

COMPUTINGCASE.ORG

THERAC-25 CASE NARRATIVE AND TEACHING TOOLS

CASE INTRODUCTION	5
<u>THERAC-25 ABSTRACT</u>	5
<u>THERAC-25: A COMPUTER CONTROLLED MEDICAL LINEAR ACCELERATOR FOR TREATING CANCER</u>	5
<u>HOW RADIATION THERAPY WORKS</u>	6
WHAT RADIATION THERAPY IS	6
WHY RADIATION THERAPY WORKS	6
WHAT A TREATMENT SESSION IS LIKE	6
<u>THE MACHINE</u>	7
<u>HOW A MEDICAL LINEAR ACCELERATOR WORKS</u>	7
GENERATING AN ELECTRON BEAM.....	7
GETTING THE BEAM INTO THE BODY	7
RADIATION ABSORBED DOSE	8
<u>HOW THERAC-25 WORKED</u>	8
A SHORT HISTORY OF THERAC.....	8
THE MACHINE IN THE ROOM	9
SWITCHING BETWEEN MODES: THE TURNTABLE.....	9
SETUP AND ACTUATION.....	11
<u>WHAT THERAC-25 SOFTWARE DID</u>	11
REAL-TIME SOFTWARE	11
DESIGN OF SOFTWARE	12
OPERATOR	12
MACHINE	12
SENSORS ON THE MACHINE	12
<u>SYSTEM SAFETY</u>	14
MACHINE-BASED SAFETY MECHANISMS	14
SOFTWARE BASED SAFETY MECHANISMS	14
SAFETY ANALYSIS OF THE SYSTEM	15
<u>PARTICIPANTS</u>	15
<u>LINEAR ACCELERATOR TREATMENT FACILITIES</u>	15
WHAT FACILITIES ARE LIKE.....	15
MACHINE SUPPORT AND MAINTENANCE	16
PRODUCTION PRESSURES	16
LIABILITY AND TRUST	16
<u>FDA</u>	17
INTRODUCTION	17
PRE-MARKET APPROVAL AND PRE-MARKET EQUIVALENCE	17
MEDICAL ERROR REPORTING AND FDA REPORTING REQUIREMENTS	18

FDA ENFORCEMENT TOOLS	18
<u>LINEAR ACCELERATOR OPERATORS</u>	19
WHAT OPERATORS DO	19
DEALING WITH DIFFICULTIES	20
PRESSURES FOR PRODUCTION	20
TRAINING AND LICENSING	20
<u>THE CANADIAN MEDICAL CORPORATION AND THERAC-25</u>	21
DEVELOPMENT OF THERAC-25	21
THERAC-25 GOES TO MARKET	22
SAFETY ANALYSIS OF THERAC-25	22
<u>ACCIDENT ACCOUNTS</u>	22
LINDA KNIGHT: JUNE 3, 1985	22
DONNA GARTNER: JULY 26, 1985	23
JANIS TILMAN: DECEMBER 1985	24
ISAAC DAHL: MARCH 22, 1986	24
DANIEL MCCARTHY : APRIL 11, 1986	25
ANDERS ENGMAN: JANUARY 17, 1987	26
<u>TEACHING SECTION</u>	27
<u>TEACHING INTRODUCTION</u>	27
<u>THERAC-25 INTRODUCTION</u>	27
THERAC-25: SAFETY IS A SYSTEM PROPERTY	27
STRUCTURE OF THE THERAC-25 CASE	28
<i>Introductory materials</i>	28
<i>The machine</i>	28
<i>The Participants</i>	28
USING THE THERAC-25 CASE IN CLASS	28
<u>PITFALLS IN TEACHING WITH THESE CASES</u>	29
THE SEARCH FOR A SINGLE CAUSE	29
SUBJECTIVITY IN THINKING ABOUT CAUSES	29
THE SEARCH TO AFFIX BLAME	30
THE RUSH TO LEGALISM	30
THE LOOSE ASCRIPTION OF "HUMAN ERROR"	30
A FIXATION ON THE TECHNICAL FAILURES (AND FIXES) IN THE SYSTEM	31
<u>ANALYSIS</u>	31
THERAC-25: A SOCIO-TECHNICAL SYSTEM	31
THE SOCIO-TECHNICAL SYSTEM	33
THE MACHINE	33
<i>Hospitals and Clinics</i>	33
<i>Canadian Medical Corporation</i>	33
<u>ETHICAL ANALYSIS</u>	33
USING THE IMPACTCS GRID TO UNDERSTAND THERAC-25	33
QUALITY OF LIFE	34
POWER	34
SYSTEM SAFETY	35

SAFETY AT THE INDIVIDUAL LEVEL.....	35
<i>The programmer</i>	35
<i>The operators</i>	36
SAFETY AT THE GROUP LEVEL.....	36
<i>The Canadian Medical Corporation</i>	36
<i>The Cancer Treatment Facilities</i>	37
SAFETY AT THE NATIONAL LEVEL.....	38
SAFETY AT THE GLOBAL LEVEL	39
PROPERTY	39
PRIVACY	39
EQUITY AND ACCESS	39
HONESTY AND DECEPTION	40
<u>EXERCISES FOR THERAC-25</u>	41
SOME INITIAL CONSIDERATIONS IN TEACHING THIS CASE	41
<u>ANALYZING THERAC-25</u>	41
GATHER DATA	42
ANALYZE THE DATA	42
CONSTRUCT AN ALTERNATIVE SCENARIO.....	42
JUDGE THE ALTERNATIVE.....	42
<u>COMPUTER CONTROL CHOICES EXERCISE</u>	44
CHOOSING THE LEVEL OF COMPUTER CONTROL	44
<u>TRACING THE CODING ERRORS TO THE HAZARDS</u>	46
<u>SOFTWARE SAFETY MYTHS</u>	47
<u>DESIGNING A REPORTING SYSTEM</u>	48
<u>ROLE PLAYING THE CASE</u>	49
<u>SUPPORTING DOCUMENTATION</u>	49
GUIDE TO THE SUPPORTING DOCUMENTS.....	49
CASE HISTORY.....	49
ALIASES	49
LEVESON EXCERPTS.....	49
PRODUCE MALFUNCTION 54.....	50
OPERATOR INTERVIEW	50
REFERENCES	50
<u>A HISTORY OF THE INTRODUCTION AND SHUT DOWN OF THERAC-25</u>	51
CMC'S FDA TESTING AND SAFETY ANALYSIS.....	51
CMC'S RESPONSE TO THE ACCIDENTS.....	52
THERAC-25 IS SHUT DOWN.....	53
CMC MEDICAL GOES INDEPENDENT.....	53
GOVERNMENT AND FDA RESPONSE TO THE ACCIDENTS.....	53
<u>ALIASES IN THERAC-25</u>	54
<u>LEVESON EXCERPTS</u>	55
GENESIS OF THE THERAC-25.....	55
TURNABLE POSITIONING.....	57
THERAC-25 SOFTWARE DESIGN	59

SAFETY ANALYSIS OF THE THERAC-25	61
THE OPERATOR INTERFACE	63
THERAC-25: TYLER, TX SOFTWARE PROBLEM.....	64
THERAC-25: YAKIMA SOFTWARE PROBLEM.....	68
<u>HOW TO PRODUCE A MALFUNCTION 54 ON A [CMC] THERAC-25 LINEAR ACCELERATOR</u>	71
<u>SUMMARY OF OPERATOR INTERVIEW</u>	72
<u>SOME USEFUL THERAC SOURCES</u>	74
ADDITIONAL THERAC SOURCES.....	75

Case Introduction

Therac-25 Abstract

Therac-25 was a new generation medical linear accelerator for treating cancer. It incorporated the most recent computer control equipment. Therac-25's computerization made the laborious process of machine setup much easier for operators, and thus allowed them to spend minimal time in setting up the equipment. In addition to making setup easier, the computer also monitored the machine for safety. With the advent of computer control, hardware based safety mechanisms were transferred to the software. Hospitals were told that the Therac-25 medical linear accelerator had "so many safety mechanisms" that it was "virtually impossible" to overdose a patient.

Normally, when a patient is scheduled to have radiation therapy for cancer, he or she is scheduled for several sessions over a few weeks and told to expect some minor skin discomfort from the treatment. The discomfort is described as being like a mild sunburn over the treated area. But in this case on safety critical software, you will find that some patients received much more radiation than prescribed.

Therac-25: A computer controlled medical linear accelerator for treating cancer

Normally, when a patient is scheduled to have radiation therapy for cancer, he or she is scheduled for several sessions over a few weeks and told to expect some minor skin discomfort from the treatment. The discomfort is described as being like a mild sunburn over the treated area.

Therac-25 was a new generation machine that incorporated the most recent computer control equipment. The machine targeted electron or X-ray beams on cancerous tissue to destroy it. Electron beams were used to treat shallow tissue, while X-ray beams could penetrate with minimal damage to treat deep tissue.

When a doctor decides that a patient needs radiation therapy, that patient is given a prescription that indicates to the medical linear accelerator operator how many rads (radiation absorbed dose) the patient should receive over how many total treatments. In addition, the prescription indicates the location where the radiation should be applied. The patient schedules a time (or times) to receive treatment.

Standard procedures then determine whether, on any particular appointment, the operator is to set up the equipment for electron or X-ray beam treatment. The patient is asked to lie in the appropriate position on the treatment table and the table is rotated to place the diseased part of the patients' body in the path of, and at the appropriate distance from, the radiation beam. The operator then does whatever mechanical setup is required and leaves the room to program the treatment data into the machine. After doing this, the operator presses the button that actuates the treatment routine. The patient is then helped off the treatment table and ushered out. After the appropriate forms have been filled out, the next patient is admitted.

Therac-25's computerization made this laborious process much easier for operators, and allowed them to spend minimal time in setting up the equipment. Operators were thus freed to spend more time talking with and helping the patient.

In addition to making setup easier, the computer also monitored the machine for safety. Previous machines had safety devices as a part of the hardware of the machine. Among other things, these safety devices kept the machine from delivering doses of radiation that were too high. So, with the advent of computer control, these hardware based safety mechanisms were transferred to the software. Hospitals were told that the Therac-25 medical linear accelerator had "so many safety mechanisms" that it was "virtually impossible" to overdose a patient.

How Radiation Therapy Works

What Radiation Therapy Is

Radiation therapy for cancer is the exposure of cancerous tissue to ionizing radiation. This is usually done by what is called "external" therapy, using electron, X-rays or gamma rays to treat the tissue. This therapy may occur either before or after surgery, or in the place of surgery.

Therac-25 was a 3rd generation radiation therapy machine for external radiation therapy. It used either electron beam or X-rays to treat tissue.

Why Radiation Therapy Works

Cancer cells usually multiply faster than most other cells in the body. Tissue composed of these quickly-dividing cells can be shrunken by disabling its genetic material. By doing this, ionizing radiation interferes with the cancerous tissue's ability to grow.

Unfortunately, the radiation makes no distinction between cancerous cells and other rapidly dividing body tissues. Skin and hair are some of the most noticeably hurt tissues after treatment, and treatment may produce skin lesions and hair loss. These tissues have cells that rapidly divide and the radiation halts their development. But they are usually able to recover from this assault and return to normalcy. Nevertheless, skin lesions and hair loss are not an unusual side effect of radiation therapy.

What a Treatment Session is Like

Radiation therapy is usually done in a series of sessions occurring over several weeks, allowing the effect of the radiation to build up over time. The treating doctor will determine the specific number of treatments, the dosage at each treatment, and the schedule. During treatment, the doctor will usually see the patient once a week to check on general health, side effects, and the progress of the treatment.

Before the series of treatments occurs, a radiation therapy technician will outline the specific area to be treated with a marking pen, indelible ink or silver nitrate.

Depending on the body area to be treated, the patient would need to remove his or her clothing and put on a hospital gown. After going to the radiation therapy room, they would then either lie on a treatment table or sit in a special chair (Therac-25 had a table). The marks on the skin are used to guide the


machine operator in locating the precise area to be treated. Once the machine is set up, the operator leaves the room for a control room nearby. This protects the operator from prolonged exposure to low-level radiation that might scatter from the machine (an operator may treat as many as 30 patients in a day). From there, the operator will turn on the treatment machine while he/she watches. With the Therac-25, this was accomplished by means of a television camera and monitor. During radiation therapy, the treatment machine makes a buzzing noise. Treatments are typically brief and painless, normally lasting 1 to 5 minutes. Total time in the treatment room will usually be 5 to 15 minutes.

The Machine

How a Medical Linear Accelerator Works

Generating an Electron Beam

Early radiation therapy machines used a radioactive source like cobalt to produce the ionizing radiation needed to treat cancerous tissue. Some machines still use an active radiation source. But most radiation therapy today is done with a linear accelerator. In principle, a linear accelerator works just



like the computer monitor you are probably using to read this web page. The electrons are accelerated by the gun in the back of the monitor and directed at the inside of the screen, where phosphors absorb the electrons and produce light. A medical linear accelerator produces a beam of electrons about 1,000 times more powerful than the standard computer monitor. The longer a linear accelerator is, the higher the energy of the beam it can produce. The innovation of Therac 25 was that the designers found a way to fold the beam back and forth so a very long accelerator could be fit into a smaller space. Thus powerful beams could be produced, but within a reasonable amount of space

Getting the Beam into the Body

Patients can be treated directly with the resulting electron beam, as long as the beam is spread out by scanning magnets to produce a safe level of radiation. The medical linear accelerator spreads and directs the beam at the appropriate place for treatment. The picture below shows a typical medical linear accelerator in operation.

But a difficulty with the electron beam is that it diffuses rapidly in tissue and cannot reach deeper tissue for treatment. The picture below is a simulation (produced by the Stanford Linear Accelerator

Center) of an electron beam traveling through air and entering human tissue. You can see the beam quickly diffuses and therefore does not penetrate deeply.



To solve this problem, Therac-25 and many other machines can switch to a mode in which X-ray photons are used for treatment. These penetrate much more deeply without harming intervening tissue. To do this, the electron beam is greatly increased in intensity and a metal foil followed by a beam “flattener” is placed in the path of the electron beam. This transforms the electron beam into an X-ray (called photons in some literature). This process is inefficient and requires a high intensity electron beam to produce enough X-ray intensity for treatment. Therac-25 used

a 25 MeV electron beam to produce an X-ray for treatment. 25 MeV is 25 million electron volts (eV -- an eV is the energy needed to move one electron through a potential of one volt).

Therac-25 was what was called a dual-mode machine. It could produce the low energy electron beams for surface treatment and it could also produce a very high intensity electron beam that would be transformed into an X-ray by placing the metal foil in the path of the beam. The serious danger in a dual mode machine is that the high-energy beam might directly strike the patient if the foil and flattener were not placed in its way.

Radiation Absorbed Dose

Although MeVs are used to measure the strength of the electron beam, the measure used for therapeutic uses is the radiation absorbed dose (rad). This is a measure of the radiation that is absorbed by tissue in a treatment. Standard single radiation treatments are in the range of 200 rads. 500 rads is the accepted level of radiation that, if the entire body is exposed to it, will result in the death of 50% of the cases. The unprotected electron beam in the Therac-25 is capable of producing between 15,000 and 20,000 rads in a single treatment. The unprotected beam is never aimed directly at a patient. It is either spread to a safe concentration by scanning magnets or turned into X-rays and reduced by a beam flattener.

How Therac-25 worked

A Short History of Therac

There were two previous versions of Therac machines, each produced by CMC in collaboration with a French company, CGR. Therac 6 and Therac 20 (each named for the MeV they could produce) were based on earlier design from CGR. By the time Therac-25 was released for sale, CMC had 13 years of experience with production of medical linear accelerators. Therac-25 was based on these previous versions. Its main innovations were (1) a “double pass” electron beam so the machine could produce more energy in less space, and (2) the addition of extensive computer control of the machine. This latter innovation allowed CMC to move much of the checking for hazardous conditions into the software.

The Therac-25's ancestors, Therac-20 and Therac-6, had used a minicomputer (a DEC PDP-11) to add some convenience to the standard hardware of a medical linear accelerator. They both could work

without computer control. CMC determined to make its new model, Therac-25, a tightly-coupled combination of software and hardware. Therac-25 software was not written from scratch, but was built up from components that were borrowed from the earlier versions of Therac.

The Machine in the Room

Therac-25 is not just a machine, but an installation consisting of the machine, the PDP-11 that controlled the machine, the shielded room the machine sits in, and the monitoring and control station.




Figure 1. Typical Therac-25 facility

The control console and printer etc. are all located outside the heavily shielded treatment room. Thus, when pressing the key to begin the treatment, the operator does not have any direct access to the machine or the patient. All the occurrences in the treatment room must be observed through the TV monitor and the intercom. The intercom works both ways, that is, the patient can hear the operator (if the operator presses a switch) and the operator can hear the patient. The patient, however, cannot see anything outside the treatment room, while the operator can look in using the TV monitor.

Switching Between Modes: The Turntable

Therac-25 is a dual mode machine. This means that it can treat the patient with relatively low energy electron beams or with X-ray beams. In addition, Therac-25 had a “field light” position that allowed a standard light beam to shine in the path of treatment to help the operator in setting up the machine.

Thus there were three modes in which the Therac-25 could operate: electron beam and X-ray for treatment, and field light for setup.

Even though they are relatively low energy, the electron beams are too powerful in their raw form to treat the patient. They need to be spread thinly enough to be the right level of energy. To do this, Therac-25 placed what are called scanning magnets in the way of the beam. The spread of the beam (and also its power) could be controlled by the magnetic fields generated by these magnets. Thus for electron beam therapy, the scanning magnets needed to be placed in the path of the beam.

X-ray treatment requires a very high intensity electron beam (25 MeV) to strike a metal foil. The foil then emits X-rays (photons). This X-ray beam is then “flattened” by a device below the foil, and the X-ray beam of an appropriate intensity is then directed to the patient. Thus, X-ray therapy requires the foil and the flattener to be placed in the path of the electron beam.

