

A

Medical Devices: The Therac-25 Story

A.1 Introduction

Between June 1985 and January 1987, a computer-controlled radiation therapy machine, called the Therac-25, massively overdosed six people. These accidents have been described as the worst in the 35-year history of medical accelerators [284].

A detailed accident investigation, drawn from publicly available documents, can be found in Leveson and Turner [187]. The following account is taken from this report and includes both the factors involved in the overdoses themselves and the attempts by the users, manufacturers, and governments to deal with them. Because this accident was never officially investigated, only partial information on the Therac-25 software development, management, and quality control procedures is available. What is included below has been gleaned from law suits and depositions, government records, and copies of correspondence and other material obtained from the U.S. Food and Drug Administration (FDA), which regulates these devices.

A.2 Background

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 1970s, Atomic Energy of Canada Limited (AECL)¹ and a French company called CGR went into business together building linear accelerators. The products of this 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 MeV, 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. We know that some of the old Therac-6 software routines were reused in the Therac-20 and that CGR developed the initial software.

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.

The business relationship between AECL 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-1970s, AECL had 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 economical to produce. Using this double-pass concept, AECL designed the Therac-25, a dual-mode linear accelerator that can deliver either photons at 25 MeV or electrons at various energy levels.

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 build-up 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).

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 computer. However, AECL designed the Therac-25 to take advantage of computer control from the outset; they 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 the electron-beam scanning plus mechanical interlocks for policing the machine and ensuring safe operation. The Therac-25 relies more on software for these functions. AECL took advantage of the com-

¹ AECL was an arms-length entity, called a crown corporation, of the Canadian government. Since the time of the incidents related in this paper, AECL Medical, a division of AECL, was privatized and is now called Theratronics International, Ltd. Currently, the primary business of AECL is the design and installation of nuclear reactors.

puter’s abilities to control and monitor the hardware and decided not to duplicate all the existing hardware safety mechanisms and interlocks.

Some software for the machines was interrelated or reused. In a letter to a Therac-25 user, the AECL quality assurance manager said, “The same Therac-6 package was used by the AECL 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” [187]. The reuse of Therac-6 design features or modules may explain some of the problematic aspects of the Therac-25 software 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.

AECL produced the first hardwired prototype of the Therac-25 in 1976, and the completely computer-controlled commercial version was available in late 1982.

Turntable Positioning. The Therac-25 turntable design plays an important role in the accidents. The upper turntable (see Figure A.1) 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, but rather is used to facilitate correct positioning of the patient. Because the accessories appropriate to each mode are physically attached to the turntable, proper operation of the Therac-25 is heavily dependent on the turntable position, which is monitored by three microswitches.

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 are used to 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 (or photon) therapy, only one energy level is available: 25 MeV. Much greater electron-beam current is required for X-ray mode (some 100 times greater than that for electron therapy) [284] 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; thus, to get a reasonable treatment dose rate out of the flattener, a very high input dose rate is required. If the machine should produce a photon beam with the beam flattener not in position, a high output dose to the patient 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




FIGURE A.1
Upper turntable assembly.

(and for checking the turntable position) so that a target, flattening filter, and X-ray ion chamber are directly in the beam path. With the target in place, 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 third or field light turntable 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 of the traditional hardware interlocks.

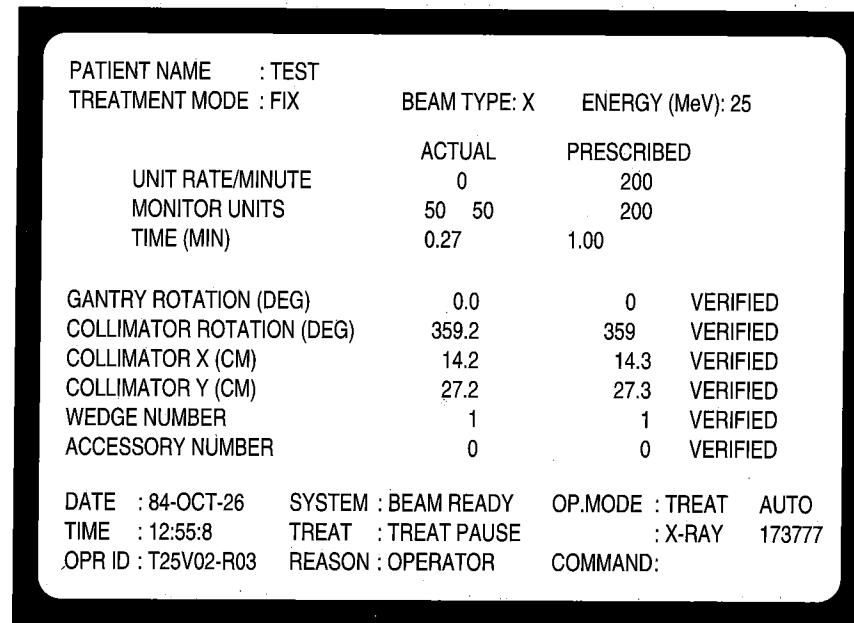


FIGURE A.2
Operator interface screen layout.

The Operator Interface. The description of the operator interface here applies to the version of the software used during the accidents. Changes made as a result of an FDA recall are described later.

The Therac-25 operator controls the machine through 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 console to enter the patient identification, treatment prescription (including mode or beam type, 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.2 shows the screen layout.

When the system was first built, operators complained that it took too long to enter the treatment plan. In response, AECL modified the software before the first unit was installed: Instead of reentering the data at the keyboard, operators could simply use a carriage return to copy the treatment site data [225]. A quick

series of carriage returns would thus complete the data entry. This modification was to figure in several of the 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 other, not so serious, was a *treatment pause*, which only required a single key command to restart the machine. If a *treatment pause* occurred, the operator could press the $\text{\textcircled{P}}$ key to “proceed” and resume treatment quickly and conveniently. The previous treatment parameters remained in effect, and no reset was required. This 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 [sic] 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 one of the accidents testified that she had become insensitive to machine malfunctions. Malfunction messages were commonplace and 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 was a lot of them.

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

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.

Hazard Analysis. In March 1983, AECL 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 about the computer and its software:

1. 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.
2. Program software does not degrade due to wear, fatigue, or reproduction process.
3. 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 the basic assumptions above, it considers hardware failures only. 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.

A.3 Events

Eleven Therac-25s were installed: five in the United States and six in Canada. Six accidents occurred between 1985 and 1987, when the machine was finally recalled to make extensive design changes. These changes include adding hardware safeguards against software errors.

Related problems were found in the Therac-20 software, but they were not recognized until after the Therac-25 accidents because the Therac-20 includes hardware safety interlocks. Thus, no injuries resulted.

A.3.1 Kennestone Regional Oncology Center, June 1985

Details of this accident in Marietta, Georgia, are sketchy because it was never investigated. There was no admission that the injury was caused by the Therac-25 until long after the occurrence, despite claims by the patient that she had been injured during treatment, the obvious and severe radiation burns the patient suffered, and the suspicions of the radiation physicist involved.

