You will now be considering and assessing it from a risk perspective rather than ethics, but keep thinking about how the two fields are related and have overlaps.

Sources:

- https://en.wikipedia.org/wiki/Therac-25
- https://hackaday.com/2015/10/26/killed-by-a-machine-the-therac-25/
- https://www.youtube.com/watch?v=41Gv-zzICIQ

Summary

What was the Therac-25

The Therac-25 was the latest in a generation of radiation therapy machine, ostensibly a "cancer zapper". Machines of this class use beams of x-rays or electrons to target and kill specific areas of tumour cell, potentially deep deep inside the body. While there is always going to be a certain amount of collateral cell damage, but like chemotherapy the hope is that more cancerous material will be killed rather than healthy.

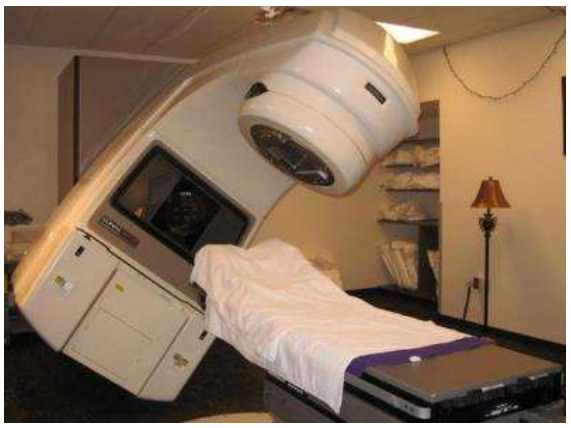

The device was originally made up of an electron beam which could run in a low-power or high-power mode, and a turntable that positioned different targets for the beam to strike before it reached the patient depending on the type of treatment needed.

- X-Ray Treatment Mode - The beam was in **high-power mode**, and the turntable would be set to cause it to hit a tungsten target that both converts the beam to X-Rays, and disperses them over the treatment area.

- Direct Electron Treatment Mode – The beam was in **low-power mode**, and the turntable would be set to cause it to be dispersed over the treatment area using magnets.

These electron and x-ray beams therefore need to be highly regulated and controlled, an inappropriately aimed beam or beam with an incorrect level of power could be highly damaging if not fatal.

The original editions of the Therac included physical safety interlocks to prevent patients being exposed to unsafe radiation such as from a direct hit of the high-powered electron beam used for X-Rays.

One example is that if the high-powered electron beam was selected to be fired at a patient, without the X-Ray target in place between the patient and the beam, the electric circuit that was created by that arrangement would result in a fuse blowing and disconnecting power from the Therac.

For the Therac-25, these physical safety features were removed from the hardware, and instead it was left up to the newly attached PDP-11 computer to control the configuration of the beam and turntable and monitor for any unsafe configurations. The computer was faster to run the motors on the device and set it up for the procedure, something that hospital staff and administrators loved for simplicity and speed and perceived accuracy. Programming code was a new thing, and once in place it was assumed not to fail.

What unfortunately happened

For six patients between 1986 and 1987, something went wrong with this configuration setup. The Therac-25 exposed them to massive overdoses of radiation, killing four patients and leaving two with lifelong injuries.

When things went wrong, the patients under treatment were reporting feeling tremendous amounts of heat and burning. In some cases, the machine would stop with an error “Malfunction 54”, which the operators only knew as either too much or too little energy had been released. The error could be cleared, and then the beam restarted.

The supervising hospital physicist would report to the vendor AECL and their local medical regulator that an overdose happened. Initially AECL denied that the Therac-25 was capable of delivering an overdose due to the amount of software protections in place that would throw errors for any problem. If anything the machine was so safe it would deliver less than the required radiation not more.

However, there was that much confidence in the correct operation of the computer-controlled system, that initially it was seen as impossible for this to have happened.

What turned out to be happening

After the second of the incidents that occurred at the East Texas Cancer Center in Tyler Texas, the staff physicist Fritz Hager was determined to get to the bottom of the issue. He and a radiotherapy technician worked through the night and weekend to try and reproduce the specific error “Malfunction 54” that was not mentioned in the manuals.

What they eventually found was that if a user would move the cursor using the arrow keys, select “X-Ray Mode”, and the turntable would begin turning to align the X-Ray target as well as set the electron beam to high-power. This would take approximately 8 seconds.