The final mode of operation for Therac-25 is not a treatment mode at all. It is merely a light that illuminates the field on the surface of the patient’s body that will be treated with one of the treatment beams. This “field light” required placing a mirror in place to guide the light in a path approximating the treatment beam’s path. This allowed accurate setup of the machine before treatment. Thus, for field light setup, the mirror needed to be placed in the path where one of the treatment beams would eventually go.




Figure B. Upper turntable assembly

In order to get each of these three assemblies (scanning magnets or X-ray target or field light mirror) in the right place at the right time, the Therac-25 designer placed them on a turntable. As the name suggests, this is a rotating assembly that has the items for each mode placed on it. The turntable is rotated to the correct position before the beam is started up. This is a crucial piece

of the Therac-25 machine, since incorrect matching of the turntable and the mode of operation (e.g. scanning magnets in place but Electron beam turned on high for X-ray) could produce potentially fatal levels of radiation.

Setup and Actuation

The Therac-25 operator sets up the patient on the table using the field light to target the beam. In doing this, treatment parameters must be entered into the machine directly in the treatment room.

He or she then leaves the room and uses the computer console to confirm the treatment parameters (electron or X-ray mode, intensity, duration, etc.). The parameters initially entered in the treatment room appear on the console and the operator simply presses return to confirm each one.

The computer then makes the appropriate adjustments in the machine (moving the turntable, setting the scanning magnets, setting beam intensity etc.). This takes several seconds to do. If the operator notices an error in the input parameters, he or she can, during the setup, edit the parameters at the console without having to start all over again from inside the treatment room.

When the computer indicates that the setup has been done correctly, the operator presses the actuation switch. The computer turns the beam on and the treatment begins. There are three possible outcomes at this point, and they all depend on sensors on the machine. If the sensors indicate no trouble, the treatment concludes successfully. If the sensors indicate a minor problem, like the beam being slightly out of tune, the computer turns the beam off immediately. The operator can then press a “proceed” key to retry the treatment up to 5 times. If the sensors indicate a more serious malfunction, like the beam being significantly stronger or weaker, the computer turns the beam off immediately and requires the machine to be completely setup from the beginning.

What Therac-25 Software Did

Real-time Software

The software that ran the Therac-25 was real-time software. What does that mean?

Real-time software is software that interacts with the world on the world’s schedule, not the software’s. For instance, software to keep a radio tuner on the signal of a drifting station could take two approaches. It might simply update the signal every 0.1 seconds, searching for the strongest signal within some bandwidth. Another approach is to include a sensor that detects when the signal loses strength and only then search for a stronger signal nearby. This latter approach is real-time. It senses the world and responds to changes in the world when those changes occur.

This sort of software (even the simple system just described) is difficult to write and maintain. First, it involves the software in reading and responding to sensors about the state of “the world.” With Therac-25, these sensors indicated things like the intensity of the beam, the position of various parts of the machine (e.g. the turntable) and commands entered at the console by the operator. Sensors, of course, can go bad, or give incorrect readings. When they do, the software needs to be able to detect these problems and respond accordingly, or at least fail in a graceful manner that doesn’t endanger life.

In addition, when real-time software has to monitor more than one thing, changes in one area may occur while the software is responding to changes in another. This is like the situation of trying to divide your limited attention to all the things you need to monitor when you are driving a car. While you are watching a red light up ahead, a car may have slipped into your blind spot without you seeing it.

So, Therac software needed to track and respond to several things in real-time without dropping any important balls. What those things are is described in the next section

Design of Software

The main tasks for which the software is responsible include:

Operator

1. Monitoring input and editing changes from an operator
2. Updating the screen to show current status of machine
3. Printing in response to an operator commands

Machine

1. monitoring the machine status
 2. placement of turntable
 3. strength and shape of beam
 4. operation of bending and scanning magnets
5. setting the machine up for the specified treatment
6. turning the beam on
7. turning the beam off (after treatment, on operator command, or if a malfunction is detected)


The Therac-25 software is designed as a real-time system and implemented in machine language (a low level and difficult to read language). The software segregated the tasks above into critical tasks (e.g. setup and operation of the beam) and non-critical tasks (e.g. monitoring the keyboard). A scheduler handled the allocation of computer time to all the processes except those handled on an interrupt basis (e.g. the computer clock and handling of computer-hardware-generated errors).

As explained above, the difficulty with this kind of software is the handling of things that might be occurring simultaneously. For example, the computer might be setting the magnets for a particular treatment already entered (which can take 8 seconds) while the operator has changed some of the parameters on the console screen. If this change is not detected an incorrect treatment can be given. More dangerous is the possibility that the change only affects the portion of the software that handles beam intensity, while the portion of the software that checks turntable position is left thinking that the old treatment parameters are still in effect.

Sensors on the Machine

The sensors in the machine reported on, among other things, the placement of the turntable and the strength and shape of the beam. In the diagram below, you can see the “transmission monitors” directly below the metal foils designed to produce X-rays. A different monitor was required for X-rays

than for the electron beam, and so these monitors (they were ion chambers) were attached to the turntable underneath either the X-ray foil of the electron beam scanning magnets. No monitor was placed below “field light assembly” and so no measurement can be made of a beam in this position. But then, no beam is supposed to be turned on in this position, only a light.



Monitoring of the position of the turntable is done by sensors at the turntable (in the diagram above, in the place where the foils are shown).

System Safety

Machine-Based Safety Mechanisms




Figure 1. Typical Therac-25 facility

As the diagram indicates, the Therac-25 linear accelerator was isolated in a heavily shielded room. This shielding protected the operator (who might do as many as 30 treatments in one day) from the low-level radiation that might scatter from the machine. In addition, the machine itself was shielded in many ways to reduce the amount of scattered radiation it would emit. CMC was particularly proud of this innovation in machine shielding, and even published a paper in a technical journal on its design.

Software Based Safety Mechanisms

Previous versions of Therac (Therac-6 and Therac-20) used software to make the hand operation of the machine more convenient. But Therac-25 was completely software controlled. In addition and safety checking was made the job of the software many of the hardware safety interlocks were removed. Thus, the safe operation of the machine became almost completely the responsibility of the software.

For example, intensity of the beam is monitored by ion chambers placed on the turntable. There were two different ion chambers, one located beneath the scanning magnets that spread the electron beam and one located beneath the foil that turned a high intensity electron beam into X-rays. These chambers monitored the amount of radiation that was being delivered to the patient in each mode (electron beam or X-ray) and each could measure the beam intensity only within the expected range from the beam with which it was paired. If the chamber detected a dose that was different from that assigned to the patient, the software immediately suspended treatment.

If the difference was a minor amount or if the beam intensity was measured as hardly there, the software might allow the operator to retry the treatment up to 5 times before shutting down completely. This retry facility was added to the software because it was a regular occurrence for the beam to be slightly “out of tune” and for the software to suspend treatment.

If the beam intensity was detected to be quite different from the assigned intensity, the software shut the machine down completely and required all the treatment parameters to be entered again.

Safety Analysis of the System

In 1983, just after CMC made the Therac-25 commercially available, CMC performed a safety analysis of the machine using Fault Tree Analysis. This involves calculating the probabilities of the occurrence of varying hazards (e.g. an overdose) by specifying which causes of the hazard must jointly occur in order to produce the hazard.

In order for this analysis to work as a safety analysis, one must first specify the hazards (not always easy), and then be able to specify the all possible causal sequences in the system that could produce them. It is certainly a useful exercise, since it allows easy identification of single-point-of-failure items and the identification of items whose failure can produce the hazard in multiple ways. Concentrating on items like these is a good way to begin reducing the probabilities of a hazard occurring.

In order to be useful, a Fault Tree Analysis needs to specify all the likely events that could contribute to producing a hazard. In addition, if one knows the specific probabilities of all the contributing events, one can produce a reasonable estimate of the probability of the hazard occurring.

Since much of the software had been taken from the Therac-6 and Therac-20 systems, and since these software systems had been running many years without detectable errors, the analysts assumed there were no design problems in the software. The analysts did consider software failures like "computer selects wrong mode" but assigned them probabilities like 4×10^{-9} . These sorts of probabilities are likely assigned based on the remote possibility of random errors produced by things like electromagnetic noise. They do not take into account the possibility of design flaws in the software

Participants

Linear Accelerator Treatment Facilities

What Facilities are like

Cancer treatment facilities are often housed in large hospitals, but some are stand-alone cancer treatment centers. Those associated with hospitals are more likely to be non-profit, while those that stand alone are more likely to be for-profit organizations. Financial pressures are likely to be strong at both for-profit and not-for-profit organizations, but they will have slightly different regulatory structures.

During the time of Therac-25 (the mid 80s) a well equipped treatment facility might have 3 different machines. The machines would be capable of producing different kinds of radiation, different

strengths of beam, and capable of different kinds of exposure to the patient. Each of these machines would cost, for the machine alone, between 1 and 2 million dollars. In addition, special housing for each machine is needed, with shielding in the walls, adequate power supply, video and intercom links, etc.

Operators would be needed to run each machine. For larger facilities, a supervisor of the operators, with more training and experience might be needed. In addition, at least one MD specialist in cancer radiation therapy (a Radiation Oncologist) would be required. Finally, a medical physicist would be needed to maintain and check the machines regularly. Some facilities contract out the services of a medical physicist. Finally, all the support personnel for these specialists (nurses, secretaries, administrative staff, people to handle billing and paperwork, janitorial staff, etc.) are required.

Machine Support and Maintenance

Medical Linear Accelerators do age over time, and older machines often produce more errors. Five to ten years is a reasonable life span for a machine. Thus, simply to maintain a set of three medical linear accelerators, an institution can expect to spend 1 to 2 million dollars every third year.

Sometimes errors can be resolved and machine kept longer using software upgrades or upgrades or retrofits of machine parts. The companies that sell linear accelerators charge maintenance contracts that can include different levels of support. Because of monetary constraints, sometimes facilities are forced to choose between software updates, manuals, and training for operators and physicists. All this is in addition to the price of the machine itself.

Production Pressures

Production pressures are always present when an expensive medical technology is being used. These very expensive machines need to treat enough patients to pay for themselves over their lifetime. And in for-profit medical facilities the additional pressure of generating a profit is added to this production pressure. Another kind of production pressure is generated because of concern for the patient. Patients' schedules require treatments on certain days and it disrupts the patients' lives and slows down their treatment to have to reschedule them for another day while the machine is being checked out.

These production pressures generate the desire to "push patients through." If a machine gives only a portion of the prescribed dose, an operator will often repeat the treatment with enough radiation to add up to the total prescribed dose. Of course, because of liability issues and concerns for patient welfare, this can only be done when it is thought safe.

One of the advantages of the significant computerization of the Therac 25 machine was that setup for treatment could be done much more quickly. This allowed the operator more time to speak with the patient and interact with them about their health concerns. In addition, this increased efficiency allowed more patients to be scheduled during a day. Thus, more patients could be treated, but the atmosphere was not reduced to that of a factory.

Liability and Trust

Facilities that run medical linear accelerator are surely concerned about liability for injury to patients that might occur. Insurance, for medical providers, is quite expensive and errors in treatment can

result in lawsuits, which in turn produce increases in insurance premiums. Standard practice in litigation is to “sue everyone with deep pockets.” This means that even if an error is the result of poor design of a linear accelerator, the facility itself will be sued simply because they were involved: they have insurance and thus “deep pockets.”

But it is in the interest of facilities to reduce errors without the threat of lawsuits. When a treatment must be restarted several times because of errors, it may reduce patient confidence in the facility. This can mean patients moving to another facility with which they are more comfortable.

Finally, medical professionals are in their business because they want to help people and have the knowledge and skill to do so. So a primary motivation of medical professionals is patient welfare.

FDA

Introduction

The *Food and Drug Administration* (FDA) was created when Congress passed the *Food and Drugs Act* in 1906. This act was the first of a series of laws and amendments that gave the FDA jurisdiction over the regulation of foods and patent medicines. In 1938, Congress strengthened and expanded the FDA, to include the regulation of therapeutic and medical devices within its jurisdiction.

The FDA's *Bureau of Medical Devices and Diagnostic Products* was created in 1974, and soon operated in conjunction with the *Medical Devices Amendments* of 1976. The amendments helped to clarify the logistics of the regulation of medical devices, and required the FDA to "ensure their safety and effectiveness."

Radiation had been recognized as a health hazard since before World War I, and the FDA monitored the health risks that radiation emitting products posed to America's workers and consumers. As FDA's responsibilities for monitoring radiological devices grew, a bureau within the FDA called the *Center for Devices and Radiological Health* (CDRH) was established.

In 1980 the FDA's budget had swelled to over \$320 million, with a staff of over 7,000. Many bureaus controlled areas such as biological drugs, consumer products, public health standards, and veterinary medicines.

Pre-Market Approval and Pre-Market Equivalence

FDA approved medical devices before they “went to market.” This was called Pre-Market Approval and was a somewhat complex process. In the FDA Pre-market Approval scheme, devices were organized into three classes, as established by the 1976 *Medical Device Amendments*.

4. Class I devices, "general controls provide reasonable assurance of safety and effectiveness," for example bedpans and tongue depressors
5. Class II devices, such as syringes and hearing aids, "require performance standards in addition to general controls"
6. Class III devices like heart valves and pacemakers are required to undergo pre-market approval as well as complying with general controls

In addition to classifying devices as Class I, II, or III, FDA approved devices for market in one of two ways:

8. Proof of Pre-market Equivalence to another device on the market, termed 501(k)
9. OR Pre-market Approval (Rigorous Testing)

If a company could show Pre-market Equivalence (proof that a new product was equivalent to one already on the market), the new product could be approved by FDA without extensive, costly, rigorous testing. In 1984 about 94% of medical devices came to market through Pre-market Equivalence.

If a product was not equivalent to one that was already on the market, FDA required that the product go through testing to gain Pre-market Approval. In 1984 only about 6% of medical devices were required to go through this testing.

Thus, it was clearly in the interest of medical device producers to show that their product had pre-market equivalence. The Therac-25, brought to market in 1983, was classified as a Class II medical device. Since Canadian Medical Company (CMC), designed the Therac-25 software based on software used in the earlier Therac-20 and Therac-6 models, Therac-25 was approved by FDA under Pre-market Equivalency.

Medical Error Reporting and FDA Reporting Requirements

A 1983 General Accounting Office (GAO) report criticized the FDA's "adverse experience warning system" as inadequate. FDA had published reports about potential hazards, including reports in their own newsletter, *The FDA Consumer*. The FDA implemented the mandatory medical-device reporting rule after Congress passed the *Medical Device Reporting Legislation* in 1984. This rule required manufacturers to report injuries and problems that could cause injuries or death.

Before 1986, users of medical devices (hospitals, doctors, independent facilities) were not required to report problems with medical devices. Instead, under the medical device reporting rule, manufacturers of these devices were required to report problems. The idea was that manufacturers would be the first to hear about any problems with the devices they made and that therefore reports would be timely. In addition, manufacturers would be most likely to have the correct information needed about a device to help resolve difficulties.

FDA Enforcement Tools

In the mid-1980s, the FDA's main enforcement tools for medical devices already on the market were publicity. The FDA could not force a recall, it could only recommend one. The CDRH (*Center for Devices and Radiological Health* monitors radiological devices) issues its public warnings and advisories in the *Radiological Health Bulletin*. Before issuing a public warning or advisory, the FDA could negotiate with manufacturers in private (and in the case of Therac 25, with regulatory agencies in Canada). In response to reports of problems with a medical device, the FDA could, in increasing order of severity:

1. Ask for information from a manufacturer.
2. Require a report from the manufacturer.
3. Declare a product defective and require a *corrective action plan* (CAP).
4. Publicly recommend that routine use of the system on patients be discontinued.
5. Publicly recommend a recall.

In deciding on the response to a problem with a device, FDA needed to consider:

- Safety of the public.
- Safety of users of the device.
- Need for medical treatment with the device.
- Impact of the decision on the individual manufacturer.
- Impact of the decision on the medical device industry.

Linear Accelerator Operators

What Operators Do

Operators are the primary persons involved in the actual administration of radiation therapy. The treating doctor (usually called a Radiation Oncologist) is responsible for prescribing and planning the treatment and for weekly checkups on the health of the patient. The Linear Accelerator Operator is responsible primarily for seeing that the prescribed treatment is carried out appropriately when the patient shows up for a treatment.

Operators are thus, usually responsible for treatment done with one (or a small set of) machines. The schedule is maintained by others, and this places the operator in the position of a “production assistant” making sure that all those persons scheduled for treatment on a particular day get treated. In addition, they have a responsibility to the patient to operate the machine safely and to treat the patient kindly and with respect. This mix of goals is not unusual in medical practice.

The Therac-25 operator greets the patient on arrival, escorts them into the treatment room and sets up the patient on the treatment table using the field light to target the beam. This may involve marking the patient’s skin for the pattern of radiation that is required. The operator then enters treatment parameters into the machine directly in the treatment room. He or she then leaves the room and uses the computer console to confirm the treatment parameters (electron or X-ray mode, intensity, duration, etc.). The computer then makes the appropriate adjustments in the machine (moving the turntable, setting the scanning magnets, setting beam intensity etc.). This takes several seconds to do. If the operator notices an error in the input parameters, he or she can, during the setup, edit the parameters at the console without having to start all over again from inside the treatment room.

When the computer indicates that the setup has been done correctly, the operator presses the actuation switch. The computer turns the beam on and the treatment begins. When treatment is over, the operator checks with the patient, updates records on that patient and then admits the next patient into the treatment room.

One of the advantages of the significant computerization of the Therac-25 machine was that setup for treatment could be done much more quickly. This allowed the operator more time to speak with the patient and interact with them about their health concerns. In addition, this increased efficiency allowed more patients to be scheduled during a day. Thus, more patients could be treated, but the atmosphere was not reduced to that of a factory.

Dealing with difficulties

If a treatment resulted in a suspend or cancellation by the machine, the operator had several choices. For some machine errors, the operator could simply press the “retry” button and attempt the therapy over again. If only half the prescribed dose had been introduced (e.g. the beam was a lower intensity or cut off early) the rest of the dose might be applied in a second, immediate, treatment.

If the error was more significant, many hospitals and facilities would have a medical physicist on call. The physicist could be called in to look at the machine immediately. For facilities without a full time physicist, contract service was usually provided. This required scheduling (but usually within the same day as the problem).