After undergoing a lumpectomy to remove a malignant breast tumor, a 61-year-old woman was receiving follow-up radiation treatment to nearby lymph nodes on a Therac-25 at the Kennestone facility in Marietta. The Therac-25 had been operating at Kennestone for about six months; other Therac-25s had been operating, apparently without incident, since 1983.

On June 3, 1985, the patient was set up for a 10 MeV electron treatment to the clavicle area. When the machine turned on, she felt a “tremendous force

of heat . . . this red-hot sensation.” When the technician came in, she said, “You burned me.” The technician replied that that was impossible. Although there were no marks on the patient at the time, the treatment area felt “warm to the touch.”

It is unclear exactly when AECL learned about this incident. Tim Still, the Kennestone physicist, said that he contacted AECL to ask if the Therac-25 could operate in electron mode without scanning to spread the beam. Three days later the engineers at AECL called the physicist back to explain that improper scanning was not possible.

In an August 19, 1986 letter from AECL to the FDA, the AECL quality assurance manager said, “In March of 1986 AECL received a lawsuit from the patient involved . . . This incident was never reported to AECL prior to this date, although some rather odd questions had been posed by Tim Still, the hospital physicist.” The physicist at a hospital in Tyler, Texas, where a later accident occurred, reported, “According to Tim Still, the patient filed suit in October 1985 listing the hospital, manufacturer and service organization responsible for the machine. AECL was notified informally about the suit by the hospital, and AECL received official notification of a lawsuit in November 1985.”

Because of the lawsuit (filed November 13, 1985), someone at AECL must have known about the Marietta accident—although no investigation occurred at this time. FDA memos point to the lack of a mechanism in AECL to follow up reports of suspected accidents [187].

The patient went home, but shortly afterward she developed a reddening and swelling in the center of the treatment area. Her pain had increased to the point that her shoulder “froze,” and she experienced spasms. She was admitted to a hospital in Atlanta, but her oncologists continued to send her to Kennestone for Therac-25 treatments. Clinical explanation was sought for the reddening of the skin, which at first her oncologist attributed to her disease or to normal treatment reaction.

About two weeks later, the Kennestone physicist noticed that the patient had a matching reddening on her back as though a burn had gone right through her body, and the swollen area had begun to slough off layers of skin. Her shoulder was immobile, and she was apparently in great pain. It was now obvious that she had a radiation burn, but the hospital and her doctors could provide no satisfactory explanation.

The Kennestone physicist later estimated that the patient received one or two doses of radiation in the 15,000 to 20,000 rad (radiation absorbed dose) range. He did not believe her injury could have been caused by less than 8,000 rads. To understand the magnitude of this, consider that typical single therapeutic doses are in the 200 rad range. Doses of 1,000 rads can be fatal if delivered to the whole body; in fact, 500 rads is the accepted figure for whole-body radiation that will cause death in 50 percent of the cases. The consequences of an overdose to a smaller part of the body depend on the tissue’s radio-sensitivity. The director of radiation oncology at the Kennestone facility explained their confusion about the accident as due to the fact that they had never seen an overtreatment of that magnitude before [308].

Eventually, the patient’s breast had to be removed because of the radiation burns. Her shoulder and arm were paralyzed, and she was in constant pain. She had suffered a serious radiation burn, but the manufacturer and operators of the machine refused to believe that it could have been caused by the Therac-25. The treatment prescription printout feature of the computer was disabled at the time of the accident, so there was no hardcopy of the treatment data. The lawsuit was eventually settled out of court.

From what we can determine, the accident was not reported to the FDA until *after* further accidents in 1986. The reporting requirements for medical device incidents at that time applied only to equipment manufacturers and importers, not users. The regulations required that manufacturers and importers report deaths, serious injuries, or malfunctions that could result in those consequences, but health-care professionals and institutions were not required to report incidents to manufacturers. The comptroller general of the U.S. Government Accounting Office (GAO), in testimony before Congress on November 6, 1989, expressed great concern about the viability of the incident-reporting regulations in preventing or spotting medical device problems. According to a 1990 GAO study, the FDA knew of less than 1 percent of deaths, serious injuries, or equipment malfunctions that occurred in hospitals [36]. The law was amended in 1990 to require health-care facilities to report incidents to the manufacturer and to the FDA.

At this point, the other Therac-25 users were also unaware that anything untoward had occurred and did not learn about any problems with the machine until after subsequent accidents. Even then, most of their information came through personal communication among themselves.

A.3.2 Ontario Cancer Foundation, July 1985

The second in this series of accidents occurred about seven weeks after the Kennestone patient was overdosed. At that time, the Therac-25 at the Ontario Cancer Foundation in Hamilton, Ontario (Canada), had been in use for more than six months. On July 26, 1985, a forty-year-old patient came to the clinic for her twenty-fourth Therac-25 treatment for carcinoma of the cervix. The operator activated the machine, but the Therac shut down after five seconds with an HTILT error message. The Therac-25’s console display read NO DOSE and indicated a TREATMENT PAUSE.

Since the machine did not suspend and the control display indicated no dose was delivered to the patient, the operator went ahead with a second attempt at treatment by pressing the Ⓢ key (the *proceed* command), expecting the machine to deliver the proper dose this time. This was standard operating procedure, and Therac-25 operators had become accustomed to frequent malfunctions that had no untoward consequences for the patient. Again, the machine shut down in the same manner. The operator repeated this process four times after the original attempt—the display showing NO DOSE delivered to the patient each time. After the fifth pause, the machine went into treatment suspend, and a hospital service

technician was called. The technician found nothing wrong with the machine. According to a Therac-25 operator, this scenario also was not unusual.

After the treatment, the patient complained of a burning sensation, described as an “electric tingling shock” to the treatment area in her hip. Six other patients were treated later that day without incident. She came back for further treatment on July 29 and complained of burning, hip pain, and excessive swelling in the region of treatment. The patient was hospitalized for the condition on July 30, and the machine was taken out of service.

AECL was informed of the apparent radiation injury and sent a service engineer to investigate. The U.S. FDA, the then Canadian Radiation Protection Bureau (RPB),² and users were informed that there was a problem, although the users claim that they were never informed that a patient injury had occurred. Users were told that they should visually confirm the proper turntable alignment until further notice (which occurred three months later).

The patient died on November 3, 1985, of an extremely virulent cancer. An autopsy revealed the cause of death as the cancer, but it was noted that had she not died, a total hip replacement would have been necessary as a result of the radiation overexposure. An AECL technician later estimated the patient had received between 13,000 and 17,000 rads.

Manufacturer's Response

AECL could not reproduce the malfunction that had occurred, but suspected a transient failure in the microswitch used to determine the turntable position. During the investigation of the accident, AECL hardwired the error conditions they assumed were necessary for the malfunction and, as a result, found some turntable positioning design weaknesses and potential mechanical problems.

The computer senses and controls turntable position by reading a 3-bit signal about the status of three microswitches in the turntable switch assembly. Essentially, AECL determined that a 1-bit error in the microswitch codes (which could be caused by a single open-circuit fault on the switch lines) could produce an ambiguous position message to the computer. The problem was exacerbated by the design of the mechanism that extends a plunger to lock the turntable when it is in one of the three cardinal positions: The plunger could be extended when the turntable was way out of position, thus giving a second false position indication. AECL devised a method to indicate turntable position that tolerated a 1-bit error so that the code would still unambiguously reveal correct position with any one microswitch failure.