If during these 8 seconds the user used the arrow keys to switch the machine to electron beam mode, the turntable would not switch to the correct position, instead being left in an unknown state with the electron beam set to a dangerous level.

This was due to a race condition in the software, where the code was essentially assuming that no-one would try to make changes to the configuration while the turn table was still rotating.

An operator in another facility reproduced this behavior on their Therac-20, which you will remember had a safety interlock fuse that was removed on the Therac-25. In that facility the safety fuse blew, that would have prevented the electron beam from energizing.

During the investigation of the incidents, there were two related causal issues. First that all physical safety interlocks that had prevented the previous generations of the Therac from being incorrectly setup for a patient were removed from the Therac-25, with control given over to the PDP-11 computer attached to the device. Then the software that the computer runs to control the setup of the device's radiation exposure contained undetected bugs.

What it "seems" the later investigation found

While the vendor AECL never officially released the source code, reports of investigations showed that the software that controlled the system and provided the only safety functions seemed to be written by a programme with little experience in real-time systems. There were few comments, and no proof that timing analysis and been performed.

There was allegedly no testing of the Therac-25 hardware and software together before the unit was assembled at a hospital, with the "testing" hours counted as only the time when a hospital staff operator was using the machine on a patient.

Of more important note, when AECL had been considering the incidents reported to them from the first patient onwards, the design of the software was not considered – instead focusing purely on the hardware and assuming the software was free of bugs.

New Risk Management Workshop Material Starts Here

What you have been asked to do

You have the advantage now of seeing how this scenario actually played out after the product was launched and unfortunately used on patients.

Despite that, place yourself in the imagined position that you were asked by executive management to carry out a “Pre-Launch Risk Assessment” of the new Therac-25 treatment unit.

It is ready for delivery to hospitals sitting in boxes in a warehouse – the only final step is your review.

Your risk assessment will include a recommendation whether to proceed with delivery or not, and the completed document will be given to executive management for review and to make the final decision.

You have complete access to the device and all documentation regarding its construction, decisions made along the way, and how the product development project was undertaken (essentially everything provided in the summary above, EXCEPT that it has never been used on a patient yet).

Management has given you complete freedom to look at the risks from any perspective or category that you wish, but we suggest you may wish to consider at least some of the following areas;

- Physical patient safety systems
- Operator training and education
- Software development and testing process
- After-sales product support and incident response
- Design and manufacturing processes
- Product design and development staff health & safety

Use the following 'Risk Matrix' as necessary in your answers to the following questions

Impact →

Likelihood	Negligible	Minor	Moderate	Significant	Severe
Very Likely	Low Med	Medium	Med Hi	High	High
Likely	Low	Low Med	Medium	Med Hi	High
Possible	Low	Low Med	Medium	Med Hi	Med Hi
Unlikely	Low	Low Med	Low Med	Medium	Med Hi
Very Unlikely	Low	Low	Low Med	Medium	Medium

https://2.bp.blogspot.com/-McErI5J1EK0/Vd3bB1Jg5GI/AAAAAAAAHwg/w-7KoUDzmyc/s1600/Typical%2BRisk%2BMatrix.png

An example risk definition table that may be useful is as follows

Category	Risk	Impact	Likelihood	Treatment	Residual Impact	Residual Likelihood
Patient Safety	Short circuit due to manfacturing causes non-fatal patient injury	Significant	Very unlikely	N/A	N/A	N/A

Questions/Tasks

1. Identify risks by category and describe by impact and likelihood (using the risk matrix provided).

2. Then pick an overall risk rating that you feel is appropriate for treatment, and for risks that have that rating or higher, briefly describe treatment plans and then establish the residual risks.

Justification & Recommendation Additional Instructions;

- For the following two activities, limit the total combined size of your answer to approximately half a page of A4.

- Think about how you communicate on the presumption that the executive management may only ever read your justification and recommendation. In real-life this may be a 'standalone' paper where management may call you to present, but that can not be assumed.

3. Risk Treatment Justification: Give a short justification for why you picked that overall risk rating as the cut-off point for applying treatment plans.

4. Final Recommendation: Document a recommendation whether given the findings of your risk assessment, **and assuming all your treatment plans were accepted and implemented** by executive management, the launch of the (potentially) revised Therac-25 product should go ahead. Include rationale for your decision.