All errors (whether by the machine or by the operator) were supposed to be logged and reported. Medical Linear Accelerators do age over time, and older machines often produce more errors. Five to ten years is a reasonable life span for a machine. Close tracking of these errors by operators allows the hospital or facility to know when to replace a machine that is generating more errors than is acceptable. Even if errors are not harmful to patients, when a treatment must be restarted several times, it may reduce patient confidence in the facility.

Pressures for Production

Production pressures are always present when an expensive medical technology is being used. Machines need to treat enough patients to pay for themselves over their lifetime. And in for-profit medical facilities the additional pressure of generating a profit is added to this production pressure. Another kind of production pressure is generated because of concern for the patient. Patient schedule requires treatments on certain days and it disrupts the patients’ lives and slows down their treatment to have to reschedule them for another day while the machine is being checked out.

These production pressures generate the desire to “push patients through.” If a machine gives only a portion of the prescribed dose, an operator will often repeat the treatment with enough radiation to add up to the total prescribed dose. Sometime this repeat has been done up to twelve times to produce the appropriate treatment with a balky machine. At times, operators have been known to collaborate with medical physicists to use jumper cables to override a particular safety mechanism, if their judgment is that the override will not reduce safety.

Operators who feel that pressures for production have decreased safety can certainly report this to their supervisors (usually a supervising operator with additional training and experience). They also have been known to leave facilities because of concern over safety.

Training and Licensing

There is currently no industry-wide standard certification and education for medical linear accelerator operators. There are about 102 radiation schools in the country, ranging from certificate programs (about 12 months in length) to four-year bachelor's degrees. Licensing standards differ from state to state. In some states, operators are required to be licensed by the American Registry of Radiologic Technologists (ARRT). This licensing requires a certified educational program and regular updating of skills for re-registration of the license.

Other states, however, have designed their own tests to set minimal standards for operators, and some of these tests are much less involved than that required by ARRT. In addition, many of these states do not require continuing education of operators. There are no national standards for training or licensing.

Operators who are licensed have more professional standing to resist production pressures that they feel lead to unsafe treatment of patients. In addition, their training gives them better arguments to stand up to hospital administrations that attempt to put pressure on technicians to push large numbers of patients through treatment in spite of possible dangers.

The Canadian Medical Corporation and Therac-25

Development of Therac-25

The story of Therac-25 begins in the early 1970's when Canadian Medical Corporation (CMC) joined forces with a French company, CGR, to design and build a medical linear accelerator based on earlier CGR machines. The companies cooperated on the design and manufacture of two successful medical linear accelerators, the Therac-6 and its successor, the Therac-20. Both these machines were based on CGR designs that did not use computer control. The new machines added computer control, in addition to other innovations. The Therac-6 was the initial product of their collaboration and was designed to produce X-rays for radiation therapy. The Therac-20 was a much more powerful and versatile machine. It could produce two different kinds of radiation beams for treatment of deep and shallow tissue. CMC also produce other medical linear accelerators, including the Therac-4, a single mode electron beam machine.

Development of Therac-25

In the early 1980's, CMC developed a much more space-efficient medical linear accelerator that was just as powerful and versatile as the Therac-20. Linear accelerators are more powerful the longer they are, and CMC found a way to fold the long beam-producing mechanism for a 25 MeV machine into a smaller space. In addition, this new version was somewhat less expensive to produce, since it used a less expensive beam production device (a magnetron instead of a klystron).

Finally, CMC intended to take advantage of increasing capability of computer software to make the machine easier to operate. The new Therac-25 was the result of a convergence of the new beam-folding technology with the ease of computer control, bringing with it the bonus of lower production costs. In addition to lower production costs, the computer control allowed faster setup of the machine for each patient. This meant that more patients could be treated in one day than with non-computerized linear accelerators.

The Therac-25's ancestors, Therac-20 and Therac-6, had used a minicomputer (a DEC PDP-11) to add some convenience to the standard hardware of a medical linear accelerator. They both could work without computer control. CMC determined to make its new model, Therac-25, a tightly-coupled combination of software and hardware. By this time, its collaboration with CGR had grown stale and CMC was bringing in its new beam folding technology (and the new Therac-25) on its own.

In tightly coupling the software and the hardware, CMC could use the software to monitor the state of the machine for proper operation and for safety. Previous versions, with designs based in models that predated computer control, had included independent circuits to monitor beam scanning and had

mechanical interlocks to ensure the machine could not enter a state in which it could harm a patient. But with increased computer control, CMC decided not to duplicate this equipment in the Therac-25 (with some cost savings), and to rely on software for policing these safety issues.

Therac-25 goes to Market

In late 1982, Therac-25 was first offered to hospitals in a commercial version. It was eventually adopted by eleven institutions, six in Canada and five in the US. These included sites in Georgia, Texas, Washington State, and Hamilton, Ontario.

Safety Analysis of Therac-25

In 1983, just after CMC made the Therac-25 commercially available, CMC performed a safety analysis of the machine using Fault Tree Analysis. This involved calculating the probabilities of the occurrence of varying hazards (e.g. an overdose) by specifying which causes of the hazard must jointly occur in order to produce the hazard.

Since much of the software had been taken from the Therac-6 and Therac-20 systems, and since these software systems had been running many years without detectable errors, the analysts assumed there were no design problems in the software. The analysts did consider software failures like "computer selects wrong mode" but assigned them probabilities like 4×10^{-9} . These sorts of probabilities are likely assigned based on the remote possibility of random errors produced by things like electromagnetic noise. They do not take into account the possibility of design flaws in the software.

Accident Accounts

Linda Knight: June 3, 1985

61-year old Linda Knight had been receiving follow-up treatment at the Kennestone Regional Oncology Center (Marietta, GA) for the removal of a malignant breast tumor. On June 3, staff at Kennestone prepared Knight for electron treatment to the clavicle area, using the Therac-25 machine.

Knight had been through the process before, which was ordinarily uneventful. This time, when the machine was turned on, Knight felt a "tremendous force of heat... this red-hot sensation." When the technician re-entered the therapy room, Knight said, "you burned me." The technician replied that that was "not possible."

Back home, the skin above Knight's left breast began swelling. The pain was so great that she checked in at Atlanta's West Paces Ferry Hospital a few days after the Therac incident. For a week, doctors at West Paces Ferry continued to send Knight back to Kennestone for Therac treatment, but when the welt on her chest began to break down and lose layers of skin, Knight refused to undergo any more radiation treatment.

About two weeks later, the physicist at Kennestone noticed that Knight had a matching burn on her back, as though the burn had gone through her body. The swelling on her back had also begun to slough off skin. Knight was in great pain, and her shoulder had become immobile. These clues led the physicist to conclude that Knight had indeed suffered a major radiation burn. Knight had probably received one or two radiation doses in the 20,000-rad (radiation absorbed dose) range, well above the typical prescribed dosage of around 200-rads. The physicist called CMC and, without telling of the

accident, asked questions about the likelihood of radiation overexposure from the Therac 25 machine: Could Therac 25 operate in electron mode without scanning to spread the beam? Three days later CMC engineers called back to say this was not possible.

Linda Knight was in constant pain, lost the use of her shoulder and arm, and her left breast had to be removed because of the radiation burns.

Donna Gartner: July 26, 1985

Donna Gartner, a 40-year old cancer patient, was at the Ontario Cancer Foundation clinic in Hamilton, Ontario, Canada for her 24th Therac treatment for carcinoma of the cervix.

The Therac-25 operator activated the machine, but after 5 seconds, the Therac-25 shut down and showed an "H-tilt" error message. The computer screen indicated that no dose had been given, so the operator hit the "P" key for the "proceed" command. The Therac shut down in the same manner as before, reading "no dose," so the operator repeated the process a total of four times after the initial try.

After the fifth try, a hospital service technician was called but found no problems with the machine. Donna Gartner left the clinic and the Therac was used with six other patients that day without any incidents. However, despite the fact that the Therac had indicated that no radiation dose had been given during Donna Gartner's five therapy attempts that day, Gartner complained of a burning sensation she described as an "electric tingling shock" in the treated area of her hip.

Gartner returned for treatment three days later, on July 29, and was hospitalized for suspected radiation overexposure. She had considerable burning, pain and swelling in the treatment region of her hip.

The Hamilton clinic took the Therac-25 machine out of service and informed CMC of the incident. This was the first time CMC had heard from a clinic about an overdose problem with the Therac-25 machine. CMC sent a service engineer to investigate.

CMC reported to a range of stakeholders that there was a problem with the operation of Therac 25. The FDA, the Canadian Radiation Protection Board (the parallel Canadian agency to the FDA), and other Therac-25 users were all notified. Users were instructed to visually confirm that the Therac turntable was in the correct position for each use.

Because of the Hamilton accident, CMC issued a voluntary recall of the Therac-25 machines and the FDA audited CMC's modifications to the Therac. CMC could not reproduce the malfunction that had occurred but suspected some hardware errors in a switch that monitored the turntable position. A failure of this switch could result in the turntable being incorrectly positioned, and an unmodified electron beam striking the patient. The company redesigned the mechanism used to lock the turntable into place, redesigned the switch to detect position and its accompanying software. They then reported in November 1985 that this redesign was complete and that, given their safety analyses, the machine was now at least 10,000 times safer than before.

Donna Gartner died on November 3, 1985 from cancer. An autopsy revealed that had the cancer not killed Gartner, a total hip replacement would have been necessary because of the radiation overexposure.

Janis Tilman: December 1985

Janis Tilman was being treated with the Therac-25 machine at the Yakima Valley Memorial Hospital in Yakima, Washington. After one treatment in December 1985, her skin in the treatment area, her right hip, began to redden in a parallel striped pattern. The reddening did not immediately follow treatment with the Therac-25 because it generally takes at least several days before the skin reddens and/or swells from a radiation overexposure.

Tilman continued Therac treatment until January 6, 1986 despite the reddening, since it was not determined that the reddening was an abnormal reaction. Hospital staff monitored the skin reaction and searched unsuccessfully for possible causes for the striped marks.

The hospital sent a letter to CMC and spoke on the phone with CMC's technical support supervisor, who later sent a written response stating, "After careful consideration, we are of the opinion that this damage could not have been produced by any malfunction of the Therac-25 or by any operator error." The hospital staff dismissed the skin/tissue problem as "cause unknown," partly due to the response from CMC, and partly because they knew CMC had already installed additional safety devices to their Therac-25 machine in September 1985.

Upon investigation in February 1987, the Yakima staff found Tilman to have a chronic skin ulcer, dead tissue, and constant pain in her hip, providing further evidence for a radiation overexposure. Tilman underwent surgery and skin grafts, and overcame the incident with minor disability and some scarring related to the overdose.

Isaac Dahl: March 22, 1986

At the East Texas Cancer Center (ETCC) in Tyler, Texas, 33-year old Isaac Dahl was to receive his ninth Therac-25 radiation therapy session after a tumor had been successfully removed from his left shoulder. By this time the Therac 25 had been in successful operation at Tyler for two years, and 500 patients had been treated with it.

The Therac-25 operator left the radiation room to begin the treatment as usual. As she was typing in values, she made a mistake and used the "cursor up" key to correct it. Once the values were correct, she hit the "B" key to begin treatment, but the Therac-25 machine shut down after a moment, and the message "Malfunction 54" showed on the control room monitor. The machine indicated that only 6 of the prescribed 202 units of radiation had been delivered. The screen of the console showed that this shut down was a "treatment pause" which indicated a problem of low priority (since little radiation had been delivered). The operator hit the "P" key to proceed with the therapy, but after a moment of activity, "Malfunction 54" appeared on the Therac control screen again.

The operator was isolated from Dahl because the Therac-25 operates from within a shielded room. On this day at the ETCC, the video monitor was unplugged and the audio monitor was broken, leaving no way for the operator to know what was happening inside. Isaac Dahl had been lying on the treatment

table, waiting for the usually uneventful radiation therapy, when he saw a bright flash of light, heard a frying, buzzing sound, and felt a thump and heat like an electric shock.

Dahl, knowing from his previous 8 sessions that this was not normal, began to get up from the treatment table when the second "attempt" at treatment occurred. This time the electric-like jolt hit him in the neck and shoulder. He rolled off the table and pounded on the treatment room door until the surprised Therac-25 operator opened it. Dahl was immediately examined by a physician, who observed reddening of the skin but suspected only an electric shock. Dahl was discharged and told to return if he suffered any further complications.

The hospital physicist was called in to examine the Therac-25, but no problems were found. The Therac-25 was shut down for testing the next day, and two CMC engineers, one from Texas and one from the home office in Canada, spent a day at the ETCC running tests on the machine but could not reproduce a Malfunction 54. The home office engineer explained that the Therac-25 was unable to overdose a patient and also said that CMC had no knowledge of any overexposure accidents by Therac-25 machines. No electrical problems were found with the ETCC's Therac machine, and it was put back into use on April 7, 1986.

Isaac Dahl's condition worsened as he lost the use of his left arm and had constant pain and periodic nausea and vomiting spells. He was later hospitalized for several major radiation-induced symptoms (including vocal cord paralysis, paralysis of his left arm and both legs, and a lesion on his left lung). Dahl died in August of 1986 due to complications from the radiation overdose.

Daniel McCarthy : April 11, 1986

Technicians could find nothing wrong with the Therac-25 unit at the East Texas Cancer Center (ETCC), after the "Malfunction 54" incident that had injured Isaac Dahl. The machine was reinstated.

Four days later, Daniel McCarthy was being treated for skin cancer on the side of his face. The same Therac operator who had treated Isaac Dahl was treating McCarthy. As the operator prepared to administer the Therac treatment from the control room, she used the "cursor up" key to correct an error in the treatment settings. She then began treatment using the "B" key.

The Therac-25 shut down within a few seconds, making a noise audible through the newly repaired intercom. The Therac monitor read "Malfunction 54." The operator rushed into the treatment room and found McCarthy moaning for help. He said that his face was on fire. The hospital physicist was called. McCarthy said that something had hit the side of his face, and that he had seen a flash of light and heard a sizzling sound.

After this second accident at the hospital, the ETCC physicist took the Therac-25 out of service and called CMC. He worked with the Therac operator who had been administering treatment to both Dahl and McCarthy when the accidents occurred. The physicist and the operator were eventually able to reproduce a Malfunction 54. They found that the malfunction occurred only if the Therac-25 operator rapidly corrected a mistake.

The ETCC physicist notified CMC of this discovery and CMC was eventually able to reproduce the error. CMC advised Therac-25 users to physically remove the up-arrow key as a short-term solution.

CMC also filed a report with the United States FDA as required by law, and began work on fixing the software bug.

The FDA worked in conjunction with CMC to identify the software problem and correct it. The FDA also requested that CMC change the machine in several other ways to clarify the meaning of malfunctions error messages and to shut down treatment after any single large radiation pulse or interrupted treatment so that multiple overdoses were less likely.

Over the next three weeks Daniel McCarthy became very disoriented and then fell into a coma. He had a fever as high as 104 degrees and had suffered neurological damage. He died on May 1, 1986.

Anders Engman: January 17, 1987

Anders Engman was at the Yakima Valley Memorial Hospital on January 17, 1987 to receive three sets of radiation treatment from the Therac-25.

The first two treatments went as planned. Engman received 7 rads (radiation absorbed dose), 4 rads followed by 3 rads of radiation to take pictures of internal structure. The Therac-25 operator then entered the room and used the Therac-25's hand control to verify proper beam alignment on Engman's body. Engman's final dose of the day was to be a moderate 79-rad photon treatment.

The operator pressed a button to command the Therac to move its turntable to the proper position for treatment. Outside the treatment room, the Therac-25's control console read "beam ready," and the operator pressed the "B" key to turn the beam on. The beam activated, but the Therac-25 shut down after about 5 seconds. The console indicated that no dose had been given, so the operator pressed "P" to proceed with the treatment.

The Therac-25 shut down again, listing "flatness" as the reason for treatment pause. Engman said something over the intercom, but the operator couldn't understand him. The operator went into the treatment room to speak with Engman. Engman told the operator that he had felt a "burning sensation" in the chest. The operator's console displayed only the total dose of the two earlier treatments (7 rads). Later that day, Engman developed a skin burn over the treatment area. Four days later the burn was striped in a manner similar to that of Janis Tilman's burn after she had been treated at Yakima the year before.

CMC investigated the accident. All users were again told to visually confirm turntable setting before proceeding with any treatment. Given the information, it was suspected that the electron beam had come on when the turntable was in the field light position. CMC could not reproduce the error.

Later that week, CMC sent an engineer to Yakima to investigate. The hospital physicist had also been running tests. They eventually discovered a software flaw and fixed it. CMC engineers estimated that Engman received between 8,000 and 10,000 rads instead of the prescribed 86.

Anders Engman died in April 1987. He had been suffering from a terminal form of cancer before the Therac accident, but it was determined that his death was primarily caused by complications related to the radiation overdose, not the cancer.

Teaching Section

Teaching Introduction

Therac-25 Introduction

This introduction to the Therac 25 case is for teachers of the case. If you have been assigned to read this case you can find the case material in the case section of this web site.

Here we provide a guide to the case from the inside or from the teacher's perspective. This section provides you with

7. supporting documents (excerpts from a published analysis of the case, an interviews with an operator, a memo from a medical physicist, and references)
8. An overview of the Socio-Technical system in which the case is embedded
9. Ethical analysis of the case (done using the ImpactCS model)
10. Specific assignments that you might use in class

Therac-25: Safety is a System Property

Normally, when a patient is scheduled to have radiation therapy for cancer, he or she is scheduled for several sessions over a few weeks and told to expect some minor skin discomfort from the treatment. The discomfort is described as being on the order of a mild sunburn over the treated area. In the case you are about to read, a very abnormal thing happened to several patients: they received severe radiation burns resulting in disability, and, in 3 cases, death.

The Therac-25 was a device that targeted electron or X-ray beams on cancerous tissue to destroy it. Electron beams were used to treat shallow tissue, while photon beams could penetrate with minimal damage to treat deep tissue. Even though operators were told that there were "so many safety mechanisms" that it was "virtually impossible" to overdose a patient, this is exactly what did occur in six documented cases [Leveson].