In addition, AECL altered the software so that the computer checked for “in transit” status of the switches to keep further track of the switch operation and turntable position and to give additional assurance that the switches were working and the turntable was moving.

As a result of these improvements, AECL claimed in its report and correspondence with hospitals that “analysis of the hazard rate of the new solution indicates an improvement over the old system by at least *5 orders of magnitude* [emphasis added].” However, in its final incident report to the FDA, AECL concluded that they “cannot be firm on the exact cause of the accident but can only suspect . . . ,” which underscored their inability to determine the cause of the accident with any certainty. The AECL quality assurance manager testified that they could not reproduce the switch malfunction and that testing of the microswitch was “inconclusive.” The similarity of the errant behavior and the patient injuries in this accident and a later one in Yakima, Washington, provide good reason to believe that the Hamilton overdose was probably related to software error rather than to a microswitch failure.

Government and User Response

The Hamilton accident resulted in a voluntary recall by AECL, and the FDA termed it a Class II recall. Class II means “a situation in which the use of, or exposure to, a violative product may cause temporary or medically reversible adverse health consequences or where the probability of serious adverse health consequences is remote.” The FDA audited AECL’s subsequent modifications, and after the modifications were made, the users were told they could return to normal operating procedures.

As a result of the Hamilton accident, the head of advanced X-ray systems in the Canadian RPB, Gordon Symonds, wrote a report that analyzed the design and performance characteristics of the Therac-25 with respect to radiation safety. Besides citing the flawed microswitch, the report faulted both hardware and software components of the Therac’s design. It concluded with a list of four modifications to the Therac-25 necessary for compliance with Canada’s Radiation Emitting Devices (RED) Act. The RED law, enacted in 1971, gives government officials power to ensure the safety of radiation-emitting devices.

The modifications specified in the Symonds report included redesigning the microswitch and changing the way the computer handled malfunction conditions. In particular, treatment was to be terminated in the event of a dose-rate malfunction, giving a treatment “suspend.” This change would have removed the option to proceed simply by pressing the ⊕ key. The report also made recommendations regarding collimator test procedures and message and command formats. A November 8, 1985 letter, signed by the director of the Canadian RPB, asked that AECL make changes to the Therac-25 based on the Symonds report “to be in compliance with the RED act.”

Although, as noted above, AECL did make the microswitch changes, they did not comply with the directive to change the malfunction pause behavior into treatment suspends, instead reducing the maximum number of retries from five to three. According to Symonds, the deficiencies outlined in the RPB letter of November 8 were still pending when the next accident happened five months later.

² On April 1, 1986, the Radiation Protection Bureau and the Bureau of Medical Devices were merged to form the Bureau of Radiation and Medical Devices (BRMD).

Immediately after the Hamilton accident, the Ontario Cancer Foundation hired an independent consultant to investigate. He concluded in a September 1985 report that an independent system (beside the computer) was needed to verify the turntable position and suggested the use of a potentiometer. The RPB wrote a letter to AECL in November 1985 requesting that AECL install such an independent interlock on the Therac-25. Also, in January 1986, AECL received a letter from the attorney representing the Hamilton clinic. The letter said that there had been continuing problems with the turntable, including four incidents at Hamilton, and requested the installation of an independent system (potentiometer) to verify the turntable position. AECL did not comply: No independent interlock was installed by AECL on the Therac-25s at this time. The Hamilton clinic, however, decided to install one themselves on their machine.

A.3.3 Yakima Valley Memorial Hospital, December 1985

In this accident, as in the Kennestone overdose, machine malfunction was not acknowledged until after later accidents were understood.

The Therac-25 at Yakima, Washington, had been modified by AECL in September 1985 in response to the overdose at Hamilton. During December 1985, a woman treated with the Therac-25 developed erythema (excessive reddening of the skin) in a parallel striped pattern on her right hip. Despite this, she continued to be treated by the Therac-25, as the cause of her reaction was not determined to be abnormal until January 1986. On January 6, her treatments were completed.

The staff monitored the skin reaction closely and attempted to find possible causes. The open slots in the blocking trays in the Therac-25 could have produced such a striped pattern, but by the time the skin reaction was determined to be abnormal, the blocking trays had been discarded, so the blocking arrangement and tray striping orientation could not be reproduced. A reaction to chemotherapy was ruled out because that should have produced reactions at the other treatment sites and would not have produced stripes. When the doctors discovered that the woman slept with a heating pad, they thought maybe the burn pattern had been caused by the parallel wires that deliver the heat in such pads. The staff X-rayed the heating pad, but discovered that the wire pattern did not correspond to the erythema pattern on the patient's hip.

The hospital staff sent a letter to AECL on January 31, and they also spoke on the phone with the AECL technical support supervisor. On February 24, the AECL technical support supervisor sent a written response to the director of radiation therapy at Yakima saying, "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 letter goes on to support this opinion by listing two pages of technical reasons why an overdose by the Therac-25 was impossible, along with the additional argument that there have "apparently been no other instances of similar damage to this or other patients." The letter ends, "In

closing, I wish to advise that this matter has been brought to the attention of our Hazards Committee as is normal practice."

The hospital staff eventually ascribed the patient's skin reaction to "cause unknown." In a report written on this first Yakima incident after another Yakima overdose a year later, the medical physicist involved wrote:

At that time, we did not believe that [the patient] was overdosed because the manufacturer had installed additional hardware and software safety devices to the accelerator.

In a letter from the manufacturer dated 16-Sep-85, it is stated that "Analysis of the hazard rate resulting from these modifications indicates an improvement of at least five orders of magnitude"! With such an improvement in safety (10,000,000%) we did not believe that there could have been any accelerator malfunction. These modifications to the accelerator were completed on 5,6-Sep-85.

Even with fairly sophisticated physics support, the hospital staff, as users, did not have the ability to investigate the possibility of machine malfunction further. They were not aware of any other incidents and, in fact, were told that there had been none, so there was no reason for them to pursue the matter. No further investigation of this incident was done by the manufacturer or by any government agencies (who did not know about it).

About a year later (February 1987), after the second Yakima overdose led the hospital staff to suspect that this first injury had been due to a Therac-25 fault, the staff investigated and found that the first overdose victim had a chronic skin ulcer, tissue necrosis (death) under the skin, and was in continual pain. The damage was surgically repaired, skin grafts were made, and the symptoms relieved. The patient is alive today with minor disability and some scarring related to the overdose. The hospital staff concluded that the dose accidentally delivered in the first accident must have been much lower than in the second, as the reaction was significantly less intense and necrosis did not develop until six or eight months after exposure. Some other factors related to the place on the body where the overdose occurred also kept her from having more significant problems.

A.3.4 East Texas Cancer Center, March 1986

More is known about the Tyler, Texas, accidents than the others because of the diligence of the Tyler hospital physicist, Fritz Hager, without whose efforts the understanding of the software problems may have been delayed even further.

The Therac-25 had been at the East Texas Cancer Center (ETCC) for two years before the first serious accident, and more than 500 patients had been treated. On March 21, 1986, a male patient came into ETCC for his ninth treatment on the Therac-25, one of a series prescribed as followup to the removal of a tumor from his back.

