



Online Ethics Center
FOR ENGINEERING AND SCIENCE

Exercises for Therac-25

Author(s)

Charles Huff
William J. Frey
José A. Cruz-Cruz

Year

2003

Description

A collection of student exercises to be used when teaching this case.

Body

Some initial considerations in teaching this case

The Therac-25 case is complex and multi-layered enough to require more than a simple once over to understand. There are multiple actors, some of them representing the same entity at different times. There are closely interwoven networks of action and reaction guided by multiple and mixed motives, where the real state of the information available to an actor at any one time is unclear.

This is not, however, simply the uniqueness of the Therac-25 case, it is a property of all cases if they are studied closely enough. Finally, it is a property of the real life of technology in use. We provide here some exercises to help students grapple with the complexity of these situations.

But first a comment on simple answers. We recommend you read the section on

pitfalls before teaching this case. It outlines ways to approach this case that bring only a shallow level of understanding to the complexities. In the Therac-25 case, one of these pitfalls (single causation) leads to the tendency to fix each error one discovers with a local "patch." This usually increases the complexity of the system, provides false confidence in its safety, and does not address the design issues that led to the existence of the error in the first place. This is clearly the kind of thinking that AECL indulged in during its initial reactions to the early accidents. We recommend you help your students avoid it as they approach this case.

Bill Frey and Jose Cruz provide many other exercises associated with cases from computingcases.org.

Exercises Table of Contents

1. [Analyzing Therac-25](#)
2. [Computer Control Choices Exercise](#)
3. [Tracing the Coding Errors to the Hazards](#)
4. [Software Safety Myths](#)
5. [Designing a Reporting System](#)
6. [Role Playing the Case](#)

Analyzing Therac-25

This exercise uses a modified version of Robert Collins and Keith Miller's Paramedic Ethics procedure. Collins and Miller recommend a procedure to use in evaluating a decision. We are not here evaluating any particular decision, but we can use their method to help us understand the obligations, rights, costs, and benefits for each of the parties in the system. This exercise will require students to read the case on the website with some knowledge of the method they will be using, so they can take relevant notes as they read. Thus, the best approach to this exercise requires introducing the modified paramedic ethics procedure in one class, assigning the exercise and case as homework, and then spending the next class period discussing student's conclusions.

Alternatively, students might be assigned the case to read for homework and then introduced to the method of analysis in the subsequent class. If this approach is taken, be sure to have the case available in class (either on a computer with a projector or in printouts for each student) to aid recall.

There are several approaches to having students read the case for this exercise. You might have them read all the case section but exclude the accident reports. Once students have gone through the paramedic procedure based on their knowledge, you might then introduce them to one or more of the accidents. Does this new information change their assessments of the case? You might give some students partial information (e.g. just the background sections) and others more extensive information. This too is likely to produce differences in their analyses of the case. Alternatively, you might use small sections of the case (e.g. just the background) early in a course and add information about the case as the course progresses.

Each of these approaches are likely to produce differences in the way the case is analyzed by students. These differences help make it clear how important a comprehensive view of a case is.

Our modified paramedic ethic procedure consists of 4 phases. The basic analysis consists of phases 1 and 2, in which the basic relationships among the important stakeholders in the case are outlined. The phases that construct and judge the various alternative scenarios can be done as many times as you wish for each set of actions you think are important. To make this go faster, you might assign groups to construct and present their analysis of the duties and rights of each of the main

stakeholders presented in the case: AECL, FDA, hospitals, operators, and patients.

Gather data

1. **List the relevant stakeholders.** Start with some of the groups mentioned in the socio-technical system page. However, do not end there. Notice that our accident victims, the patients, are not included. Other important groups may also be omitted (e.g. "the public"). The ImpactCS framework provides you with a useful guide to different levels of stakeholders that you might overlook.
2. **Outline the duties and rights the stakeholders have toward each other.** This is best done with a drawing of each stakeholder with arrows indicating duties one owes to other and rights one has. Duties always have targets, one has duties to a particular person (even to oneself). Rights may appear to be free floating (e.g. not to be harmed) but they can often be translated into duties that others have toward the individual (avoid harming X). The ImpactCS framework provides a useful guide to outlining these duties and rights. Use the list of ethical issues to remind yourself of rights and duties in the range of likely ethical domains.