These massive radiation overdoses were the result of a convergence of many factors including

11. simple programming errors,
12. inadequate safety engineering,
13. poor human computer interaction design,
14. a lax culture of safety in the manufacturing organization,
15. inadequate reporting structure at the company level and as required by the U.S. government.

In presenting this case we are not interested in determining who should be blamed for these accidents. All the cases have already gone through the courts and have been settled. We are interested in helping you learn how to think about the design and use of software in safety-critical applications. What are the responsibilities of the organizations and individuals involved? What design decisions and organizational structures led to the accidents? How might different organizational systems or software design have helped avoid or minimize the harm?

As a computer scientist, you will be focussing on the software in this medical linear accelerator. And indeed there are some clear coding errors on which we can focus. However, the more difficult and dangerous problems are those in the design of the entire system, and in the way the software plays its part in that design. These system safety issues are critical to understanding this case and to understanding what it means to design safe software.

Structure of the Therac-25 Case

Our presentation of the case itself is composed of three parts: Introductory materials, a description of the machine, and overviews of the participants in the case. Together, these sections give one a good idea of the information each actor in the case had at the time of the accidents.

We reserve any analysis of this case for the teaching section. However, many of the section contain broad hints regarding the danger of the machine and the particular ways that inadequate software design might cause harm to patients.

Introductory materials

These provide some background for students to understand the case. There is a general introduction to the case, explanations of how radiation therapy works, and a section on how medical linear accelerators work.

The machine

This section provides an overview of how the Therac-25 machine itself worked. This includes a description of the turntable, the rooms in which the machine is placed, and the role of the operator in setting up the machine.

There is also a section on the design of the software. This is a high-level introduction to the issues involved in the design of the software. The excerpts from Leveson we provide in the resource section provide much more detail, down to two particular coding errors that probably caused some of the accidents.

Finally there is a section specifically on safety. The issues involved in removing the hardware interlocks are explained, as are the issues of sensing the position of the turntable and of reuse of software from older Therac machines.

The Participants

Each of four participants are presented here, along with the accounts of each accident. The perspectives of the designer/manufacturer of Therac, of the FDA, of the hospitals, and of the operators of the machines are all presented in some detail. This will allow you to assign individuals to cover the perspectives of each of these groups.

Using the Therac-25 Case in Class

To get acquainted with the case, we recommend you read as much of the case material as holds your attention. You might then turn to the analysis documents to see how we view the ethical issues in the case. Finally, you might at least look at the overview of the supporting documents we provide.

For a more practical turn, you might choose an exercise from our list of exercises. Each exercise will require the use of different supporting documents and of different pieces of the case presentation.

As you develop the exercise that you want to use in your class, you should think of how you might present it on a web page. We hope soon to provide some support to make it easy for you to construct a web page that presents your exercise, as you have modified it, to your students.

Pitfalls in Teaching with These Cases

Each of the cases we provide at this site 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 just an oddity of the cases we have selected; 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 you with the tools you will need to teach these cases in a classroom environment. One of those tools is some advice about pitfalls one is likely to find in any discussion of cases this complex.

The Search for a Single Cause

This pitfall is usually signaled by a sentence like "If only X hadn't happened, then Y would never have occurred!" There are sometimes single-point-of-failure items in a design or in a social system. And it may well be that if X hadn't happened then Y would not have happened either--AT LEAST BY WAY OF X. Finding one cause of an accident does not mean one has found all of them. There are at least two documented coding errors in the Therac-25 case, but no real assurance that these are the only ones. Email made it easy for Sanchez to verbally attack others, but other technologies would have worked well too. American Aircraft was cited for particular flaws in its inspection system, but the larger culture of safety was in need of repair.

One way to head off this tendency toward simplification is to maintain a system perspective. The levels of analysis in the ImpactCS framework help in this regard, focussing attention on various levels of scale in the system.

Subjectivity in Thinking About Causes

This pitfall is harder to appreciate in a case analysis unless one does some role playing with only certain kinds of information available to each individual in the role play. If you have read all of the Therac-25 case we present, you have a "god's eye view" of the action, like the reader of a classic novel. But the people inside the case never had the convenience of this view, and might not even be able to agree now on a single hindsight-driven account after the fact.

Once you can firmly grasp the limited nature of each individual's knowledge about what was going on at any particular time, you can then progress to further ambiguity about how people might be personally or professionally motivated to find one kind of cause for an accident more likely than another. It may be that the CMC technicians were simply handicapped by their limited knowledge

when they thought an electric shock might be the cause of the first accident in Tyler, TX. But it was also a convenient cause in that it was not CMC's problem.

A note of caution here. It is easy to become cynical and to invent conspiracies and evil motivation for everyone. Resist this temptation. Cynicism is the easy way to appear sophisticated about a case, but it is often misleading. Most people, most of the time, at least think they are not harming others by their actions. When you find the opportunity for a motivated distortion of the causes of an accident, note it, and ask yourself how that person could maintain the distortion while still thinking well of themselves.

The Search to Affix Blame

This pitfall is related to the (sometimes reasonable) search for malign motivation mentioned above. It is also a version of simplification. If we can just find a culprit to blame we feel as though we have found a satisfying explanation. Preferably, there is only one culprit, but if there are several, they should be in a conspiracy together.

This search for a psychologically satisfying explanation may be fine for coffeehouse conversation, but engineers responsible for the operation of the system should look beyond it for two reasons. First, if safety is a system property (as Safety Engineering and the ImpactCS model imply), then focussing on only one level of the system will mislead us about the complexity of the issues. Implied blaming Sanchez for the email harassment does not lead us to look for the more complex ways that the culture of information sharing made this harassment easier. Second, a focus on individual or corporate blame implies only one kind of solution: people should be more careful or they will be punished. We can blame American Aircraft for its mistakes in inspection. But this will lead us to overlook the production pressures produced by the system of military contracts. The better approach is to look at the system to determine what people should be careful about. What can be modified in the system to increase those values about which we care?

The Rush to Legalism

The Therac-25 case has already been adjudicated by courts. Awards have been made or people have settled out of court to avoid further litigation. We do not include this, or any other, case in our study so individuals can determine how to avoid being sued. We include it so we can understand how a system works to either provide safety or to make accidents more likely.

It is quite easy to make distinctions between what is legal and what is moral. Neither set is contained entirely by the other, though they do overlap. One can think of actions that are legal but not moral (e.g. racial discrimination in the US before the 1960s). In addition, it is easy to find examples of actions that are moral but not legal (any civil disobedience). But again, like the search for someone to blame, this false trail leads us to ignore both the complexity of the systems and many possible solutions that might lead to a safer system.

The Loose Ascription of "Human Error"

Certainly any system that involves humans will find itself having to cope with human error. But explaining the Therac-25 accidents as due to human error suggests there is nothing that could have been done to prevent it, other than telling people to "be more careful" (see affixing blame, above). The fields of Computer Human Interaction and of Human Factors are based on the idea of accommodating

the propensities of humans, by adapting both to their strengths and weaknesses. Thus, a finding of "human error" is useful only if it leads to a search for system change that might help to make that human error less likely in the future.

A Fixation on the Technical Failures (and Fixes) in the System

This pitfall is yet another version of simplistic thinking about the causes of accidents. Certainly technical failures in the system should be fixed. But a too narrow focus on these can lead one to think that simply fixing each one as it occurs is the appropriate action. For instance, one might think that the child safety problem with early refrigerators was that the latch could not be opened from the inside if a child was trapped inside. A narrow focus on this might lead one to add a "child detector" to the refrigerator that opens the latch when this troublesome condition occurs. This now leaves us with an additionally complicated system with more points of failure. Rethinking the design of the product allows one to see that one does not need latches to close refrigerator doors--magnets will do fine.

In a similar manner, we have several examples in the Therac-25 case of fixation on technical failures. The initial response of CMC involved pinpointing a microswitch failure as the problem. The claim that the technical fix to this produced a five order of magnitude increase in safety suggests that CMC felt this fix was the single solution. It took more than a year of negotiation with the FDA to get a plan from CMC that involved more systemic redesign issues.

One might ask the question: "Why have a dual mode medical linear accelerator in the first place?" At least one of the kinds of accidents from the Therac-25 would have been completely impossible if the machine used only a single mode. There are, of course, tradeoffs in this design decision. But negotiating design tradeoffs is standard fare for software engineers.

Analysis

Therac-25: A Socio-Technical System

The safety of the Therac-25 is not really a property of the machine alone. Accidents that go unreported contribute to (or at least fail to stop) later accidents. When the TV camera in the room is unplugged, the operator cannot see that the patient is in trouble. So safety is really a property of the entire technical and social system (socio-technical system). In a similar manner, an ethical analysis of the issues in this case requires an awareness of the entire socio-technical system.




Figure 1. Typical Therac-25 facility

The Therac-25 Medical Linear Accelerator is a large machine that sits in a room designed just for it. We think of the machine itself or the machine-in-the-room as the system. But the larger system, or the Socio-Technical system, that we need to think about includes:

16. *Hardware*: The mechanics of the machine itself, including its associated computer
17. *Software*: the operating system of the computer and the operating system of the machine
18. *Physical surroundings*: the room with its shielding, cameras, locking doors, etc.
19. *People*: operators, medical physicists, doctors, engineers, salespeople, managers at CMC, government regulators
20. *Institutions*: CMC, FDA, each medical facility, associations of operators, etc.
21. *Procedures*
 22. *Management models*: CMC's model of how risk is managed
 23. *Reporting relationships*: who was required to report accidents to whom
 24. *Documentation requirements*: for the software, for the facilities, for the FDA
 25. *Data flow*: how different parts of CMC shared information, how information was shared among agencies and organizations, how data was used by the Therac software.
 26. *Rules & norms*: what patients are "normally" told, what operator & physicist responsibilities are, expectations set for the programmer
 27. *Laws and regulations*: Reporting requirements, FDA enforcement mechanisms, medical liability law

28. *Data*: data was collected in FDA approval process, use of data in Therac software,

The following table presents some of these items in a schematic form.

The Socio-Technical System	
<p><i>The Machine</i></p> <p>6. Supporting Systems (video, audio, etc.)</p> <p>7. Hardware</p> <p>8. Software Systems</p>	<p>Hospitals and Clinics</p> <ul style="list-style-type: none"> • Doctors, Medical Physicists • Management, User Groups • Operators, Reporting Procedures
<p>Canadian Medical Corporation</p> <ul style="list-style-type: none"> • Management, Reporting Procedures, • Design Teams, Sales Staff, Support and Field Engineers 	<p>Government Medical Device Regulation</p> <ul style="list-style-type: none"> • Food and Drug Administration • Canadian Radiation Protection Bureau • Reporting Procedures

A thorough investigation of the Therac-25 case requires some grasp of most of these items. You will come across most of these items as you read this case. Setting your sights on the entire system will help you avoid the trap of finding a single point of blame. It is easy, for instance, to decide that the programmer made serious mistakes and to end one's analysis there. This is a short-sighted approach. It would miss the problems with maintenance in the cancer therapy facilities; it would miss the incomplete reporting requirements for the FDA; it would miss the inadequate and misleading testing of the Therac-25 system.

Ethical Analysis

Using the ImpactCS Grid to Understand Therac-25

The ImpactCS approach to ethical analysis was devised by a panel of ethicists, computer scientists, and social scientists. The point is that any particular computing system can be analyzed from both the perspectives of social analysis and of particular ethical issues. The grid you see below was designed by the panel to serve as an analytic tool in thinking about any system. The idea is that each of the ethical issues can be analyzed at each of the levels of social analysis.

For instance, in this case, safety is the primary concern. But we need to think about issues of safety at individual, group, national and global (or international) levels. Each of these levels brings forth different issues and different ethical concerns. In addition, the grid reminds us that safety issues are only one of many issues that might concern us in a case.

We unpack these concerns in a series of documents that you can access by clicking on the highlighted cells in the Framework below.

		Ethical Issues						
		Quality of Life	Use of Power	Safety	Property Rights	Privacy	Equity and Access	Honesty and Deception
Levels of Social Analysis	Individual							
	Group							
	National							
	Global							

Quality of Life

The decision to computerize a medical linear accelerator was, in the beginning, a quality of life consideration. CMC did not set out to make a device that would expose individual to harm. It made improvements in a device that would, in theory, allow greater access to a medical technology and would increase the quality of care those patients received. But in the process, a lax culture of safety in the organization led to a system design that was unsafe and not well tested. In this case, the values of safety and increased quality of life for consumers need not have been in conflict. But in practice, they became so.

Power

Perrow suggests that most risk analysis procedures are really a way for *some* people to think clearly about the risks to which they will subject *other* people. People who are doing a risk or safety analysis are usually those hired by the company to protect itself from risk. There are mixed motives here: by protecting themselves from risk, they also protect the safety of those in using their products.

Ford Motor Company made itself infamous by explicitly paring risk to the company (in dollars lost from lawsuits) to risk that consumer faced (from inadequate design of gas tanks in the Pinto). They decided that it would cost less to pay the lawsuits than to fix the car. Here the calculations were all financial. But it is at least up for debate whether all companies make decisions in this manner. In many, the motives are mixed: protection of the company and safety of the consumer.

But CMC's priority seems odd even in the light of self-protection. Its risk analysis seemingly was not done to protect the company, but to certify their already strongly held belief that the machine was safe. This sort of unfounded optimism regarding technology at least provides them with the defense of ignorance. But this defense is less persuasive when offered by those with power over other's well-being. Often, when individuals or corporation are given more power, we are also more likely to hold them more responsible for their actions.

At any rate, this case is clearly an issue of who has power to enforce the acceptance of risks on others. This power may be economic (as in the case of CMC) or political (as in the FDA).

But at the individual level, the power may simply be positional--acquired because you happen to be the software engineer assigned to a particular project. This is what Huff has called unintentional power--the power that a designer has over the users of a product. Someone with unintentional power uses it without intending benefit or harm to the ultimate user of the product. This is another case of the

defense of ignorance. In this case, the defense is harder to believe, since the software controlled a potentially lethal radiation beam. But the intention of the programmer, or of the operator, was not to harm patients, or even to place themselves in the position where they could harm them. But taking a job as a programmer entails this unintentional, positional power. It is better to know this than to ignore it.

These sorts of power differentials exist at all levels of the social analytical framework. And a careful ethical analysis of power will ask what duties go along with that power, and what rights are held by those with less power.

System Safety

We cannot present here a full analysis of system safety and instead refer the interested reader to items in our bibliography and to various web sites that address these issues. You can see what noted system safety expert Nancy Leveson has to say about this case in the excerpt from her article that we provide in the supporting documentation.

However, we have at least learned that in order to understand the safety issues properly, we must look at them at several levels of social complexity, just as the ImpactCS framework suggests.

Safety at the Individual Level

The programmer

Certainly the single individual who did the programming for Therac 25 had responsibilities as a computing professional. To whom were these responsibilities owed? An obvious first responsibility is to the organization that employed him. Another party to whom responsibility is owed is to the eventual users of the linear accelerator: the patients. We can certainly add to these two (e.g. to the profession, to the machine operators), but let's take each of these for the purposes of this analysis.

The programmer's responsibilities to his employer were more than simply to do as directed to by the other designers of the system (or by whoever his immediate superiors were). He had a responsibility to make his superiors aware of the dangers inherent in doing safety interlocks only in the software. Whether this danger was obvious to him or not is an interesting question. Even today many computing professionals place more confidence in the safety of software than is likely. Software safety was little understood at the time the Therac-25 system was designed.

The point here is that a computer professional is responsible to the employer for using the best available methods to solve the software problems with which he or she is confronted. There are a variety of professional decisions the programmer made in the Therac-25 design that suggest he was lax in this responsibility (e.g., using unprotected memory, improper initialization, lack of appropriate testing, etc.). Thus, as an employee, he fell short of the mark in providing his employer with professional work. We cannot know whether this shortcoming was one of lack of knowledge or of poor execution.

In addition to responsibilities to his employer, the programmer clearly had a responsibility to the users of the technology he was designing. In the context of safety, his responsibility was to design software

that minimized the likelihood of harm from a dangerous medical device. This obligation to “do no harm” need not mean that software should never be paired with medical linear accelerators. From the perspective of the operator we interviewed this pairing was a positive benefit in making setup and treatment easier. But it does mean that, to the extent it was within the professional control of the programmer, he should have designed the system to do no harm while providing this positive good. Again, whether the failure to do this was a result of a lack of knowledge or of poor execution, we cannot know.

To sum up, the programmer had clear responsibilities to both his employer and to the users of the device. He clearly failed in these responsibilities. If we were interested in blame, we could not tell the amount of blame to assign here. We know nothing of the programmer’s background or training. Thus we cannot know if the programmer knew how poor the software design and testing was. Nor do we have any idea of the internal company dynamics that may have resulted in the lack of testing.

The operators

We noted in the case write-up that the operators of the linear accelerators had a complex combination of responsibilities. Chief among these are responsibilities to their employer and to the patients.

Just like the programmer of the system, we know little about the background and training of the operators in this case. But we can at least specify their responsibilities. They were responsible to their employers to operate the machine efficiently, getting all the scheduled treatment done in any particular day. They also had a responsibility to their employer to look after the machine so it could be maintained properly. Finally they had a responsibility to their employer to operate the machine carefully and not to place patient in danger.

From published accounts of operator’s comments, from our interview of a Therac-4 operator, and from comment’s made in court documents, it seems clear that none of the operators felt they were placing their patients in any danger when they pressed the button. Thus, we can rule out intentional negligence. But what happened? Leveson suggests that the interface on the console made operators tolerant of error messages, and readily rewarded them for pressing the “proceed” button whenever a minor error appeared. The interface made no distinction between life threatening errors and minor errors, except that major errors would not allow a “proceed.” Given this, it is hard to see how the operators might be responsible for the errors, even though they were the ones to press the key.

An interesting issue arises because of the current move among operators to become more professionalized. As operators are better trained, are certified, and are more aware of the workings of the machine, they gain the prestige – but they also gain responsibility. As they become well trained enough to foresee such errors, their responsibility for them will increase.