This treatment was to be a 22 MeV electron beam treatment of 180 rads on

the upper back and a little to the left of his spine, for a total of 6,000 rads over six and a half weeks. He was taken into the treatment room and placed face down on the treatment table. The operator then left the treatment room, closed the door, and sat at the control terminal.

The operator had held this job for some time, and her typing efficiency had increased with experience. She could quickly enter prescription data and change it conveniently with the Therac's editing features. She entered the patient's prescription data quickly, then noticed that she had typed "x" (for X-ray) when she had intended "e" (for electron) mode. This was a common mistake as most of the treatments involved X-rays, and she had gotten used to typing this. The mistake was easy to fix; she merely used the $\text{\textcircled{1}}$ key to edit the mode entry.

Because the other parameters she had entered were correct, she hit the return key several times and left their values unchanged. She reached the bottom of the screen, where it was indicated that the parameters had been VERIFIED and the terminal displayed BEAM READY, as expected. She hit the one-key command, $\text{\textcircled{B}}$ for *beam on*, to begin the treatment. After a moment, the machine shut down and the console displayed the message MALFUNCTION 54. The machine also displayed a TREATMENT PAUSE, indicating a problem of low priority. The sheet on the side of the machine explained that this malfunction was a "dose input 2" error. The ETCC did not have any other information available in its instruction manual or other Therac-25 documentation to explain the meaning of MALFUNCTION 54. An AECL technician later testified that "dose input 2" meant that a dose had been delivered that was either too high or too low. The messages had been expected to be used only during internal company development.

The machine showed a substantial underdose on its dose monitor display—6 monitor units delivered whereas the operator had requested 202 monitor units. She was accustomed to the quirks of the machine, which would frequently stop or delay treatment; in the past, the only consequences had been inconvenience. She immediately took the normal action when the machine merely paused, which was to hit the $\text{\textcircled{P}}$ key to proceed with the treatment. The machine promptly shut down with the same MALFUNCTION 54 error and the same underdose shown by the dosimetry.

The operator was isolated from the patient, since the machine apparatus was inside a shielded room of its own. The only way that the operator could be alerted to patient difficulty was through audio and video monitors. On this day, the video display was unplugged and the audio monitor was broken.

After the first attempt to treat him, the patient said that he felt as if he had received an electric shock or that someone had poured hot coffee on his back: He felt a thump and heat and heard a buzzing sound from the equipment. Since this was his ninth treatment, he knew that this was not normal. He began to get up from the treatment table to go for help. It was at this moment that the operator hit the $\text{\textcircled{P}}$ key to proceed with the treatment. The patient said that he felt like his arm was being shocked by electricity and that his hand was leaving his body. He went to the treatment room door and pounded on it. The operator was

shocked and immediately opened the door for him. He appeared visibly shaken and upset.

The patient was immediately examined by a physician, who observed intense reddening of the treatment area, but suspected nothing more serious than electric shock. The patient was discharged and sent home with instructions to return if he suffered any further reactions. The hospital physicist was called in, and he found the machine calibration within specifications. The meaning of the malfunction message was not understood. The machine was then used to treat patients for the rest of the day.

In actuality, but unknown to anyone at that time, the patient had received a massive overdose, concentrated in the center of the treatment location. After-the-fact simulations of the accident revealed possible doses of 16,500 to 25,000 rads in less than 1 second over an area of about 1 cm.

Over the weeks following the accident, the patient continued to have pain in his neck and shoulder. He lost the function of his left arm and had periodic bouts of nausea and vomiting. He was eventually hospitalized for radiation-induced myelitis of the cervical cord causing paralysis of his left arm and both legs, left vocal cord paralysis (which left him unable to speak), neurogenic bowel and bladder, and paralysis of the left diaphragm. He also had a lesion on his left lung and recurrent herpes simplex skin infections. He died from complications of the overdose five months after the accident.

User and Manufacturer Response

The Therac-25 was shut down for testing the day after this accident. One local AECL engineer and one from the home office in Canada came to ETCC to investigate. They spent a day running the machine through tests, but could not reproduce a Malfunction 54. The AECL engineer from the home office reportedly explained that it was not possible for the Therac-25 to overdose a patient. The ETCC physicist claims that he asked AECL at this time if there were any other reports of radiation overexposure and that AECL personnel (including the quality assurance manager) told him that AECL knew of no accidents involving radiation overexposure by the Therac-25. This seems odd since AECL was surely at least aware of the Hamilton accident that had occurred seven months before and the Yakima accident, and, even by their account, learned of the Georgia lawsuit around this time (which had been filed four months earlier). The AECL engineers then suggested that an electrical problem might have caused the burn.

The electric shock theory was checked out thoroughly by an independent engineering firm. The final report indicated that there was no electrical grounding problem in the machine, and it did not appear capable of giving a patient an electrical shock. The ETCC physicist checked the calibration of the Therac-25 and found it to be satisfactory. He put the machine back into service on April 7, 1986, convinced that it was performing properly.

A.3.5 East Texas Cancer Center, April 1986

Three weeks later, on April 11, 1986, another male patient was scheduled to receive an electron treatment at ETCC for a skin cancer on the side of his face. The prescription was for 10 MeV. The same technician who had treated the first Tyler accident victim prepared this patient for treatment. Much of what follows is from the operator's deposition.

As with her former patient, she entered the prescription data and then noticed an error in the mode. Again she used the edit \oplus key to change the mode from X-ray to electron. After she finished editing, she pressed the RETURN key several times to place the cursor on the bottom of the screen. She saw the BEAM READY message displayed and turned the beam on.

Within a few seconds the machine shut down, making a loud noise audible via the (now working) intercom. The display showed MALFUNCTION 54 again. The operator rushed into the treatment room, hearing her patient moaning for help. He began to remove the tape that had held his head in position and said something was wrong. She asked him what he felt, and he replied, "fire" on the side of his face. She immediately went to the hospital physicist and told him that another patient appeared to have been burned. Asked by the physicist to describe what had happened, the patient explained that something had hit him on the side of the face, he saw a flash of light, and he heard a sizzling sound reminiscent of frying eggs. He was very agitated and asked, "What happened to me, what happened to me?"

This patient died from the overdose on May 1, 1986, three weeks after the accident. He had disorientation, which progressed to coma, fever to 104°F, and neurological damage. An autopsy showed an acute high-dose radiation injury to the right temporal lobe of the brain and the brain stem.

User and Manufacturer Response

After this second Tyler accident, the ETCC physicist immediately took the machine out of service and called AECL to alert them to this second apparent overexposure. The physicist then began a careful investigation of his own. He worked with the operator, who remembered exactly what she had done on this occasion. After a great deal of effort, they were eventually able to elicit the MALFUNCTION 54 message. They determined that data entry speed during editing was the key factor in producing the error condition: If the prescription data was edited at a fast pace (as is natural for someone who has repeated the procedure a large number of times), the overdose occurred. It took some practice before the physicist could repeat the procedure rapidly enough to elicit the MALFUNCTION 54 message at will.

