



Online Ethics Center
FOR ENGINEERING AND SCIENCE

Therac-25 Case Narrative

Author(s)

Charles Huff

Year

2003

Description

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

Body

Therac Table of Contents

Case Narrative

[Therac-25 Abstract](#)

[Therac-25 Introduction](#)

[How Radiation Therapy Works](#)

The Machine

[Basic Principles](#)

[Machine Design](#)

[Software Design](#)

[System Safety](#)

The Participants

[Hospital](#)

[FDA](#)

[Operators](#)

[AECL](#)

Accident Accounts

[Linda Knight](#)

[Donna Gartner](#)

[Janis Tilman](#)

[Isaac Dahl](#)

[Daniel McCarthy](#)

[Anders Engman](#)

Case Narrative

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 Introduction

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 shrunk 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


Basic Principles

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 [Schematic diagram of a typical medical accelerator used in cancer radiotherapy.] 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).

Image



[A simulated cross section view of radiation dispersing upon entering the body.]

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.

Machine Design

How Therac-25 worked

A Short History of Therac

There were two previous versions of Therac machines, each produced by AECL 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, AECL 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 AECL 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. AECL 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

[A cross section drawing of a Therac-25 facility, including technological devices and electronic switches.]

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.

Image




Figure B. Upper turntable assembly

[A cross section drawing of the Therac-25 upper turntable components.]

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.

Software Design

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

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

Machine

- monitoring the machine status
- placement of turntable
- strength and shape of beam
- operation of bending and scanning magnets
- setting the machine up for the specified treatment
- turning the beam on
- 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 [A schematic diagram of a typical medical accelerator used in cancer radiotherapy.] 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, on

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).



Machine-Based Safety Mechanisms




Figure 1. Typical Therac-25 facility

[A cross section drawing of a Therac-25 facility, including technological devices and electronic switches.]

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. AECL 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 AECL made the Therac-25 commercially available, AECL 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.

The 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*.

1. Class I devices, "general controls provide reasonable assurance of safety and effectiveness," for example bedpans and tongue depressors
2. Class II devices, such as syringes and hearing aids, "require performance standards in addition to general controls"
3. 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:

1. Proof of Pre-market Equivalence to another device on the market, termed 501(k)
2. 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 (AECL), 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 Atomic Energy Canada, Limited and Therac-25

Early Therac Machines

The story of Therac-25 begins in the early 1970's when Atomic Energy Canada, Limited (AECL) 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. AECL 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, AECL 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 AECL 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, AECL 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. AECL 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 AECL was bringing in its new beam folding technology (and the new Therac-25) on its own.

In tightly coupling the software and the hardware, AECL 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, AECL 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 AECL made the Therac-25 commercially available, AECL 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 AECL 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 AECL 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 AECL of the incident. This was the first time AECL had heard from a clinic about an overdose problem with the Therac-25 machine. AECL sent a service engineer to investigate.

AECL 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, AECL issued a voluntary recall of the Therac-25 machines and the FDA audited AECL's modifications to the Therac. AECL 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 AECL and spoke on the phone with AECL'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 AECL, and partly because they knew AECL 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 AECL 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 AECL 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 AECL. 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 AECL of this discovery and AECL was eventually able to reproduce the error. AECL advised Therac-25 users to physically remove the up-arrow key as a short-term solution. AECL 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 AECL to identify the software problem and correct it. The FDA also requested that AECL 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.

AECL 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. AECL could not reproduce the error.

Later that week, AECL sent an engineer to Yakima to investigate. The hospital physicist had also been running tests. They eventually discovered a software flaw and fixed it. AECL 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.

[Continue to Exercises for Therac-25](#)

Notes

Material developed as part of ComputingCases.org by Dr. Charles Huff of St. Olaf College.

Resource Type

Case Study / Scenario
Educational Activity Description

Parent Collection

Therac-25

Topics

Catastrophes, Hazards, Disasters

Product Liability

Safety

Discipline(s)

Computer Sciences

Computer, Math, and Physical Sciences

Research Ethics