Safety at the Group Level

There are two organizations at this level whose actions need to be thought about: the treatment facilities and the Canadian Medical Corporation.

The Canadian Medical Corporation

With regard to safety in this case, CMC’s responsibility in making a medical linear accelerator are to a range of individuals: their shareholders, their employees, the governments of Canada and the United

States, to the facilities that bought the machine, and finally to the patients who were treated by them. Responsibility to shareholders and employees are similar, and for this analysis will be considered the same.

Before we look at these specific responsibilities, we will need to understand some of the technical issues involved in the analysis of a system for safety. In this instance, technical knowledge is required to make ethical judgments.

CMC claimed to do a safety analysis of its machine, but in fact the analysis only shows the likelihood of the system failing because a part wears out. There was apparently no systematic search for design flaws in the software until after the FDA required an analysis. Unfortunately, a system can be highly reliable but thereby reliably kill people because of a design flaw. This confusion of reliability analysis and safety analysis is a critical failing on the part of CMC.

Some indication of the motivations behind CMC's inadequate safety analyses can be gleaned from the way CMC appeared to use probabilities in its analysis. These probabilities seemed to be assigned to quantify and to prove the safety of the system, rather than to identify design flaws. For example, after redesigning the logic to track the microswitches that indicated the position of the turntable, CMC apparently used a sort of Fault Tree Analysis to assert that the safety of the system had been improved by at least 5 orders of magnitude. This astonishing claim of improvement is applied to the safety of the entire machine. This use of probabilities from a Fault Tree Analysis can effectively hide critical design flaws by inflating the perception of reliability and discouraging additional search for design flaws. This hiding of design flaws was a tragic, if unintentional side effect of the improper use of this analysis.

Thus, in failing to look systematically for design flaws in its software, CMC left itself (and its employees and shareholders) open to liability claims from injured consumers. This is clearly also a failure of its responsibility to patients and to the facilities who bought the Therac-25 machine and who were assured there was no way it could hurt patients. This failure must be perceived in the light of prevailing standards (or lack thereof) in system safety at the time of the design and release of Therac-25.

The Cancer Treatment Facilities

The cancer treatment centers are primarily consumers of a product that is tested, maintained, and certified by others. This product was sold to them with assurances that it could not hurt patients. And the facilities do not have the responsibility or the capability to independently check these systems for safety.

But they do have responsibility for the safe operation and low level maintenance of the machines once they are in operation. It was clear that at least one facility fell down in this respect. In the first Tyler accident, the video monitor to the room was unplugged and the intercom was out of order. This would not have been a problem if there were no accidents – but there were. One difficulty with the safe operations of systems is that standard maintenance can become tedious and not seem a necessary component in the safe operation of a regularly used system. In this case, an individual might have been spared a second overdose if the basic communication systems had been working.

We should note, however, the extraordinary efforts of the medical physicist at Tyler in determining the cause of the overdose. This individual effort was supported by the Tyler facility and made possible by the facility's decision to have a full time physicist on staff. Some evidence of this support comes from the facility's decision to report the accident to the FDA even though there was no requirement that they do so. Note that the facility decided that their responsibility extended beyond the requirements of the law.

Thus, most facilities had relatively minimal responsibilities in this case and most seemed to fulfill them. The facilities had little power to resolve the problem and depended on CMC and on the FDA's approval process to protect them and their patients. Perhaps in this dependence they were too optimistic, but it is difficult to see what other choices they might have had.

Safety at the National Level

At the time of the Therac-25 accidents, the *Center for Devices and Radiological Health* (CDRH) of the FDA was responsible for the oversight of the immense market in radiation based therapy and diagnostics. As we have seen, most (94% in 1984) devices for the market were approved by "pre-market equivalence" and thus not subjected to stringent testing. The CDRH could not have handled the load of testing all these devices.

Since the rules for FDA are set by congress, FDA's rules need to be analyzed from the perspective of the responsibilities of congress. But FDA implementation of those rules is under its control. Thus we can ask if the CDRH (as a center in FDA) should have allowed Therac-25 to be approved under pre-market equivalence. Without more information this is difficult to determine. The CDRH did seem to vigorously follow the case once it became aware of the Tyler accidents, though there is some evidence that they were reluctant to quickly halt the used of the Therac-25 when the problems became evident. This reluctance may be because of their responsibility to not place an undue burden on manufacturers in their caution regarding a product. This tension between responsibilities to manufacturers/industry and responsibilities to patients is always present in decisions by the FDA. Hindsight makes this one seems easy to decide.

One of the problems in this case is that the FDA depended on CMC to notify it of accidents that had occurred. They did not hear directly from the hospitals when the accidents happened. CMC had, at best, a mixed record of notifying the FDA of problems. So perhaps facilities should have been required to directly report accidents (they are today). But FDA could not make this requirement; it could only enforce existing law. Thus, perhaps it was a responsibility of congress to enact this law. The counter-argument is that congress should allow the market to work out these issues. But in this instance, at least, the market was too slow to save the individuals who were killed or injured.

The entanglement of different ethical issues become very clear at this level of analysis. Different constituencies will value different things (e.g. personal privacy vs business freedom). These choices among different values are as severe at the other levels (e.g. operator's responsibility to employer and patient) but not as easily seen to the outside observer. Choice and balance among these values becomes inescapable, however, at the political level.

Safety at the Global Level

Communication between the Canadian Radiation Protection Board and the FDA seemed to work pretty well in this case. These two agencies had responsibilities to their respective governments, to industry in their countries, and to patients in their countries. CMC's communication with FDA did not seem to be hampered by its international flavor. However, this is a case of two relatively similar countries and cultures interacting with each other. Similarities in legal standards and in government oversight made this case easier. This might be an even less happy story if we had been dealing with widely different cultures of business or legal systems in the two countries.

Property

At one point in its interactions with the user groups, CMC found itself being asked for the source code used in the Therac 25 software. CMC claimed that it had proprietary rights to the software and would not make it public. This is another case of two values coming into conflict. A concern for safety suggests that it would be helpful to open the source code to inspection by the FDA or its agents or by the user groups. But to force this openness would violate the property rights of the owner of the software, CMC. One suspects that CMC's refusal to open the code to inspection is a defensive move based on avoiding liability rather than an attempt to protect the value of the intellectual property. But if close inspection showed the software to be poorly designed, the value of the software would surely diminish. Is this a case in which we want to uphold property rights? There may in fact be some case here for an "open software" standard to protect public safety.

The ImpactCS grid suggests we identify several levels of social analysis for each ethical issue. In this case we have interaction among those levels. The public (ideally, represented by the FDA) was placed at great risk by the software. What claims are there that the cancer treatment facilities and the FDA can make regarding the value of their being able to inspect the design and logic of the software? What claims can patients (or their surviving families) make regarding the validity of the claim by CMC to keep its software a trade secret? So, in order to understand the property issue correctly, we need to look at claims made at many levels: national, organizational, and individual.

Privacy

Privacy issues may well be raised by this case as one begins to recommend some sort of a national reporting mechanism for medical devices. In order to accurately report on any accident, sensitive medical data about individuals would need to be collected by treatment facilities and made available to national agencies. These national agencies might, in turn, make this data available internationally. In our case data about the accidents was shared by agencies of the Canadian and US governments.

It seems possible to make the data on patient records related to medical accidents anonymous – to separate the data from the identity of the patient. But one would have to think carefully about how to do this. The quickest solution, attaching the patient's medical history, is likely collecting too much data and violating the privacy of the patient. Collecting only as much data as is needed is a reasonable requirement.

Equity and Access

Equity and access issues are also raised by this case. The whole point of the design of Therac-25 was to make a medical linear accelerator that was less expensive to produce and thus likely less expensive

(and more available) to consumers. This is often an effect of the free market on the price of technology. If CMC could make a less expensive, but equally useful linear accelerator, it would sell more of them and they would be more easily available to the public. CMC would, doubtless, make money in the process. This is Adam Smith's invisible hand at work: decisions to make a better, less expensive product are good both for the manufacturer and for the consumer.

On the other hand, increasing regulation and oversight will impose increasing costs on providers of medical devices. These regulations may be seen as necessary, given the track record of companies like CMC. But they still increase the development costs to the company and the cost of the product. This, in turn, reduces the availability of the devices.

Again, we have here an issue of balance between competing goods: safety of the consumer and the increased access of the consumer to life-saving medical technology.

Honesty and Deception

Honesty and Deception issues are central to this case. In several places, CMC representatives made claims of safety for the Therac-25 device that in retrospect seem at least exaggerated.

How did this occur? These claims were made by individuals (salespersons, engineers), on the behalf of an organization (CMC). Individuals making claims like these have some responsibility to check on their accuracy. But salespeople have little expertise to evaluate this information and are thus more dependent on the organization. Engineers, including software engineers, have the capability of evaluating the claims though they may be allowed little time in which to do so. Again, we find a balance between the engineer's responsibilities to the company (use time efficiently) and to the consumer (evaluate carefully claims made about the product). Because of their special expertise, it is precisely the role of a professional to balance these conflicting responsibilities, and not to neglect responsibility to the consumer.

What organizational responsibilities might there be regarding claims of safety in medical devices? CMC representatives in several instances made claims that no overdoses had occurred with the Therac-25 machine, when there was clear evidence that someone at CMC must have heard of several previous accidents. This suggests that there may have been some internal miscommunication within CMC. Some portions of the organization may have known about the lawsuit regarding radiation harm but not have had the time or seen the need to inform other parts of the organization. For instance, those in the legal division, hearing of the lawsuit, may have assumed that the engineers were aware of the issue and that there was no immediate need to contact them. This sort of miscommunication is a daily matter in organizations, even small one (think of the miscommunication that occurs in your family).

Thus one part of the organization may have been making claims that no accidents similar to the reported ones had occurred based on the information available to them. Still, when the stakes are as high as they were in this case, organizations have a special responsibility to transmit safety critical information as quickly and as accurately as possible. This occurred sometimes within CMC (the FDA was notified quickly of the Hamilton accident) but not all the time.

Exercises for Therac-25

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 CMC indulged in during its initial reactions to the early accidents. We recommend you help your students avoid it as they approach this 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: CMC, 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.

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

Judge the Alternative

- Judge the alternative's effect on each stakeholders' opportunities and vulnerabilities and on each stakeholders' duties and rights.
- **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?

- **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 range of human-computer control possibilities on (448) to locate Therac-25. Recommend and argue for a change in level. What would be required to move 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 set 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.

29. 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?
30. 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.
31. Are the temporary fixes recommended by CMC adequate to remove the hazard?
32. 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 CMC home office and to other sites and then getting resolutions of the problems communicated back to the sites. In some cases CMC 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 (CMC, 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 (CMC, 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.

Supporting Documentation

Guide to the Supporting Documents

We have provided a set of supporting documents that should help the teacher of this case provide depth and analysis for much of the case. Please be aware that some of these items are copyrighted (e.g. the Leveson excerpts). Other items have been changed slightly from their originals to protect the identity of some individuals.

Case History

This is an overview of the case from the design of Therac-25 to the eventual shut-down of the machine pending redesign. It serves as a short history and guide to the case to give you your bearings.

The case narrative materials provide only information up until the time of the accidents. This nicely puts students in the decision makers seat, but one is left wondering what decisions actually were made by the main actors. This document provides answers to those questions.

Aliases

We use aliases for many names in this case. We do so because we do not want to focus on who the players are in this case, but rather on the predicaments in which they found themselves. However, if you must know who is who, we provide a key.

Leveson Excerpts

These excerpts are from the article Leveson published in 1993 in *IEEE Computer*. You can find a more current version of the article at her web site at <http://sunnyday.mit.edu/papers/therac.pdf>. We selected these excerpts because we felt they described the most critical issues in the case. They go into

much more detail on the software problems, the design of the machine and software, and the interface on the VT100 terminal.

Therac History: An overview of the history and physical design of the Therac-25.

The TurnTable. A close look at how the turntable was constructed and how its position was monitored. The turntable position was a critical issue in all the overdoses.

Software Design. An overview of the design of the software in Therac-25. Particular attention is given to how real-time issues resulted in race conditions.

Safety Analysis. An overview of the several different safety analyses that were done on the Therac-25 system.

Interface. A description of the operator interface on the VT100 operator console. Particular attention is given to the difficulties with error messages and with editing.

Tyler Software Problem. A description of the software problem that resulted in overdoses at Tyler, TX.

Yakima Software Problem. A description of the software problem that resulted in overdoses at Yakima, and possibly Hamilton, Ontario.

Produce Malfunction 54

This is a transcription of the memo that the medical physicist at the Tyler, Texas produced upon discovering how to produce the “malfunction 54.” Malfunction 54, produced in this way would deliver a dose 25,000 rads of 25 MeV electrons in less than two seconds. The standard therapeutic dose is about 200 rads at any one time. A dose of 500 rads over the entire body is considered lethal to 50% of individuals who receive it. Two persons were killed from the malfunction 54 overdose. One died in 5 months, the other within one month.

Operator Interview

There are two documents in this section. Both are derived from an interview we conducted with an operator of a Therac machine. This person was trained as a linear accelerator operator just before the transition to computer controlled linear accelerators in the mid-1980s. One of the first machines she worked on was a Therac-4. This machine had similar computer controls to the Therac-25, but it was not dual mode, and so all the accidents based on dual mode operation were not a possibility. In her interview she speaks of the training (or lack thereof) that operators receive, of the financial pressures on hospitals and cancer treatment facilities, and of the production pressures that operators experience.

First, there is the verbatim interview. This is somewhat closely transcribed, and so shows some of the standard awkwardness of spoken language transcribed into written language.

In addition, there is a summary of the interview. This provides most of the kernel without the need to wade through the chaff of the entire interview.

References

In addition to an attempt at a comprehensive bibliography, we provide annotations to those references that we think would be the most useful to students of this case.

A History of the Introduction and Shut Down of Therac-25

Therac-25 was released on the market in 1983. In 1987, its treatments with the eleven machines in operation was suspended. Those machines were refitted with the safety devices required by the FDA and remained in service. No more accidents were reported from these machines. At about that time, the division of CMC that designed and manufactured Therac-25 became an independent company.

The major innovations of Therac-25 were the double pass accelerator (allowing a more powerful accelerator to be fitted into a small space, at less cost) and the move to more complete computer control. The move to computer control allowed operators to set up the machine more quickly, giving them more time to speak with patients and making it possible to treat more patients in a day. Along with the move to computer control, most of the safety checks for the operation of the machine were moved to software and hardware safety interlocks removed.

CMC's FDA Testing and Safety Analysis

Before release of Therac-25 on the US market, CMC obtained approval to market it from the FDA. This approval was obtained by declaring what FDA called pre-market equivalence. Since the software was based on software already in use, and the linear accelerator was a minor modification of existing technology, designation of Therac-25 as equivalent to this earlier technology meant that Therac-25 bypassed the rigorous FDA testing procedures. In 1984, 94% of medical devices entered the market in this manner. This declaration of pre-market equivalence seems optimistic in that most of the safety mechanisms were moved into the software, a major change from previous versions of the machine.

In 1983, just after CMC made the Therac-25 commercially available, CMC performed a safety analysis of the machine using Fault Tree Analysis. This involves calculating the probabilities of the occurrence of varying hazards (e.g. an overdose) by specifying which causes of the hazard must jointly occur in order to produce the hazard.

In order for this analysis to work as a Safety Analysis, one must first specify the hazards (not always easy), and then be able to specify the all possible causal sequences in the system that could produce them. It is certainly a useful exercise, since it allows easy identification of single-point-of-failure items and the identification of items whose failure can produce the hazard in multiple ways. Concentrating on items like these is a good way to begin reducing the probabilities of a hazard occurring.

In addition, if one knows the specific probabilities of all the contributing events, one can produce a reasonable estimate of the probability of the hazard occurring. This quantitative use of Fault Tree Analysis is fraught with difficulties and temptations, as CMC's approach shows.

In order to be useful, a Fault Tree Analysis needs to specify all the likely events that could contribute to producing a hazard. Unfortunately, CMC's analysis left out consideration of the software in the system almost entirely. Since much of the software had been taken from the Therac-6 and Therac-20 systems, and since these software systems had been running many years without detectable errors, the analysts assumed there were no design problems in the software. The analysts considered software failures like "computer selects wrong mode" but assigned them probabilities like 4×10^{-9} .

These sorts of probabilities are likely assigned based on the remote possibility of random errors produced by things like electromagnetic noise. They do not at all take into account the possibility of

design flaws in the software. This shows a major difficulty with Fault Tree Analysis as it is often practiced. If the only items considered are "failure" items (e.g. wear, fatigue, etc.) a Fault Tree Analysis really only gives one a reliability for the system.

CMC's Response to the Accidents

In July of 1985, CMC was notified that a patient in Hamilton had been overdosed. CMC sent a service engineer to the site to investigate. CMC also informed the United States Food and Drug Administration (FDA), and the Canadian Radiation Protection Board (CRPB) of the problem. In addition they notified all users of the problem and issued instructions that operators should visually confirm hardware settings before each treatment. CMC could not reproduce the malfunction, but its engineers suspected that a hardware failure in a microswitch was at fault. They redesigned the hardware and claimed that this redesign improved the safety of the machine by five orders of magnitude. After modifications were made in the installed machines, CMC notified sites that they did not need to manually check the hardware settings anymore.

In November of 1985, CMC heard of another incident in Georgia. The patient in that incident (Linda Knight) filed suit that month based on an overdose that occurred in June. There is no evidence that CMC followed up this case with the Georgia hospital. Though this information was clearly received by CMC, there is no evidence that this information, was communicated internally to engineers or others who responded to later accidents.

In January of 1986, CMC heard from a hospital in Yakima, Washington that a patient had been overdosed. The CMC technical support supervisor spoke with the Yakima hospital staff on the phone, and contacted them by letter indicating that he did not think the damage they reported was caused by the Therac-25 machine. He also notified them that there have "apparently been no other instances of similar damage to this or other patients."