The next day, an engineer from AECL called and said that he could not reproduce the error. After the ETCC physicist explained that the procedure had to be performed quite rapidly, AECL could finally produce a similar malfunction

on its own machine. Two days after the accident, AECL said it had measured the dosage (at the center of the field) to be 25,000 rads. An AECL engineer explained that the frying sound heard by the patients was the ion chambers being saturated.

In one lawsuit that resulted from the Tyler accidents, the AECL quality control manager testified that a "cursor up" problem had been found in the service (maintenance) mode at other clinics in February or March of 1985 and also in the summer of 1985. Both times, AECL thought that the software problems had been fixed. There is no way to determine whether there is any relationship between these problems and the Tyler accidents.

Related Therac-20 Problems

The software for both the Therac-25 and Therac-20 "evolved" from the Therac-6 software. Additional functions had to be added because the Therac-20 (and Therac-25) operate in both X-ray and electron mode, while the Therac-6 has only X-ray mode. CGR modified the software for the Therac-20 to handle the dual modes. When the Therac-25 development began, AECL engineers adapted the software from the Therac-6, but they also borrowed software routines from the Therac-20 to handle electron mode, which was allowed under their cooperative agreements.

After the second Tyler, Texas, accident, a physicist at the University of Chicago Joint Center for Radiation Therapy heard about the Therac-25 software problem and decided to find out whether the same thing could happen with the Therac-20. At first, the physicist was unable to reproduce the error on his machine, but two months later he found the link.

The Therac-20 at the University of Chicago is used to teach students in a radiation therapy school conducted by the center. The center's physicist, Frank Borger, noticed that whenever a new class of students started using the Therac-20, fuses and breakers on the machine tripped, shutting down the unit. These failures, which had been occurring ever since the school had acquired the machine, might happen three times a week while new students operated the machine and then disappear for months. Borger determined that new students make many different types of mistakes and use "creative methods" of editing parameters on the console. Through experimentation, he found that certain editing sequences correlated with blown fuses and determined that the same computer bug (as in the Therac-25 software) was responsible. The physicist notified the FDA, which notified Therac-20 users [148].

The software error is just a nuisance on the Therac-20 because this machine has independent hardware protective circuits for monitoring the electron beam scanning. The protective circuits do not allow the beam to turn on, so there is no danger of radiation exposure to a patient. While the Therac-20 relies on mechanical interlocks for monitoring the machine, the Therac-25 relies largely on software.

The Software “Bug”

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 some 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 demonstrating the overall software design flaws. First the software design is described and then the errors believed to be involved in the Tyler accidents and perhaps others.

Therac-25 Software Development and Design. AECL 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 hardcopy 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.

There are two basic operational modes: treatment mode and service mode. Treatment mode controls the normal treatment process. In service mode, the unit can be operated with some of the operational and treatment interlocks bypassed, and additional operational commands and characteristics may be selected. Service mode is entered only through the use of a password at the service keyboard.

The manufacturer describes the Therac-25 software as having a stand-alone, real-time treatment operating system. The system does not use 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. Cycles are allocated to the critical and noncritical tasks using a preemptive scheduler.

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 routine
- A scanning interrupt service routine

- 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 sequencing 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 as 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 (such as keyboard processor or 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 displays appropriate messages on the CRT display. It decodes some information and checks the setup verification.

Noncritical tasks include

- Checksum 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 communication interface between the other software and the operator.
- Treatment console screen processor (run periodically). This task lays out appropriate record formats for either CRT 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.
- Hand 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 AECL documentation on the modifications that the software allows concurrent access to shared memory, that there is no real synchronization aside from data that are 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.

Specific Design Errors. The following explanation of the specific software problems found at this time is taken from the description AECL provided to the FDA, but clarified somewhat. The description leaves some unanswered questions, but it is the best that can be done with the information available.

The treatment monitor task (Treat) controls the various phases of treatment by executing its eight subroutines. The treatment phase indicator variable (Tphase) is used to determine which subroutine should be executed (Figure A.3). Following the execution of a particular subroutine, Treat reschedules itself.

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 Complete flag) to determine whether the prescription data has been entered. The keyboard handler recognizes the completion of data entry and changes the Data Entry Complete variable to denote this. Once this variable is set, the Datent subroutine detects the variable’s change in status and changes the value of Tphase from 1 (Datent) 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 Complete variable has not been set, Datent leaves the value of Tphase unchanged and exits back to Treat’s mainline. Treat will then reschedule itself, essentially rescheduling the Datent subroutine.

The command line at the lower right-hand corner of the screen (see Figure A.2) is the cursor’s normal position when the operator has completed all the necessary changes to the prescription. Prescription editing is signified by moving the cursor off the command line. As the program was originally designed, the Data Entry Complete variable by itself is not sufficient because 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 variable (MEOS). 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 and 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

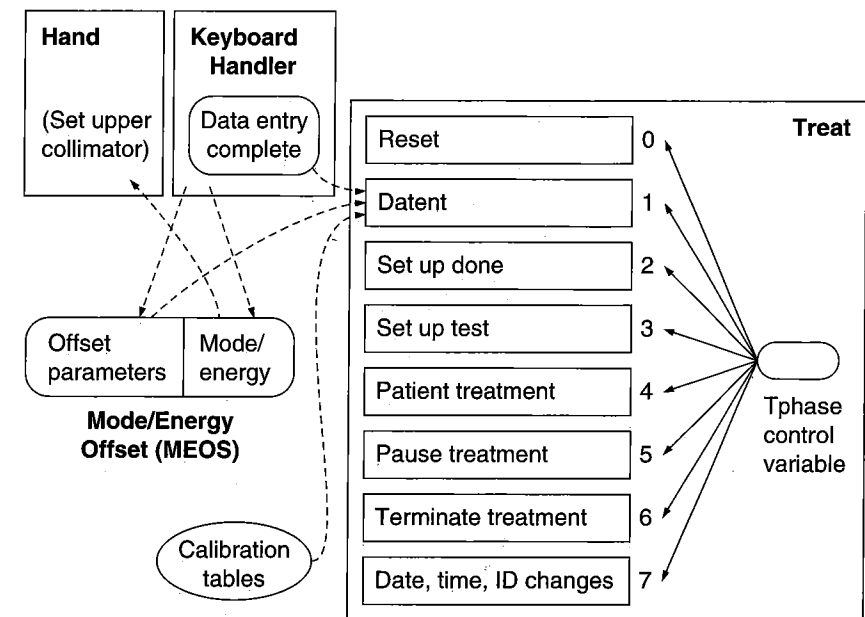


FIGURE A.3
Tasks and subroutines in the code blamed for the Tyler accidents.

energy except when the photon mode is selected, 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 Complete flag before the operator changes the data in MEOS, Datent will not detect the changes because it has already exited and will not be reentered again. The upper collimator (turntable), 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. The software appears to contain no checks to detect such an incompatibility.