Analyze the data

1. **List the relevant opportunities and vulnerabilities that each stakeholder had in the case.** This is the beginning of what Collins and Miller call a utilitarian ethical analysis. Who is being helped and harmed? What advantages or opportunities does each party receive in this case? What costs or dangers, or vulnerabilities does each party experience?
2. **Determine to what degree each stakeholder's duties were fulfilled or neglected.**
3. **Determine to what degree each stakeholder's rights were violated or protected, and by whom.**

Construct an Alternative Scenario.

1. Construct a promising alternative for some set of actions for a significant actor (e.g. reporting procedures in AECL, FDA procedures, hospital treatment procedures, safety analysis procedures by AECL). For some hints about alternative sets of actions, see the exercises about computer control choices and about reporting procedures.

Judge the Alternative

1. Judge the alternative's effect on each stakeholders' opportunities and vulnerabilities and on each stakeholders' duties and rights.
2. **Imagine each stakeholder in a negotiation** with other stakeholders about whether the alternative should be adopted or not. This certainly helps uncover disagreements about the opportunities and vulnerabilities for each party. One interesting way to stage this negotiation is to have parties that initially represent each stakeholder attempt to don a "veil of ignorance" about which stakeholder they might be when the alternative is adopted. If you might be randomly assigned to any of the stakeholder roles in the case, how would this affect your evaluation of the alternative?
3. **Rank the alternative with other alternatives** for that set of actions. An alternative does not have to be perfect, or even optimal, to be better than the others.

Reference:

Collins, W. R. & Miller, K. W. (1992). Paramedic ethics for computer professionals. *Journal of Systems and Software*, 1-20.

Computer Control Choices Exercise

EXERCISE: Use the range of human-computer control possibilities (on p. 448 in Leveson) to locate Therac-25 control levels. Recommend and argue for a change in level. What would be required to move a level up? Down?

Choosing the Level of Computer Control

In her book *Safeware: System Safety and Computers*, Nancy Leveson lists nine different levels of computer control (taken from Sheridan's analysis):

1. The operator does everything.
2. The computer tells the operator the options available.
3. The computer tells the operator the options available and suggests one.
4. The computer suggests an action and implements it if asked.
5. The computer suggests an action, informs the operator, and implements the action if not stopped in time.
6. The computer selects and implements an action if not stopped in time and

then informs the operator.

7. The computer selects and implements an action and tells the operator if asked.
8. The computer selects and implements an action and tells the operator if the designer decides the operator should be notified.
9. The computer selects and implements an action without any human involvement.

After students have explored the case, have them decide at what level the Therac-25 system is targeted. This may initially cause some confusion, since one way of looking at the system is to think that the operator tells the computer what to do and then the computer does it. Point out to them that this is true in the larger sense, but that the computer clearly has sensors and information available to it to allow it to give error messages. What do we know about the level in this control hierarchy at which those error messages are resolved?

What levels of computer control is the system using when:

- an error message is given (e.g. Malfunction 54), but the system allows the operator to press a "proceed" key to retry the treatment.
- vs. (as required by the FDA) the treatment is suspended after any error and all treatment data must be typed in over again
- or, when the operator is required to "visually check the settings" on the treatment machine
- vs. when the machine sets itself up based on the treatment data entered and then proceeds with the treatment

Once you have established levels of computer control the machine is using, ask for suggestions about how one might increase the amount of computer control. What safety issue does this bring up?

One of the best ways to analyze the effects of changes in computer control is to have already completed the basic steps in the case analysis (determining stakeholders, duties and rights, opportunities and vulnerabilities).

References:

Leveson, N. G. (1995). Safeware: System safety and computers. New York: Addison Wesley.

Sheridan, T.B. (1989). Trustworthiness of command and control systems. In J. Ranta,

(ed.) Analysis, Design, and Evaluation of Man-Machine Systems, (p. 427-431). New York: Pergamon Press.

Tracing the Coding Errors to the Hazards

The Leveson excerpts section of the resources reprints explanations from Nancy Leveson about each of the two identified coding errors in the system that resulted in overdoses to patients. Have students trace each coding error from the problematic variable or operation (e.g. a comparison) to how this resulted in an overdose.

1. What items or sections in the code you have reviewed should be labeled safety-critical? Why? How is it different from other sections of code?
2. What information is available in the design that the code is safety-critical? Assume you are inspecting the code before it is shipped and do not use information gleaned from accident reports.
3. Are the temporary fixes recommended by AECL adequate to remove the hazard?
4. What design changes would you recommend to the software, to the machine, or to the socio-technical system that might reduce the hazard?

This exercise might be done as an in-class exercise or as individual homework and then discussed in the class.