In March of 1986, CMC was notified that the Therac-25 unit in Tyler, Texas had overdosed a patient. They sent both a local Texas engineer and an engineer from their Canada home office to investigate the incident the day after it occurred. They spent a day running tests on the machine but could not reproduce the specific error. The CMC engineer suggested that perhaps an electrical problem had caused the accident. He also said that CMC knew of no accidents involving radiation overexposure with the Therac-25. An independent engineering firm checked out the electric shock theory and found that the machine did not seem capable of delivering an electric shock to a patient.

On April 11th of 1986, CMC was alerted to another overdose that had occurred in Tyler. After communication with the medical physicist at Tyler, CMC engineers were able to reproduce the overdose and the sequences leading up to it.

CMC filed a medical device report with the FDA on April 15, 1986 to notify them of the circumstances that produced the two Tyler accidents.

At this point, the FDA, having been notified of the first Tyler accident by the hospital, declared Therac-25 defective and ordered the firm to contact all sites that used the machine, investigate the problem, and submit a report called a corrective action plan. CMC contacted all sites and

recommended a temporary fix involving removing some keys from the keyboard at the computer console.

The FDA was not satisfied with the notification that CMC gave sites, and in May 1986 required CMC to re-notify all sites with more specific information about the defect in the product and the hazards associated with it. CMC was also at this time involved in meetings with a "user's group" of Therac-25 sites to help formulate its corrective action plan. After several exchanges of information among CMC and the FDA (in July, September, October, November, and December of 1986), CMC submitted a revised corrective action plan to FDA.

In January 1987, CMC was notified of another overdose occurring again at the Yakima, Washington hospital. After sending an engineer to investigate this incident, CMC concluded that there was a different software problem that allowed the electron beam to be turned on without the device that spread it to a safe concentration being placed in the beam.

Therac-25 is Shut Down

In February, 1987, the FDA and its Canadian counterpart cooperated to require all units of Therac-25 to be shut down until effective and permanent modifications were made. After another 6 months of negotiation with the FDA, CMC received approval for its final corrective action plan. This plan included numerous software fixes, the installation of independent, mechanical safety interlocks, and a variety of other safety related changes.

Several of the surviving victims or the deceased victim's families filed suit in US courts against CMC and the medical facilities using Therac-25. All of these suits were settled out of court.

CMC Medical Goes Independent

The division of CMC that designed and manufactured Therac-25 has become an independent private Canadian company. They still make radiation therapy machines.

Government and FDA response to the Accidents

The Therac-25 case pointed to significant weak links in communication between FDA, medical device manufacturers, and their customers or users. Users were not required to report injuries to any government office, or to the manufacturers of the devices that had caused injury.

A 1986 GAO study found 99% of injuries caused by medical devices were not reported to the FDA. At that time, hospitals reported only about 51% of problems to the manufacturer. The hospitals mostly reported dealing with problems themselves. Problems were mainly the result of wear and tear on machines and design flaws.

The breakdown in communication with hospitals and clinics using medical devices prevented FDA from knowing about the isolated and recurring problems with the Therac-25 until after two deaths occurred in Tyler, TX.

Even when the FDA became aware of the problem, they did not have the power to recall Therac-25, only to recommend a recall. After the Therac-25 deaths occurred, the FDA issued an article in the *Radiological Health Bulletin* (Dec. 1986) explaining the mechanical failures of Therac-25 and

explaining that "FDA had now declared the Therac-25 defective, and must approve the company's corrective action program."

After another Therac-25 overdose occurred in Washington state, the FDA took stronger action by "recommending that routine use of the system on patients be discontinued until a corrective plan had been approved and implemented" (Radiological Health Bulletin, March 1987). CMC was expected to notify Therac-25 users of the problem, and of FDA's recommendations.

After the Therac-25 deaths, the FDA made a number of adjustments to its policies in an attempt to address the breakdowns in communication and product approval. In 1990, health-care facilities were required by law to report incidents to both the manufacturer and FDA.

Aliases in Therac-25

This case is a report of a real occurrence that we have documented to the best of our ability. None of the individuals in the case were fictional. None of the individuals were composites. We have attempted to verify the truth of all the events we have reported.

However, we have used aliases in this case because the purpose of our presentation is not to blame individuals, but to see the ethical and social issues that surround the design, manufacture, and use of computing systems in the real world.

We have not attempted to hide the identity of some organizations that were widely known (e.g. the Food and Drug Administration). Therac-25 is the real name of the machine involved in this case, and we felt it useless to change that name for presentation of the case. But we have attempted to provide a modicum of anonymity to others involved.

For those who require the real names involved in this case, we offer the following key

Real Name	Alias
Katy Yarbrough	Linda Knight (Injured June 3, 1985)
Frances Hill	Donna Gartner (Injured July 26, 1985)
Not disclosed	Janis Tilman (Injured Dec. 1985)
Voyne Ray Cox	Isaac Dahl (Injured March 22, 1986)
Verdon Kidd	Daniel McCarthy (Injured April 11, 1986)
Glen Dodd	Anders Engman (Injured Jan. 17, 1987)
<i>Atomic Energy Canada, Limited (AECL). Now a vendor of Nuclear power reactors</i>	Canadian Medical Corporation (CMC)
<i>Theratronics (private company resulting from spin-off of AECL Medical, then purchased by MDS Nordion)</i>	

Leveson Excerpts

The material on this page is reprinted from N.G. Leveson, & C.S. Turner. "An Investigation of the Therac-25 Accidents." *Computer*, Vol. 26, No. 7, July 1993, pp. 18-41. Copyright © 1993 Institute of Electrical and Electronics Engineers. This material is posted here with permission of IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of St. Olaf College's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by sending a blank email message to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

Genesis of the Therac-25

Medical linear accelerators (linacs) accelerate electrons to create high-energy beams that can destroy tumors with minimal impact on the surrounding healthy tissue. Relatively shallow tissue is treated with the accelerated electrons; to reach deeper tissue, the electron beam is converted into X-ray photons.

In the early 1970's, Canadian Medical Corporation (CMC) and a French company called CGR collaborated to build linear accelerators. (CMC is an arms-length entity, called a crown corporation, of the Canadian Government.) Since the time of the incidents related in this article, CMC Medical, a division of CMC is in the process of being privatized and is now called Therapeutic Accelerators Limited. Currently CMC's primary business is the design and installation of nuclear reactors.) The products of CMC and CGR's cooperation were (1) the Therac-6, a 6 million electron volt (MeV) accelerator capable of producing X rays only and, later, (2) the Therac-20, a 20 Me V dual-mode (X rays or electrons) accelerator. Both were versions of older CGR machines, the Neptune and Sagittaire, respectively, which were augmented with computer control using a DEC PDP 11 minicomputer.

Software functionality was limited in both machines: The computer merely added convenience to the existing hardware, which was capable of standing alone. Industry-standard hardware safety features and interlocks in the underlying machines were retained. We know that some old Therac-6 software routines were used in the Therac-20 and that CGR developed the initial software.

The business relationship between CMC and CGR faltered after the Therac-20 effort. Citing competitive pressures, the two companies did not renew their cooperative agreement when scheduled in 1981. In the mid-1970's, CMC developed a radical new "double-pass" concept for electron acceleration. A double-pass accelerator needs much less space to develop comparable energy levels because it folds the long physical mechanism required to accelerate the electrons, and it is more economic to produce (since it uses a magnetron rather than a klystron as the energy source).

Using this double-pass concept, CMC designed the Therac-25, a dual-mode linear accelerator that can deliver either photons at 25 Me V or electrons at various energy levels (see Figure 1). Compared with the Therac-20, the Therac-25 is notably more compact, more versatile, and arguably easier to use. The higher energy takes advantage of the phenomenon of "depth dose": As the energy increases, the depth in the body at which maximum dose buildup occurs also increases, sparing the tissue above the target area. Economic advantages also come into play for the customer, since only one machine is required for both treatment modalities (electrons and photons).




Figure 1. Typical Therac-25 facility

Several features of the Therac-25 are important in understanding the accidents. First, like the Therac-6 and the Therac-20, the Therac-25 is controlled by a PDP 11. However, CMC designed the Therac-25 to take advantage of computer control from the outset; CMC did not build on a stand-alone machine. The Therac-6 and Therac-20 had been designed around machines that already had histories of clinical use without computer control.

In addition, the Therac-25 software has more responsibility for maintaining safety than the software in the previous machines. The Therac-20 has independent protective circuits for monitoring electron-beam scanning, plus mechanical interlocks for policing the machine and ensuring safe operation. The Therac-25 relies more on software for these functions. CMC took advantage of the computer's abilities to control and monitor the hardware and decided not to duplicate all the existing hardware safety mechanisms and interlocks. This approach is becoming more common as companies decide that hardware interlocks and backups are not worth the expense, or they put more faith (perhaps misplaced) on software than on hardware reliability.

Finally some software for the machines was interrelated or reused. In a letter to a Therac-25 user, the CMC quality assurance manager said, "The same Therac-6 package was used by the CMC software

people when they started the Therac-25 software. The Therac-20 and Therac-25 software programs were done independently, starting from a common base." Reuse of Therac-6 design features or modules may explain some of the problematic aspects of the Therac-25 software development and design. The quality assurance manager was apparently unaware that some Therac-20 routines were also used in the Therac-25; this was discovered after a bug related to one of the Therac-25 accidents was found in the Therac-20 software.

CMC produced the first hardwired prototype of the Therac-25 in 1976, and the completely computerized commercial version was available in late 1982. In March 1983, CMC performed a safety analysis on the Therac-25. This analysis was in the form of a fault tree and apparently excluded the software. According to the final report, the analysis made several assumptions:

33. Programming errors have been reduced by extensive testing on a hardware simulator and under field conditions on teletherapy units. Any residual software errors are not included in the analysis.
34. Program software does not degrade due to wear, fatigue, or reproduction process.
35. Computer execution errors are caused by faulty hardware components and by "soft" (random) errors induced by alpha particles and electromagnetic noise.

The fault tree resulting from this analysis does appear to include computer failure, although apparently judging from these assumptions, it considers only hardware failures. For example, in one OR gate leading to the event of getting the wrong energy, a box contains "Computer selects wrong energy" and a probability of 10^{-11} is assigned to this event. For "Computer selects wrong mode," a probability of 4×10^{-9} is given. The report provides no justification of either number.

Turntable Positioning

The Therac-25 turntable design is important in understanding the accidents. The upper turntable (see Figure B) is a rotating table, as the name implies. The turntable rotates accessory equipment into the beam path to produce two therapeutic modes: electron mode and photon mode. A third position (called the field-light position) involves no beam at all; it facilitates correct positioning of the patient.




Figure B. Upper turntable assembly

Proper operation of the Therac-25 is heavily dependent on the turntable position; the accessories appropriate to each mode are physically attached to the turntable. The turntable position is monitored by three microswitches corresponding to the three cardinal turntable positions: electron beam, X ray, and field light. These microswitches are attached to the turntable and are engaged by hardware stops at the appropriate positions. The position of the turntable, sent to the computer as a 3-bit binary signal, is based on which of the three microswitches are depressed by the hardware stops.

The raw, highly concentrated accelerator beam is dangerous to living tissue. In electron therapy, the computer controls the beam energy (from 5 to 25 MeV) and current while scanning magnets spread the beam to a safe, therapeutic concentration. These scanning magnets are mounted on the turntable and moved into proper position by the computer. Similarly, an ion chamber to measure electrons is mounted on the turntable and also moved into position by the computer. In addition, operator-mounted electron trimmers can be used to shape the beam if necessary.

For X-ray therapy, only one energy level is available: 25 MeV. Much greater electron-beam current is required for photon mode (some 100 times greater than that for electron therapy)[Rawlinson] to produce comparable output. Such a high dose-rate capability is required because a “beam flattener” is used to produce a uniform treatment field. This flattener, which resembles an inverted ice-cream cone, is a very efficient attenuator. To get a reasonable treatment dose rate out, a very high input dose rate is required. If the machine produces a photon beam with the beam flattener not in position, a high output dose rate results. This is the basic hazard of dual-mode machines: If the turntable is in the wrong position, the beam flattener will not be in place.

In the Therac-25, the computer is responsible for positioning the turntable (and for checking turntable position) so that a target, flattening filter, and X-ray ion chamber are directly in the beam path. With

the target in the beam path, electron bombardment produces X-rays. The X-ray beam is shaped by the flattening filter and measured by the X-ray ion chamber.

No accelerator beam is expected in the field-light position. A stainless steel mirror is placed in the beam path and a light simulates the beam. This lets the operator see precisely where the beam will strike the patient and make necessary adjustments before treatment starts. There is no ion chamber in place at this turntable position, since no beam is expected.

Traditionally, electromechanical interlocks have been used on these types of equipment to ensure safety – in this case, to ensure that the turntable and attached equipment are in the correct position when treatment is started. In the Therac-25, software checks were substituted for many traditional hardware interlocks.

Reference: J.A. Rawlinson, "Report on the Therac-25," OCFTR/OCI Physicists Meeting, Kingston, Ont., May 7, 1987.

Therac-25 Software Design

We know that the software for the Therac-25 was developed by a single person, using PDP-11 assembly language, over a period of several years. The software "evolved" from the Therac-6 software, which was started in 1972. According to a letter from CMC to the FDA, the "program structure and certain subroutines were carried over to the Therac-25 around 1976." Apparently, very little software documentation was produced during development. In a 1986 internal FDA memo, a reviewer lamented, "Unfortunately, the CMC response also seems to point out an apparent lack of documentation on software specifications and a software test plan."

The manufacturer said that the hardware and software were "tested and exercised separately or together over many years." In his deposition for one of the lawsuits, the quality assurance manager explained that testing was done in two parts. A "small amount" of software testing was done on a simulator, but most testing was done as a system. It appears that unit and software testing was minimal, with most effort directed at the integrated system test. At a Therac-25 user group meeting, the same quality assurance manager said that the Therac-25 software was tested for 2,700 hours. Under questioning by the users, he clarified this as meaning "2,700 hours of use."

The programmer left CMC in 1986. In a lawsuit connected with one of the accidents, the lawyers were unable to obtain information about the programmer from CMC. In the depositions connected with that case, none of the CMC employees questioned could provide any information about his educational background or experience. Although an attempt was made to obtain a deposition from the programmer, the lawsuit was settled before this was accomplished. We have been unable to learn anything about his background.

CMC claims proprietary rights to its software design. However, from voluminous documentation regarding the accidents, the repairs, and the eventual design changes, we can build a rough picture of it.

The software is responsible for monitoring the machine status, accepting input about the treatment desired, and setting the machine up for this treatment. It turns the beam on in response to an operator command (assuming that certain operational checks on the status of the physical machine are satisfied) and also turns the beam off when treatment is completed, when an operator commands it, or when a malfunction is detected. The operator can print out hard-copy versions of the CRT display or machine setup parameters.

The treatment unit has an interlock system designed to remove power to the unit when there is a hardware malfunction. The computer monitors this interlock system and provides diagnostic messages. Depending on the fault, the computer either prevents a treatment from being started or, if the treatment is in progress, creates a pause or a suspension of the treatment.

The manufacturer describes the Therac-25 software as having a stand-alone, real-time treatment operating system. The system is not built using a standard operating system or executive. Rather, the real-time executive was written especially for the Therac-25 and runs on a 32K PDP 11/23. A preemptive scheduler allocates cycles to the critical and noncritical tasks.

The software, written in PDP 11 assembly language, has four major components; stored data, a scheduler, a set of critical and noncritical tasks, and interrupt services. The stored data includes calibration parameters for the accelerator setup as well as patient-treatment data. The interrupt routines include:

- a clock interrupt service routing,
- a scanning interrupt service routing,
- traps (for software overflow and computer-hardware-generated interrupts),
- power up (initiated at power up to initialize the system and pass control to the scheduler),
- treatment console screen interrupt handler,
- treatment console keyboard interrupt handler,
- service printer interrupt handler,
- service keyboard interrupt handler.

The scheduler controls the sequences of all noninterrupt events and coordinates all concurrent processes. Tasks are initiated every 0.1 second, with the critical tasks executed first and the noncritical tasks executed in any remaining cycle time. Critical tasks include the following:

- The treatment monitor (Treat) directs and monitors patient setup and treatment via eight operating phases. These are called subroutines, depending on the value of the Tphase control variable. Following the execution of a particular subroutine, Treat reschedules itself. Treat interacts with the keyboard processing task, which handles operator console communication. The prescription data is cross-checked and verified by other tasks (for example, the keyboard processor and the parameter setup sensor) that inform the treatment task of the verification status via shared variables.
- The servo task controls gun emission, dose rate (pulse-repetition frequency), symmetry (beam steering), and machine motions. The servo task also sets up the machine parameters and monitors the beam-tilt-error and the flatness-error interlocks.

- The housekeeper task takes care of system-status interlocks and limit checks, and puts appropriate messages on the CRT display. It decodes some information and checks the setup verification.

-

Noncritical tasks include

- Check sum processor (scheduled to run periodically).
- Treatment console keyboard processor (scheduled to run only if it is called by other tasks or by keyboard interrupts). This task acts as the interface between the software and the operator.
- Treatment console screen processor (run periodically). This task lays out appropriate record formats for either displays or hard copies.
- Service keyboard processor (run on demand). This task arbitrates non-treatment-related communication between the therapy system and the operator.
- Snapshot (run periodically by the scheduler). Snapshot captures preselected parameter values and is called by the treatment task at the end of a treatment.
- Hard-control processor (run periodically).
- Calibration processor. This task is responsible for a package of tasks that let the operator examine and change system setup parameters and interlock limits.

It is clear from the CMC documentation on the modifications that the software allows concurrent access to shared memory, that there is no real synchronization aside from data stored in shared variables, and that the "test" and "set" for such variables are not indivisible operations. Race conditions resulting from this implementation of multitasking played an important part in the accidents.

Safety analysis of the Therac-25

The Therac-25 safety included (1) failure mode and effect analysis, (2) fault-tree analysis, and (3) software examination.

Failure mode and effect analysis. An FMEA describes the associated system response to all failure modes of the individual system components, considered one by one. When software was involved, CMC made no assessment of the "how and why" of software faults and took any combination of software faults as a single event. The latter means that if the software was the initiating event, then no credit was given for the software mitigating the effects. This seems like a reasonable and conservative approach to handling software faults.