The first thing Datent does when it is entered is to check whether the keyboard handler has set the mode and energy 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-to-analog converter during the next clock cycle. Once the parameters are all set, Datent calls the subroutine Magnet, which sets the bending magnets. The following shows a simplified pseudocode description of relevant parts of the software:

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

```

Setting the bending magnets takes about eight 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 the 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. Because Ptime clears it during its first execu-

tion, 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 and energy, went to the command line, then moved the cursor up to change the mode or energy and returned to the command line all within eight seconds. Because the magnet setting takes about eight 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 is not the only problem. Upon exit from the Magnet subroutine, the data entry subroutine (Datent) checks the Data Entry Complete 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 Complete 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, AECL 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.

The Government and User Response

The FDA does not approve each new medical device on the market: All medical devices go through a classification process that determines the level of FDA approval necessary. Medical accelerators follow a procedure called pre-market notification before commercial distribution. In this process, the firm must establish that the product is substantially equivalent in safety and effectiveness to a product already on the market. If that cannot be done to the FDA's satisfaction, a pre-market approval is required. For the Therac-25, the FDA required only a pre-market notification. After the Therac-25 accidents, new procedures for approval of software-controlled devices were adopted.

The agency is basically reactive to problems and requires manufacturers to report serious ones. Once a problem is identified in a radiation-emitting product, the FDA is responsible for approving the corrective action plan (CAP).

The first reports of the Tyler incidents came to the FDA from the State of Texas Health Department, and this triggered FDA action. The FDA investigation was well under way when AECL produced a medical device report to discuss the details of the radiation overexposures at Tyler. The FDA declared the Therac-25 defective under the Radiation Control for Health and Safety Act and ordered the firm to notify all purchasers, investigate the problem, determine a solution, and submit a corrective action plan for FDA approval.

The final CAP consisted of more than twenty changes to the system hardware and software, plus modifications to the system documentation and manuals. Some of these changes were unrelated to the specific accidents, but were improvements to the general safety of the machine. The full CAP implementation, including an extensive safety analysis, was not complete until more than two years after the Tyler accidents.

AECL made their accident report to the FDA on April 15, 1986. On that same date, AECL sent out a letter to each Therac user recommending a temporary "fix" to the machine that would allow continued clinical use. The letter (shown in its complete form) stated:

SUBJECT: CHANGE IN OPERATING PROCEDURES FOR THE THERAC 25 LINEAR ACCELERATOR

Effective immediately, and until further notice, the key used for moving the cursor back through the prescription sequence (i.e., cursor 'UP' inscribed with an upward pointing arrow) must not be used for editing or any other purpose.

To avoid accidental use of this key, the key cap must be removed and the switch contacts fixed in the open position with electrical tape or other insulating material. For assistance with the latter you should contact your local AECL service representative.

Disabling this key means that if any prescription data entered is incorrect then a 'R' reset command must be used and the whole prescription re-entered.

For those users of the Multiport option it also means that editing of dose rate, dose and time will not be possible between ports.

On May 2, 1986, the FDA declared the Therac defective, demanded a CAP, and required renotification of all the Therac customers. In the letter from the FDA to AECL, the Director of Compliance, Center for Devices and Radiological Health, wrote:

We have reviewed [AECL's] April 15 letter to purchasers and have concluded that it does not satisfy the requirements for notification to purchasers of a defect in an electronic product. Specifically, it does not describe the defect nor the hazards associated with it. The letter does not provide any reason for disabling the cursor key and the tone is not commensurate with the urgency for doing so. In fact, the letter implies the inconvenience to operators outweighs the need to disable the key. We request that you immediately renotify purchasers.

AECL promptly made a new notice to users and also requested an extension to produce a CAP. The FDA granted this request.

About this time, the Therac-25 users created a user's group and held their first meeting at the annual conference of the American Association of Physicists in Medicine. At the meeting, users discussed the Tyler accident and heard an

AECL representative present the company's plans for responding to it. AECL promised to send a letter to all users detailing the CAP.

Several users described additional hardware safety features that they had added to their own machines to provide additional protection. An interlock (that checked gun current values), which the Vancouver clinic had previously added to their Therac-25, was labeled as redundant by AECL; the users disagreed. There were further discussions of poor design and other problems that caused a 10- to 30-percent underdosing in both modes.

The meeting notes said

There was a general complaint by all users present about the lack of information propagation. The users were not happy about receiving incomplete information. The AECL representative countered by stating that AECL does not wish to spread rumors and that AECL has no policy to 'keep things quiet.' The consensus among the users was that an improvement was necessary.

After the first user's group meeting, there were two user's group newsletters. The first, dated fall 1986, contained letters from Tim Still, the Kennestone physicist, who complained about what he considered to be eight major problems he had experienced with the Therac-25. These problems included poor screen-refresh subroutines that leave trash and erroneous information on the operator console and some tape-loading problems upon startup that he discovered involved the use of "phantom tables" to trigger the interlock system in the event of a load failure instead of using a checksum. He asked the question, "Is programming safety relying too much on the software interlock routines?" The second user's group newsletter, in December 1986, further discussed the implications of the phantom table problem.

AECL produced its first CAP on June 13, 1986. The FDA asked for changes and additional information about the software, including a software test plan. AECL responded on September 26 with several documents describing the software and its modifications but no test plan. They explained how the Therac-25 software evolved from the Therac-6 software and stated that "no single test plan and report exists for the software since both hardware and software were tested and exercised separately and together over many years." AECL concluded that the current CAP improved "machine safety by many orders of magnitude and virtually eliminates the possibility of lethal doses as delivered in the Tyler incident."

An FDA internal memo dated October 20 commented on these AECL submissions, raising several concerns:

Unfortunately, the AECL response also seems to point out an apparent lack of documentation on software specifications and a software test plan.

... concerns include the question of previous knowledge of problems by

AECL, the apparent paucity of software quality assurance at the manufacturing facility, and possible warnings and information dissemination to others of the generic type problems.

. . . As mentioned in my first review, there is some confusion on whether the manufacturer should have been aware of the software problems prior to the ARO's [Accidental Radiation Overdoses] in Texas. AECL had received official notification of a law suit in November 1985 from a patient claiming accidental over-exposure from a Therac-25 in Marietta, Georgia. . . . If knowledge of these software deficiencies were known beforehand, what would be the FDA's posture in this case?