Software Safety Myths

In her book Safeware: System Safety and Computers (p. 26) Nancy Leveson lists seven myths regarding the safety of software.

1. The cost of computers is lower than that of analog or electromechanical devices.
2. Software is easy to change.
3. Computers provide greater reliability than the devices they replace.
4. Increasing software reliability will increase safety.
5. Testing software and formal verification of software can remove all the errors.
6. Reusing software increases safety.
7. Computer reduce risk over mechanical systems.

After having the class explore the Therac 25 case, ask students to evaluate the

truth of each of these statements as they pertain to the case. This can be done either as part of a homework assignment, with class discussion after papers are turned in, or as a class discussion followed by individual papers. Alternatively, you might combine these two approaches and have students turn in a paper and then revise it (or write a short postscript) based on class discussion.

Reference

Leveson, N. G. (1995). *Safeware: System safety and computers*. New York: Addison Wesley.

Designing a Reporting System

A life cycle approach to software requires some way to gather reports in the field of the operation of the software and feed those reports back into maintenance and updating of the software. One of the clear difficulties in the Therac-25 case was the process of getting the right information back from the field to the AECL home office and to other sites and then getting resolutions of the problems communicated back to the sites. In some cases AECL was only notified by lawsuit months after an incident. In other cases, information languished at the home office that might have been useful to sites where the machine was being used.

In this exercise, you will ask your class to design a reporting system and to evaluate its impact on the various stakeholders in the case. In her book *Safeware: System Safety and Computers* (p. 88), Nancy Leveson lists four requirements of a successful reporting system:

1. Explicit delegation of responsibility for reporting. Who should report accidents and to whom? What about other errors or malfunctions? What kind of deadlines and penalties should be imposed? Whose responsibility should it be for imposing deadlines and penalties (e.g. the company, the FDA)?
2. Protection and incentives for informants. If hospitals or manufacturers are required to report errors, incidents, or accidents, there is likely to be some resistance to reporting all errors because of liability issues. What sort of protection and incentives might be given to increase accuracy? Who else within the system other than an official representative might be a useful informant?
3. Procedures for analyzing incidents and identifying causal factors. When an accident or error is reported, who should investigate the facts? How should the

person or panel identify causal factors?

4. Procedures for using reports and generating corrective actions. When causal factors have been identified, who should be notified of the analysis? What requirements and deadlines should there be for generating corrective actions?

Use these requirements to design a reporting system that might help to reduce the risk to patients. Make sure to address all four points requirements in a successful system. This exercise might be done as an in-class exercise or as individual homework and then discussed in the class.

A more time consuming but interesting alternative is to have teams from representing various stakeholders (AECL, the hospitals, the patients, the FDA) design their preferred reporting system as homework and then have these systems presented in class on the same day. Class discussion after these presentations might be a general comparison or some sort of a negotiation among the various parties.

References

Leveson, N. G. (1995). *Safeware: System safety and computers*. New York: Addison Wesley.

Wahlstrom, B., & Swaton, E. (1991). *Influence of organization and management on industrial safety*. Technical report, International Institute for Applied systems Analysis.

Role Playing the Case

Have students read the case, including the background materials. Do not allow student to read any of the accident reports. Assign particular groups to prepare to defend the viewpoint of each of the participants in the case (AECL, FDA, Hospital, Operator). In class, give each group the description of the two Tyler incidents. Also give to them the explanation of the Tyler code and why it produced Malfunction 54. This is what each participant knew shortly after the Tyler accidents.

Allow each group 15 minutes to produce a proposal regarding what should be done. Keep this part of the assignment vague enough to allow them to propose a wide variety of remedies if they desire.

Allow each group 3 minutes to propose its remedy and each group 3 minutes to

comment after hearing all the proposals.

Class discussion can initially center on which proposals are better. Use your knowledge of the case to present the Yakima accidents and ask them which of their proposals would have helped prevent that case. This will allow you to point out the larger issues involved in designing for safety: safety is a system property and not just a property of the software itself.

Notes

Exercises and other materials from ComputingCases.org developed by Dr. Charles Huff of St. Olaf College.

Resource Type

Case Study / Scenario
Educational Activity Description
Historical / Non-Fiction Case

Parent Collection

Therac-25

Topics

Catastrophes, Hazards, Disasters
Ethics and Society
Product Liability
Research and Practice
Safety

Discipline(s)

Computer Sciences
Computer, Math, and Physical Sciences
Research Ethics