Fault-tree analysis. An FMEA identifies single failures leading to Class I hazards. To identify multiple failures and quantify the results, CMC used fault-tree analysis. An FTA starts with a postulated hazard-- for example, two of the top events for the Therac-25 are high dose per pulse and illegal gantry motion. The immediate causes for the event are then generated in an AND/OR tree format, using a basic understanding of the machine operation to determine the causes. The tree generation continues until all branches end in the "basic events." Operationally, a basic event is sometimes defined as an event that can be quantified (for example, a resistor fails open).

CMC used a "generic failure rate" of 10^{-4} per hour for software events. The company justified this number as based on the historical performance of the Therac-25 software. The final report on the

safety analysis said that many fault trees for the Therac-25 have a computer malfunction as a causative event, and the outcome of quantification is therefore dependent on the failure rate chosen for software.

Leaving aside the general question of whether such failure rates are meaningful or measurable for software in general, it seems rather difficult to justify a single figure of this sort for every type of software error or software behavior. It would be equivalent to assigning the same failure rate to every type of failure of a car, no matter what particular failure is considered.

The authors of the safety study did note that despite the uncertainty that software introduces into quantification, fault-tree analysis provides valuable information in showing single and multiple failure paths and the relative importance of different failure mechanisms. This is certainly true.

Software examination. Because of the difficulty of quantifying software behavior, CMC contracted for a detailed code inspection to "obtain more information on which to base decisions." The software functions selected for examination were those related to the Class I software hazards identified in the FMEA: electron-beam scanning, energy selection, beam shutoff, and dose calibration.

The outside consultant who performed the inspection included a detailed examination of each function's implementation, a search for coding errors, and a qualitative assessment of its reliability. The consultant recommended program changes to correct shortcomings, improve reliability, or improve the software package in a general sense. The final safety report gives no information about whether any particular methodology or tools were used in the software inspection or whether someone just read the code looking for errors.

Conclusions of the safety analysis. The final report summarizes the conclusions of the safety analysis:

The conclusions of the analysis call for 10 changes to Therac-25 hardware; the most significant of these are interlocks to back up software control of both electron scanning and beam energy selection.

Although it is not considered necessary or advisable to rewrite the entire Therac-25 software package, considerable effort is being expended to update it. The changes recommended have several distinct objectives: improve the protection it provides against hardware failures; provide additional reliability via cross-checking; and provide a more maintainable source package. Two or three software releases are anticipated before these changes are completed.

The implementation of these improvements including design and testing for both hardware and software is well under way. All hardware modifications should be completed and installed by mid 1989, with final software updates extended into late 1989 or early 1990.

The recommended hardware changes appear to add protection against software errors, to add extra protection against hardware failures, or to increase safety margins. The software conclusions included the following:

The software code for Beam Shut-Off, Symmetry Control, and Dose Calibration was found to be straight-forward and no execution path could be found which would cause them to perform incorrectly. A few improvements are being incorporated, but no additional hardware interlocks are required.

Inspection of the Scanning and Energy Selection functions, which are under software control, showed no improper execution paths; however, software inspection was unable to provide a high level of confidence in their reliability. This was due to the complex nature of the code, the extensive use of variables, and the time limitations of the inspection process. Due to these factors and the possible clinical consequences of malfunction, computer-independent interlocks are being retrofitted for these two cases.

Given the complex nature of this software design and the basic multitasking design, it is difficult to understand how any part of the code could be labeled "straightforward" or how confidence could be achieved that "no execution paths" exist for particular types of software behavior. However, it does appear that a conservative approach-- including computer-independent interlocks-- was taken in most cases. Furthermore, few examples of such safety analyses of software exist in the literature. One such software analysis was performed in 1989 on the shutdown software of a nuclear power plant, which was written by a different division of CMC.¹ Much still needs to be learned about how to perform a software-safety analysis.

Reference

1. W. C. Bowman et al. "An Application of Fault Tree Analysis to Safety-Critical Software at Ontario Hydro," Conf. Probabilistic Safety Assessment and Management, 1991.

The Operator Interface

The Therac-25 operator controls the machine with a DEC VT100 terminal. In the general case, the operator positions the patient on the treatment table, manually sets the treatment field sizes and gantry rotation, and attaches accessories to the machine. Leaving the treatment room, the operator returns to the VT100 console to enter the patient identification, treatment prescription (including mode, energy level, dose, dose rate, and time), field sizing, gantry rotation, and accessory data. The system then compares the manually set values with those entered at the console. If they match, a "verified" message is displayed and treatment is permitted. If they do not match, treatment is not allowed to proceed until the mismatch is corrected. Figure A. shows the screen layout.

```

PATIENT NAME:TEST
TREATMENT MODE: FIX
BEAM TYPE: X ENERGY(KeV):
A      1
25

UNIT RATE/MINUTE      ACTUAL      PERSCRIBED
MONITOR UNITS          50  50      200
TIME (MIN)             0.27      1.00

GANTRY ROTATION (DEG)  0.0        0          VERIFIED
COLLIMATOR ROTATION (DEG) 359.2      359        VERIFIED
COLLIMATOR X(CM)       14.2       143        VERIFIED
COLLIMATOR Y(CM)       27.2       273        VERIFIED
WEDGE NUMBER           1          1          VERIFIED
ACCESSORY NUMBER       0          0          VERIFIED

DATE: 84-OCT-26        SYSTEM: BEAM READY      OP.MODE: TREAT      AUTO
TIME: 12:55.8         TREAT: TREAT PAUSE     X-RAY              173777
OPR ID: T25V02-RO3    REASON: OPERATOR      COMMAND:

```

Figure A. Operator interface screen layout

When the system was first built, operators complained that it took too long to enter the treatment plan. In response, the manufacturer modified the software before the first unit was installed so that, instead of reentering the data at the keyboard, operators could use a carriage return to merely copy the treatment site data [Miller]. A quick series of carriage returns would thus complete data entry. This interface modification was to figure in several accidents.

The Therac-25 could shut down in two ways after it detected an error condition. One was a treatment suspend, which required a complete machine reset to restart the machine. If a treatment pause occurred, the operator could press the "P" key to "proceed" and resume treatment quickly and conveniently. The previous treatment parameters remained in effect, and no reset was required. This convenient and simple feature could be invoked a maximum of five times before the machine automatically suspended treatment and required the operator to perform a system reset.

Error messages provided to the operator were cryptic, and some merely consisted of the word "malfunction" followed by a number from 1 to 64 denoting an analog/digital channel number. According to an FDA memorandum written after one accident:

The operator's manual supplied with the machine does not explain nor even address the malfunction codes. The [Maintenance] Manual lists the various malfunction numbers but gives no explanation. The materials provided give no indication that these malfunctions could place a patient at risk.

The program does not advise the operator if a situation exists wherein the ion chambers used to monitor the patient are saturated, thus are beyond the measurement limits of the instrument. This software package does not appear to contain a safety system to prevent parameters being entered and intermixed that would result in excessive radiation being delivered to the patient under treatment.

An operator involved in an overdose accident testified that she had become insensitive to machine malfunctions. Malfunction messages were commonplace – most did not involve patient safety. Service technicians would fix the problems or the hospital physicist would realign the machine and make it operable again. She said, "It was not out of the ordinary for something to stop the machine...It would often give a low dose rate in which you would turn the machine back on...They would give messages of low dose rate, V-tilt, H-tilt, and other things; I can't remember all the reasons it would stop, but there [were] a lot of them." The operator further testified that during instruction she had been taught that there were "so many safety mechanisms" that she understood it was virtually impossible to overdose a patient.

A radiation therapist at another clinic reported an average of 40 dose-rate malfunction, attributed to underdoses, occurred on some days.

Reference: E. Miller, "The Therac-25 Experience," Proc. Conf. State Radiation Control Program Directors, 1987.

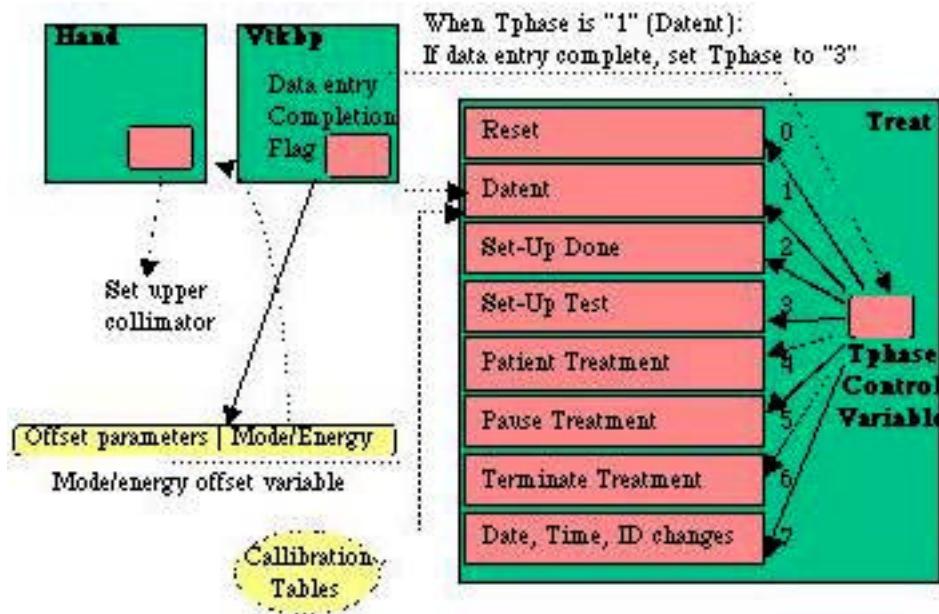
Therac-25: Tyler, TX Software Problem

A lesson to be learned from the Therac-25 story is that focusing on particular software bugs is not the way to make a safe system. Virtually all complex software can be made to behave in an unexpected fashion under certain conditions. The basic mistakes here involved poor software-engineering practices and building a machine that relies on the software for safe operation.

Furthermore, the particular coding error is not as important as the general unsafe design of the software overall. Examining the part of the code blamed for the Tyler accidents is instructive, however, in

showing the overall software design flaws. The following explanation of the problem is from the description CMC provided for the FDA, although we have tried to clarify it somewhat. The description leaves some unanswered questions, but it is the best we can do with the information we have.

As described in the sidebar on Therac-25 development and design, the treatment monitor task (Treat) controls the various phases of treatment by executing its eight subroutines (see Figure 2). The treatment phase indicator variable (Tphase) is used to determine which subroutine should be executed. Following the execution of a particular subroutine, Treat reschedules itself.



Tasks and subroutines in the code blamed for the Tyler accidents

Figure 2

One of Treat’s subroutines, called Datent (data entry), communicates with the keyboard handler task (a task that runs concurrently with Treat) via a shared variable (Data-entry completion flag) to determine whether the prescription data has been entered. The keyboard handler recognizes the completion of data entry and changes the Data-entry completion variable to denote this. Once the Data-entry completion variable is set, the Datent subroutine detects the variable’s change in status and changes the value of Tphase from 1 (Data Entry) to 3 (Set-Up Test).

In this case, the Datent subroutine exits back to the Treat subroutine, which will reschedule itself and begin execution of the Set-Up Test subroutine. If the Data-entry completion variable has not been set, Datent leaves the value of Tphase unchanged and exits back to Treat’s main line. Treat will then reschedule itself, essentially rescheduling the Datent subroutine.

The command line at the lower right corner of the screen is the cursor’s normal position when the operator has completed all necessary changes to the prescription. Prescription editing is signified by cursor movement off the command line. As the program was originally designed, the Data-entry completion variable by itself is not sufficient since it does not ensure that the cursor is located on the command line. Under the right circumstances, the data-entry phase can be exited before all edit changes are made on the screen.

The keyboard handler parses the mode and energy level specified by the operator and places an encoded result in another shared variable, the 2-byte mode/energy offset (MEOS) variable. The low-order byte of this variable is used by another task (Hand) to set the collimator/turntable to the proper position for the selected mode/energy. The high-order byte of the MEOS variable is used by Datent to set several operating parameters.

Initially, the data-entry process forces the operator to enter the mode and energy, except when the operator selects the photon mode, in which case the energy defaults to 25 MeV. The operator can later edit the mode and energy separately. If the keyboard handler sets the data-entry completion variable before the operator changes the data in MEOS, Datent will not detect the changes in MEOS since it has already exited and will not be reentered again. The upper collimator, on the other hand, is set to the position dictated by the low-order byte of MEOS by another concurrently running task (Hand) and can therefore be inconsistent with the parameters set in accordance with the information in the high-order byte of MEOS. The software appears to include no checks to detect such an incompatibility.

```

Datent:
  if mode/energy specified then
    begin
      calculate table index
      repeat
        fetch parameter
        output parameter
        point to next parameter
      until all parameters set
      call Magnet
      if mode/energy changed then return
    end
  if data entry is complete then set Tphase to 3
  if data entry is not complete then
    if reset command entered then set Tphase to 0
  return

Magnet:
  Set bending magnet flag
  repeat
    Set next magnet
    Call Ptime
    if mode/energy has changed then exit
  until all magnets are set
  return

Ptime:
  repeat
    if bending magnet flag is set then
      if editing taking place then
        if mode/energy has changed then exit
  until hysteresis delay has expired
  Clear bending magnet flag
  return

```

Figure 3. Datent, Magnet, and Ptime subroutines

The first thing that Datent does when it is entered is to check whether the mode/energy has been set in MEOS. If so, it uses the high-order byte to index into a table of preset operating parameters and places them in the digital-to-analog output table. The contents of this output table are transferred to the digital-analog converter during the next clock cycle. Once the parameters are all set, Datent calls the subroutine Magnet, which sets the bending magnets. Figure 3 is a simplified pseudocode description of relevant parts of the software.

Setting the bending magnets takes about 8 seconds. Magnet calls a subroutine called Ptime to introduce a time delay. Since several magnets need to be set, Ptime is entered and exited several times. A flag to indicate that bending magnets are being set is initialized upon entry to the Magnet subroutine and cleared at the end of Ptime. Furthermore, Ptime checks a shared variable, set by the keyboard handler, that indicates the presence of any editing requests. If there are edits, then Ptime clears the bending magnet variable and exits to Magnet, which then exits to Datent. But the edit change variable is checked by Ptime only if the bending magnet flag is set. Since Ptime clears it during its first execution, any edits performed during each succeeding pass through Ptime will not be recognized. Thus, an edit change of the mode or energy, although reflected on the operator's screen and the mode/energy offset variable, will not be sensed by Datent so it can index the appropriate calibration tables for the machine parameters.

Recall that the Tyler error occurred when the operator made an entry indicating the mode/energy, went to the command line, then moved the cursor up to change the mode/energy, and returned to the command line all within 8 seconds. Since the magnet setting takes about 8 seconds and Magnet does not recognize edits after the first execution of Ptime, the editing had been completed by the return to Datent, which never detected that it had occurred. Part of the problem was fixed after the accident by clearing the bending-magnet variable at the end of Magnet (after all the magnets have been set) instead of at the end of Ptime.

But this was not the only problem. Upon exit from the Magnet subroutine, the data-entry subroutine (Datent) checks the data-entry completion variable. If it indicates that data entry is complete, Datent sets Tphase to 3 and Datent is not entered again. If it is not set, Datent leaves Tphase unchanged, which means it will eventually be rescheduled. But the data-entry completion variable only indicates that the cursor has been down to the command line, not that it is still there. A potential race condition is set up. To fix this, CMC introduced another shared variable controlled by the keyboard handler task that indicates the cursor is not positioned on the command line. If this variable is set, then prescription entry is still in progress and the value of Tphase is left unchanged.

Therac-25: Yakima Software Problem

The software problem for the second Yakima accident is fairly well established and different from that implicated in the Tyler accidents. There is no way to determine what particular software design errors were related to the Kennestone, Hamilton, and first Yakima accidents. Given the unsafe programming practices exhibited in the code, it is possible that unknown race conditions or errors could have been responsible. There is speculation, however, that the Hamilton accident was the same as this second Yakima overdose. In a report of a conference call on January 26, 1987, between the CMC quality assurance manager and Ed Miller of the FDA discussing the Yakima accident, Miller notes:

This situation probably occurred in the Hamilton, Ontario, accident a couple of years ago. It was not discovered at that time and the cause was attributed to intermittent interlock failure. The subsequent recall of the multiple microswitch logic network did not really solve the problem.

The second Yakima accident was again attributed to a type of race condition in the software, this one allowed the device to be activated in an error setting (a "failure" of a software interlock). The Tyler accidents were related to problems in the data-entry routines that allowed the code to proceed to Set-Up Test before the full prescription had been entered and acted upon. The Yakima accident involves problems encountered later in the logic after the treatment monitor Treat reaches Set-Up Test.

The Therac-25's field-light feature permits very precise positioning of the patient for treatment. The operator can control the Therac-25 right at the treatment site using a small hand control offering certain limited functions for patient setup, including setting gantry, collimator, and table motions.

Normally, the operator enters all the prescription data at the console (outside the treatment room) before the final setup of all machine parameters is completed in the treatment room. This gives rise to an "unverified" condition at the console. The operator then completes the patient setup in the treatment room, and all relevant parameters now "verify". The console displays the message "Press set button" while the turntable is in the field-light position. The operator now presses the set button on the hand control or types "set" at the console. That should set the collimator to the proper position for treatment.