. . . The materials submitted by the manufacturer have not been in sufficient detail and clarity to ensure an adequate software quality assurance program currently exists. For example, a response has not been provided with respect to the software part of the CAP to the CDRH's [FDA Center for Devices and Radiological Health] request for documentation on the revised requirements and specifications for the new software. In addition, an analysis has not been provided, as requested, on the interaction with other portions of the software to demonstrate the corrected software does not adversely affect other software functions.

The July 23 letter from the CDRH requested a documented test plan including several specific pieces of information identified in the letter. This request has been ignored up to this point by the manufacturer. Considering the ramifications of the current software problem, changes in software QA attitudes are needed at AECL.

AECL also planned to retain the malfunction codes, but the FDA required better warnings for the operators. Furthermore, AECL had not planned on any quality assurance testing to ensure exact copying of software, but the FDA insisted on it. The FDA further requested assurances that rigorous testing would become a standard part of AECL's software modification procedures.

We also expressed our concern that you did not intend to perform the protocol to future modifications to software. We believe that the rigorous testing must be performed each time a modification is made in order to ensure the modification does not adversely affect the safety of the system.

AECL was also asked to draw up an installation test plan to ensure that both hardware and software changes perform as designed when installed.

AECL submitted CAP Revision 2 and supporting documentation on December 22, 1986. They changed the CAP to have dose malfunctions suspend treatment and included a plan for meaningful error messages and highlighted dose error messages. They also expanded their diagrams of software modifications and expanded their test plan to cover hardware and software.

A.3.6 Yakima Valley Memorial Hospital, January 1987

On Saturday, January 17, 1987, the second patient of the day was to be treated for a carcinoma. This patient was to receive two film verification exposures of 4 and 3 rads plus a 79-rad photon treatment (for a total exposure of 86 rads.)

Film was placed under the patient and 4 rads were administered. After the machine paused to open the collimator jaws further, the second exposure of 3 rads was administered. The machine paused again.

The operator entered the treatment room to remove the film and verify the patient's precise position. He used the hand control in the treatment room to rotate the turntable to the field light position, which allowed him to check the alignment of the machine with respect to the patient's body in order to verify proper beam position. He then either pressed the *set* button on the hand control or left the room and typed a SET command at the console to return the turntable to the proper position for treatment; there is some confusion as to exactly what transpired. When he left the room, he forgot to remove the film from underneath the patient. The console displayed "beam ready," and the operator hit the **B** key to turn the beam on.

The beam came on, but the console displayed no dose or dose rate. After five or six seconds, the unit shut down with a pause and displayed a message. The message "may have disappeared quickly"; the operator was unclear on this point. However, since the machine merely paused, he was able to push the **P** key to proceed with treatment.

The machine paused again, this time displaying FLATNESS on the reason line. The operator heard the patient say something over the intercom, but could not understand him. He went into the room to speak with the patient, who reported "feeling a burning sensation" in the chest. The console displayed only the total dose of the two film exposures (7 rads) and nothing more.

Later in the day, the patient developed a skin burn over the entire treatment area. Four days later, the redness developed a striped pattern matching the slots in the blocking tray. The striped pattern was similar to the burn a year earlier at this same hospital, which had first been ascribed to a heating pad and later officially labeled by the hospital as "cause unknown."

AECL began an investigation, and users were told to confirm the turntable position visually before turning on the beam. All tests run by the AECL engineers indicated that the machine was working perfectly. From the information that had been gathered to that point, it was suspected that the electron beam had come on when the turntable was in the field light position. But the investigators could not reproduce the fault condition.

On the following Thursday, AECL sent in an engineer from Ottawa to investigate. The hospital physicist had, in the meantime, run some tests himself. He placed a film in the Therac's beam and then ran two exposures of X-ray parameters with the turntable in field light position. The film appeared to match the film that was left (by mistake) under the patient during the accident.

After a week of checking the hardware, AECL determined that the “incorrect machine operation was probably not caused by hardware alone.” After checking the software, AECL engineers discovered a flaw (described below) that could explain the erroneous behavior. The coding problems explaining this accident are completely different from those associated with the Tyler accidents.

Preliminary dose measurements by AECL indicated that the dose delivered under these conditions—that is, when the turntable is in the field light position—is on the order of 4,000 to 5,000 rads. After two attempts, the patient could have received 8,000 to 10,000 instead of the 86 rads prescribed. AECL again called users on January 26 (nine days after the accident) and gave them detailed instructions on how to avoid this problem. In an FDA internal report on the accident, the AECL quality assurance manager investigating the problem is quoted as saying that the software and hardware changes to be retrofitted following the Tyler accident nine months earlier (but which had not yet been installed) would have prevented the Yakima accident.

The patient died in April from complications related to the overdose. He had a terminal form of cancer, but a lawsuit was initiated by his survivors alleging that he died sooner than he would have and endured unnecessary pain and suffering due to the radiation overdose. The suit, like all the others, was settled out of court.

The Yakima Software “Bug”

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, unknown race conditions or errors could have been responsible for them. 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 AECL 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 involved problems encountered later in the logic after the treatment monitor Treat reaches Set Up Test.

The Therac-25’s field light feature allows very precise positioning of the patient for treatment. The operator can control the machine right at the treatment

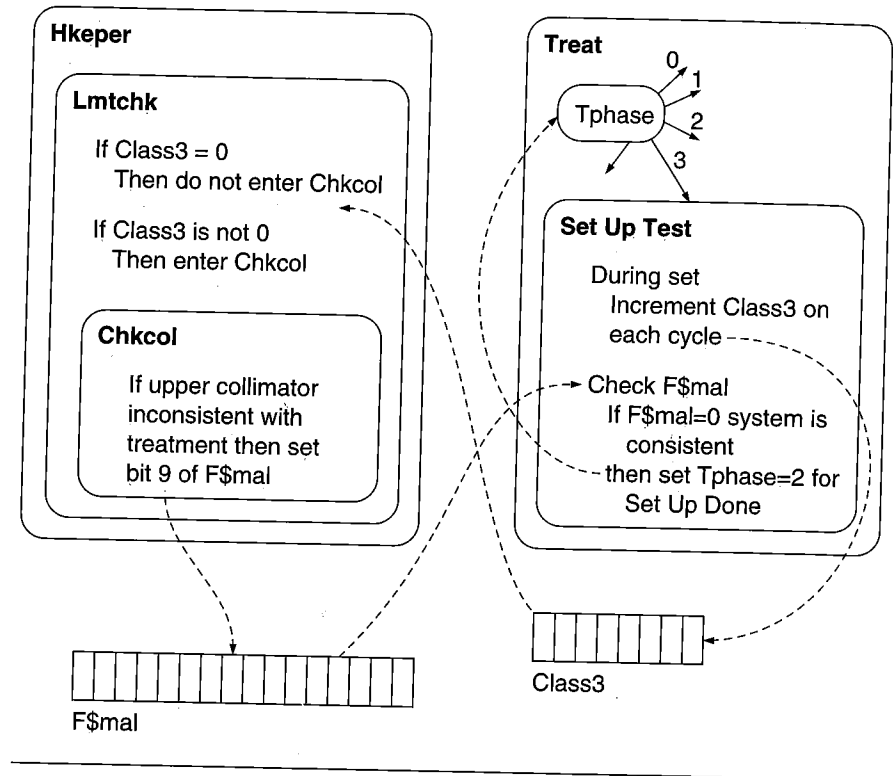


FIGURE A.4
The Yakima software flaw.

site using a small hand control that offers 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 patient setup in the treatment room, and all relevant parameters now VERIFY. The console displays a message to 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.

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 (Figure A.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 software does not inhibit the beam.

After setting the Class3 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-to-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 Treat to decide whether to reschedule itself or to proceed to Set Up Done.

During machine setup, Set Up Test will be executed several hundred times because 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 one byte, it can only contain a maximum value of 255 decimal. Thus, on every 256th pass through the Set Up Test code, the variable will overflow and have a zero value. That means that on every 256th pass through Set Up Test, the upper collimator will not be checked and an 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 that the upper collimator was still in the 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.

The technical “fix” implemented for this particular software flaw is described by AECL 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.

Manufacturer, Government, and User Response

On February 3, 1987, after interaction with the FDA and others, including the user’s group, AECL announced to its customers

1. A new software release to correct both the Tyler and Yakima software problems

2. A hardware single-pulse shutdown circuit
3. A turntable potentiometer to independently monitor turntable position
4. A hardware turntable interlock circuit