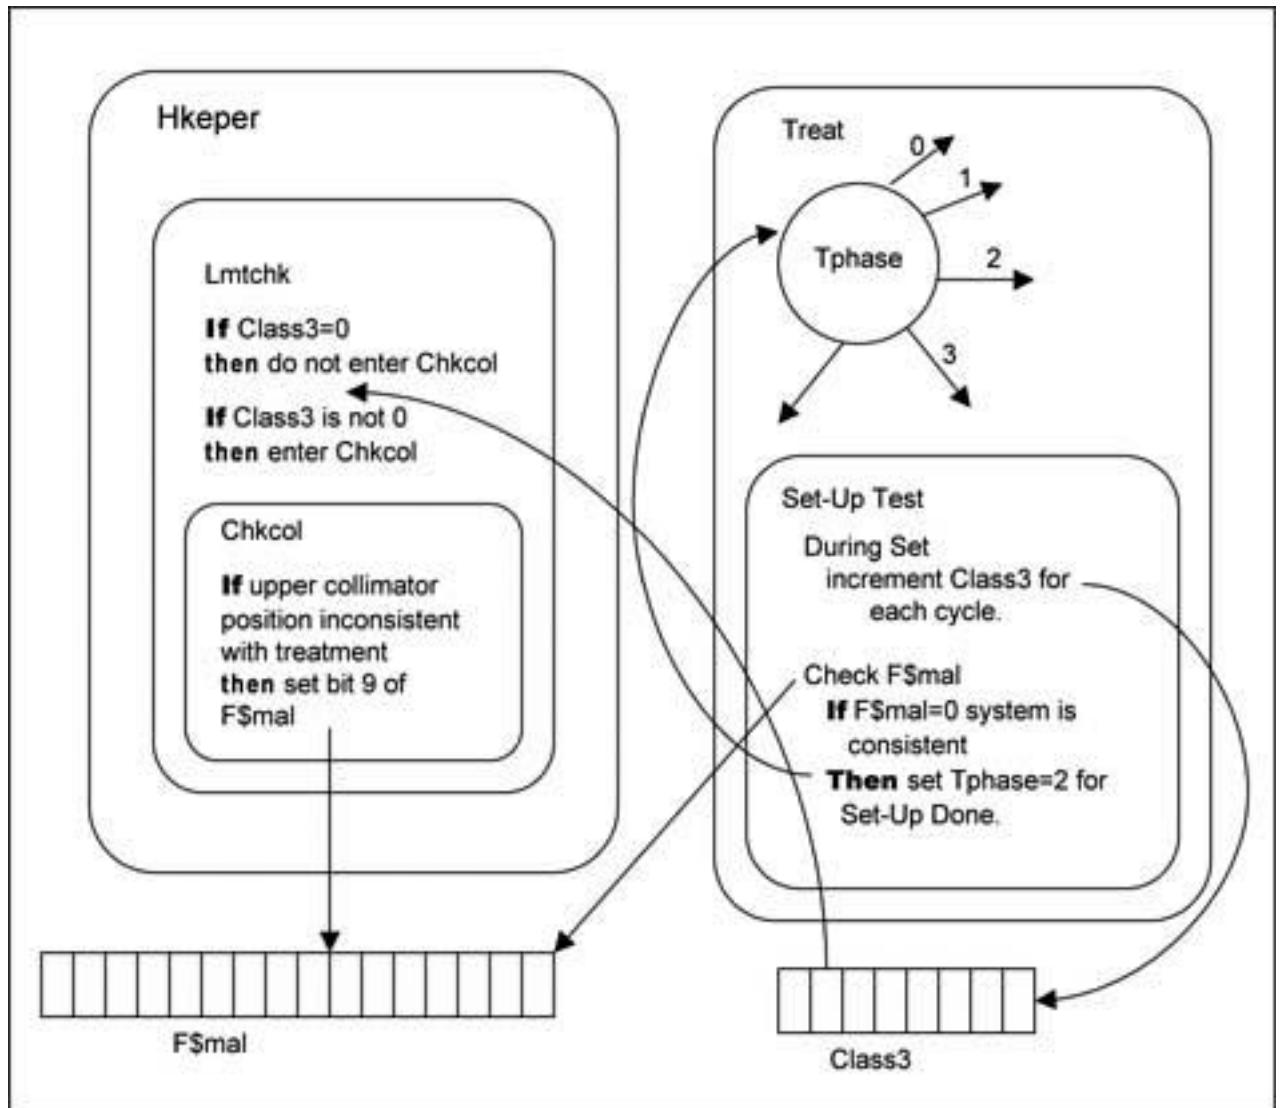


Figure 4. Yakima software flaw

In the software, after the prescription is entered and verified by the Datent routine, the control variable Tphase is changed so that the Set-Up Test routine is entered (see Figure 4). Every pass through the Set-Up Test routine increments the upper collimator position check, a shared variable called Class3. If Class3 is nonzero, there is an inconsistency and treatment should not proceed. A zero value for Class3 indicates that the relevant parameters are consistent with treatment, and the beam is not inhibited.

After setting the Class 3 variable, Set-Up Test next checks for any malfunctions in the system by checking another shared variable (set by a routine that actually handles the interlock checking) called F\$mal to see if it has a nonzero value. A nonzero value in F\$mal indicates that the machine is not ready for treatment, and the Set-Up Test subroutine is rescheduled. When F\$mal is zero (indicating that everything is ready for treatment), the Set-Up Test subroutine sets the Tphase variable equal to 2, which results in next scheduling the Set-Up Done subroutine, and the treatment is allowed to continue.

The actual interlock checking is performed by a concurrent Housekeeper task (Hkeper). The upper collimator position check is performed by a subroutine of Hkeper called Lmtchk (analog/digital limit checking). Lmtchk first checks the Class3 variable. If Class3 contains a nonzero value, Lmtchk calls the Check Collimator (Chkcol) subroutine. If Class3 contains zero, Chkcol is bypassed and the upper collimator position check is not performed. The Chkcol subroutine sets or resets bit 9 of the F\$mal shared variable, depending on the position of the upper collimator (which in turn is checked by the Set-Up Test subroutine of Datent so it can decide whether to reschedule itself or proceed to Set-Up Done).

During machine setup, Set-Up Test will be executed several hundred times since it reschedules itself waiting for other events to occur. In the code, the Class3 variable is incremented by one in each pass through Set-Up test. Since the class3 variable is 1 byte, it can only contain a maximum value of 255 decimal. Thus, on every 256th pass through the Set-Up Test code, the variable overflows and has a zero value. That means that on every 356th pass through Set-Up Test, the upper collimator fault will not be detected.

The overexposure occurred when the operator hit the "set" button at the precise moment that Class3 rolled over to zero. Thus Chkcol was not executed, and F\$mal was not set to indicate the upper collimator was still in field-light position. The software turned on the full 25 MeV without the target in place and without scanning. A highly concentrated electron beam resulted, which was scattered and deflected by the stainless steel mirror that was in the path.

CMC described the technical "fix" implemented for this software flaw as simple: The program is changed so that the Class3 variable is set to some fixed nonzero value each time through Set-Up Test instead of being incremented.

How to Produce a Malfunction 54 on a [CMC] Therac-25 Linear Accelerator

This statement was written by the East Texas Cancer Center physicist after he discovered how to reproduce the Malfunction 54 error

Enter the room and set up the machine for an electron beam treatment by selecting a field size and installing the trimmers. Press the set button. Leave the room and close the door. At the control console proceed to the patient set-up display. For Mode enter "X". The machine will default to 25 MeV and go to dose rate of 250 rads/min. Use return key to go to dose. Enter 200. Use return key to go to time. Enter 0.8 min. Use the return key to rapidly advance to the bottom of the display. Immediately use the up arrow to move from the bottom of the display. You are now in the edit mode. Use the up arrow to go to the top of the display and change the mode "X" to "E" for electrons. Change the energy from 25 to 10. Use the return key to go back down to the bottom of the display. Wait for the "beam ready" message then type "B" return. The unit will have no indications on dose rate or dose 1 or dose 2 for about 3 to 4 seconds. Then the dose rate will flash 550 to 575 for one cycle and return to zero. Dose 1 and Dose 2 will count to -6. A malfunction 54 message will appear at the bottom of the display. You have just delivered a dose of approximately 25,000 rads of 25 MeV electrons in less than two seconds.

Summary of Operator Interview

The following article is the result of an interview we conducted with a Registered Therapy Technologist who has extensive experience operating medical linear accelerators. This individual currently manages a Radiation Therapy Department at a major United States hospital, and trains technicians to operate radiation therapy machinery. For privacy purposes, the true identity of this person will remain anonymous, and for the remainder of the article, we will refer to our interviewee as "Susan."

Susan operated a Therac-4 linear accelerator machine in the mid 1980's. At the time, Susan had recently graduated and was working at a University where the radiation therapy technology was fairly advanced. She enjoyed operating CMC's Therac machine because it was one of the first computerized linear accelerators. Looking back, Susan remembered that while operating the machines, she did not think much about whether there could be computer software "bugs" in the system. The technology was new, and she remembered trusting the machine's components and its designers.

When recalling the advantages of the new computerized machine, Susan reported being able to move more patients through during the day. She also remembered feeling good about the extra time she had to talk with patients when she was working with a computerized machine.

Susan learned about the Therac-25 incidents while attending a national radiation therapy conference in 1990. A radiation therapist who was also a lawyer gave a lecture on the Therac-25 accidents. He handed out newspaper articles about the incidents and spoke about how many times the therapists involved in the accidents attempted to resume treatment in spite of the error messages they received from the computer. The lecture focused on the question of how many attempts to resume treatment is too many? The lecturer and the participants discussed the possibility of establishing institutional policies and limits on the number of times an operator could resume treatment after having received an error message, such as the cryptic "malfunction 54" messages that the operator received during the two fatal accidents in Texas.

The problem, Susan reported, is that back in 1990, and today, there are no industry-wide standards or rules for these types of situations. Susan felt that she had been lucky to have always worked where there was a physicist available to provide help with the many error messages operators received. She also felt that in other clinics, where this kind of assistance is not available, there was, and still is, a great deal more pressure on therapists to just keep going despite the error messages. An operator might attempt, for example, to deliver the prescribed dose in 12 increments instead of 1 by continually clearing the faults generated by the computer. Susan stated that this type of activity happens all the time in medical radiation therapy, particularly in clinics where there is more pressure from the administration to keep patients moving through quickly.

Although Susan had been working with a CMC Therac machine at the time of the accidents, she did not remember receiving warning notices from CMC about the Therac-related accidents. Susan believes that this is one aspect of the industry that has changed, possibly, in part due the Therac-25 accidents. At the present time Susan receives notices from the manufacturers of the linear accelerators used at her hospital whenever there is a linear accelerator malfunction, or even if there is a malfunction that almost occurred, but was prevented.

Perhaps part of the reason that Susan did not hear of the Therac-25 incidents until much later was that the hospital where she worked got rid of the Therac-4, moved their facilities, and bought a new set of linear accelerators. Susan estimated the average life of the linear accelerator to be between 5 and 10 years. After that, she said, the accelerator tends to act somewhat like an old car in which the engine light is coming on all the time. The accelerator's computer generates many faults that can become a nuisance to the operators and to the patients. Responsible operators will continue to report these faults to the physicist, when one is available, and eventually, the machine is replaced.

Susan feels that one of the biggest problems in her industry today is the lack of rigorous industry-wide standard certification and education for operators. Susan reported that there are about 102 radiation schools in the country, and that there are also different types of schools. Students are able to receive a certificate from a certificate program, usually about 12 months in length. Students are also able to receive a four-year bachelor's degree from certain schools. The American Registry of Radiologic Technologists (ARRT) provides a test that graduates of these programs may then take in order to be considered licensed entry level technicians. The ARRT also requires that therapists maintain their training through continuing education. Therapists must have 24 credits in two years before they may re-register their licenses.

In spite of the fact that the ARRT provides these guidelines for licensure, many states in the U.S. do not require hospitals or clinics to hire licensed radiation therapists. Some states require very basic exams, but, according to Susan, that in essence means that in many states anyone off the street could learn how to operate a machine, take one of these basic exams, and then be qualified to operate radiation therapy machines.

Susan and many of her colleagues continue to fight for mandatory standard certification of radiation therapists. The safety of patients depends on all of the elements of their systems of treatment working together correctly. The more operators are trained to know about the process, the more they will be able to help prevent accidents. Well-trained operators can double-check radiation dose prescriptions and question doctors when something does not seem right. With the benefit of extensive training, operators have a better sense of when it is alright to over-ride a fault message from the computer.

Well trained technicians will also be better equipped to stand up to hospital administrations that attempt to put pressure on technicians to push large numbers of patients through treatment in spite of possible dangers. Though Susan does not feel this kind of pressure from her own administration, she knows that other technicians in other clinics definitely do, especially at "free-standing" clinics that operate for profit. Susan is aware that at these clinics there is a tremendous amount of pressure put on machine operators to get patients through treatment.

Susan also described incidents in which technicians left institutions because they didn't feel that the institutions' radiation therapy practices were safe for patients. Because there is no federal law regulating how many times an operator can re-attempt therapy after the computer displays a fault or shuts down, some operators allegedly use jumper cables that continuously override their computer's emergency shut down mechanism. Susan cited a lack of regulation, lack of training, and lack of adequate funding as reasons for these procedures.

Another issue in the radiation therapy industry that worries Susan is the fact that linear accelerator manufacturers charge large fees for operator training sessions, software upgrades, and machine maintenance contracts. When a radiation therapy machine is purchased, it comes with many binders full of information provided by the company. The clinic is given the option to buy service contracts and send physicists and operators to the company headquarters for training. Susan reported that in many clinics where money is tight, administrators are forced to choose between machine servicing contracts, software upgrades, and training.

According to Susan, mistakes are still made in the radiation therapy treatment of patients. Much of the information and calibration is still done by human beings and subject to human error. As an instructor, Susan teaches her students to anticipate every angle of the treatment, and then to check, and re-check their work. Susan also mentioned that while she teaches her students not to trust wholly in the machinery and its software, operators are largely dependent on manufacturers and hospital physicist teams to keep the machines running correctly.

Susan has a positive outlook regarding the radiation therapy industry. She knows that thousands of patients benefit greatly from radiation therapy technology. While Susan continues to push for operator certification legislation, she focuses on training her own staff well. Susan and her administration also focus heavily on quality patient care.

When asked if she thought it would be important for the designers of the software that runs the machines to know what it is like to do her job, Susan's reply was an emphatic yes, though she doubted many of the software designers of her machinery had spent much time observing a radiation treatment facility.

Some Useful Therac Sources

1. N. Leveson, C. Turner, "An Investigation of the Therac-25 Accidents," In *Ethics and Computing: Living Responsibly in a Computerized World*, by K. W. Bowyer. Los Alamitos, CA: IEEE Computer Society Press, 1996. First Published in *Computer*, Vol. 26. No. 7, July 1993, pp. 18-41.

The classical report on the Therac 25 system from someone who was involved in the cases as an expert witness. Some portions of this article are provided in the resources sections of this case. This piece is packed with detail and will take some time for even a careful reader to fully grasp. This reprint of the article is in an excellent volume that contains many other cases that computer science instructors will find helpful.

2. N. G. Leveson. *Safeware: System Safety and Computers*. Reading: Addison-Wesley Publishing Company, 1995

An excellent source on software systems and their role in the safety of larger industrial and military systems. This large text pays dividends to the extent that one reads closely. Instructors who want to teach the Therac 25 system from a safety perspective will be well served if they have recourse to this volume.

3. E.J. Joyce, "Malfunction 54: Unraveling Deadly Medical Mystery of Computerized Accelerator Gone Awry," *American Medical News*, Oct. 3, 1986, pp. 1,13-17

An accessible and journalistic approach to the Therac case. Has pictures of the people involved in the case and of the machine. This is part of a series of articles in *American Medical News* that

is quite helpful.

4. J. Reason, *Human Error*. Cambridge: Cambridge University Press, 1990.

Reason is an expert on human-machine interactions. His book contains excellent analyses of system safety, and many detailed studies of system accidents, with references. His analysis of how systems should be built to take advantage of human abilities (rather than to emphasize their shortcomings) is well worth the price of the book.

5. C. Perrow. *Normal Accidents: Living with High-Risk Technologies*. Princeton, NJ: Princeton University Press, 1984.

A classic text in system safety. Perrow's argument is that technological systems are becoming so complex to build and maintain that catastrophic accidents will become normal.

Additional Therac Sources

C.A. Bowsher, "Medical Devices: The Public Health at Risk," US Gov't General Accounting Office Report GAO/T-PEMD-90-2. 046987/139922, November 6, 1989.

E. Chelimsky, "Medical Devices: Early Warning of Problems is Hampered by Severe Underreporting," GAO/PEMD-87-1, December, 1986.

J. Jacky, "Safety-Critical Computing: Hazards, Practices, Standards, and Regulation." In *Computerization and controversy: Value Conflicts and Social Choices*, edited by C. Dunlop and R. King. San Diego: Academic Press, Inc., 1991. First published in *The Sciences*, September/ October, 1989.

E.J. Joyce, "Accelerator Linked to 5th Radiation Overdose," *American Medical News*, Feb. 6, 1987, pp. 1,49-50.

E.J. Joyce, "Software 'Bug' Discovered in Second Linear Accelerator," *American Medical News*, Nov. 7, 1986, pp.20-21.

M. Kivel, "FDA Monitoring Correction of Therac Radiation Therapy Units," *Radiological Health Bulletin*, Vol.XX, No. 8, Jan. 7, 1987, pp. 1-2.

M. Kivel, "Therac-25 Accelerator Purchasers Advised to Discontinue Routine Use," *Radiological Health Bulletin*, Vol. XXI, No. 3, March 1987, pp. 1-2.

D. Lacasse, "Kanata Firm Works to Prevent Repeat Errors; Complexity of Software Used in 'Safety-Critical' Environments Makes It Impossible to Ensure Absolute Reliability," *The Ottawa Citizen*, Final Edition, Sept. 15, 1991, p. E6.

E. Miller, "The Therac-25 Experience," *Proc. Eighteenth Annual National Conf. on Radiation Control*, Jan. 1987, pp. 101-105.

R. Pear, "Group Asking U.S. for New Vigilance in Patient Safety," *New York Times*, Nov. 30, 1999. W.

Plummer, "A Computer Glitch Turns Miracle Machine into Monster for Three Cancer Patients," People Weekly, Vol. 26, Nov. 24, 1986, pp. 48-51.

J.A. Rawlinson. "Report on the Therac-25," OCFRF/OCI Physicists Meeting, Kingston, Ont., Canada, May 7, 1987.

B. Steinhardt, "Medical Devices: FDA Can Improve Oversight of Tracking and Recall Systems," GAO/HEHS-98-211, September, 1998.

R.C. Thompson, "Faulty Therapy Machines Cause Radiation Overdoses," FDA Consumer, Vol. 21, No.10, Dec-Jan. 1987, p.37-38.