The second item, a hardware single-pulse shutdown circuit, essentially acts as a hardware interlock to prevent overdosing by detecting an unsafe level of radiation and halting beam output after one pulse of high energy and current. This interlock effectively provides an independent way to protect against a wide range of potential hardware failures and software errors. The third item, a turntable potentiometer, was the safety device recommended by several groups after the Hamilton accident.

After the second Yakima accident, the FDA became concerned that the use of the Therac-25 during the CAP process, even with AECL’s interim operating instructions, involved too much risk to patients. The FDA concluded that the accidents demonstrated that the software alone could not be relied upon to assure safe operation of the machine. In a February 18, 1987, internal FDA memorandum, the Director of the Division of Radiological Products wrote:

It is impossible for CDRH to find all potential failure modes and conditions of the software. AECL has indicated the “simple software fix” will correct the turntable position problem displayed at Yakima. We have not yet had the opportunity to evaluate that modification. Even if it does, based upon past history, I am not convinced that there are not other software glitches that could result in serious injury.

... We are in the position of saying that the proposed CAP can reasonably be expected to correct the deficiencies for which they were developed (Tyler). We cannot say that we are reasonable [sic] confident about the safety of the entire system to prevent or minimize exposure from other fault conditions.

On February 6, 1987, Ed Miller of the FDA called Pavel Dvorak of Canada’s Health and Welfare to advise him that the FDA would recommend that all Therac-25s be shut down until permanent modifications could be made. According to Miller’s notes on the phone call, Dvorak agreed and indicated that Health and Welfare would coordinate their actions with the FDA.

AECL responded on April 13 with an update on the Therac-25 CAP status and a schedule of the nine action items pressed by the users at a user’s group meeting in March. This unique and highly productive meeting provided an unusual opportunity to involve the users in the CAP evaluation process. It brought together all concerned parties in one place and at one time so that a course of action could be decided upon and approved as quickly as possible. The attendees included representatives from

- The manufacturer (AECL)
- All users, including their technical and legal staffs
- The FDA and the Canadian Bureau of Radiation and Medical Devices

- the Canadian Atomic Energy Control Board
- the Province of Ontario
- the Radiation Regulations Committee of the Canadian Association of Physicists

According to Gordon Symonds, from the Canadian BRMD, this meeting was very important to the resolution of the problems, since the regulators, users, and manufacturer arrived at a consensus in one day.

At this second user's meeting, the participants carefully reviewed all the six known major Therac-25 accidents to that date and discussed the elements of the CAP along with possible additional modifications. They came up with a prioritized list of modifications they wanted included in the CAP and expressed concerns about the lack of independent evaluation of the software and the lack of a hardcopy audit trail to assist in diagnosing faults.

The AECL representative, who was the quality assurance manager, responded that tests had been done on the CAP changes, but that the tests were not documented and that independent evaluation of the software "might not be possible." He claimed that two outside experts had reviewed the software, but he could not provide their names. In response to user requests for a hardcopy audit trail and access to source code, he explained that memory limitations would not permit including such options and that source code would not be made available to users.

On May 1, AECL issued CAP Revision 4 as a result of the FDA comments and the user's meeting input. The FDA response on May 26 approved the CAP subject to submission of the final test plan results and an independent safety analysis, distribution of the draft revised manual to customers, and completion of the CAP by June 30, 1987. The FDA concluded by rating this a Class I recall: a recall in which there is a reasonable probability that the use of, or exposure to, a violative product will cause serious adverse health consequences or death [35].

AECL sent more supporting documentation to the FDA on June 5, 1987, including the CAP test plan, a draft operator's manual, and the draft of the new safety analysis. This time the analysis included the software in the fault trees but used a "generic failure rate" of 10^{-4} for software events. This number was justified as being based on the historical performance of the Therac-25 software. The final report on the safety analysis states that many of the fault trees had a computer malfunction as a causative event, and the outcome for quantification was therefore dependent on the failure rate chosen for the software. Assuming that all software errors are equally likely seems rather strange.

A close inspection of the code was also conducted during this safety analysis to "obtain more information on which to base decisions." An outside consultant performed the inspection, which included a detailed examination of the implementation of each function, a search for coding errors, and a qualitative assessment of the software's reliability. No information is provided in the final safety report about whether any particular methodology or tools were used in the software inspection or whether someone just read the code looking for errors.

AECL planned a fifth revision of the CAP to include the testing and final safety analysis results. Referring to the test plan at this, the final stage of the CAP process, an FDA reviewer said,

Amazingly, the test data presented to show that the software changes to handle the edit problems in the Therac-25 are appropriate prove the exact opposite result. A review of the data table in the test results indicates that the final beam type and energy (edit change) has no effect on the initial beam type and energy. I can only assume that either the fix is not right or the data was entered incorrectly. The manufacturer should be admonished for this error. Where is the QC [Quality Control] review for the test program? AECL must: (1) clarify this situation, (2) change the test protocol to prevent this type of error from occurring, and (3) set up appropriate QC control on data review.

A further FDA memo indicated:

[The AECL quality assurance manager] could not give an explanation and will check into the circumstances. He subsequently called back and verified that the technician completed the form incorrectly. Correct operation was witnessed by himself and others. They will repeat and send us the correct data sheet.

At the American Association of Physicists in Medicine meeting in July 1987, a third user's meeting was held. The AECL representative described the status of the latest CAP and explained that the FDA had given verbal approval and that he expected full implementation by the end of August 1987. He went on to review and comment on the prioritized concerns of the last meeting. Three of the user-requested hardware changes had been included in the CAP. Changes to tape load error messages and checksums on the load data would wait until after the CAP was done. Software documentation was described as a lower priority task that needed definition and would not be available to the FDA in any form for over a year.

On July 6, 1987, AECL sent a letter to all users to update them on the FDA's verbal approval of the CAP and to delineate how AECL would proceed. Finally, on July 21, 1987, AECL issued the final and fifth CAP revision. The major features of the final CAP are these:

- All interruptions related to the dosimetry system will go to a treatment suspend, not a treatment pause. Operators will not be allowed to restart the machine without reentering all parameters.
- A software single-pulse shutdown will be added.
- An independent hardware single-pulse shutdown will be added.
- Monitoring logic for turntable position will be improved to ensure that the turntable is in one of the three legal positions.
- A potentiometer will be added to the turntable. The output is used to monitor exact turntable location and provide a visible position signal to the operator.

- Interlocking with the 270-degree bending magnet will be added to ensure that the target and beam flattener are in position if the X-ray mode is selected.
- Beam-on will be prevented if the turntable is in the field light or any intermediate position.
- Cryptic malfunction messages will be replaced with meaningful messages and highlighted dose-rate messages.
- Editing keys will be limited to *cursor up*, *backspace*, and *return*. All other keys will be inoperative.
- A motion-enable footswitch (a type of deadman switch) will be added. The operator will be required to hold this switch closed during movement of certain parts of the machine to prevent unwanted motions when the operator is not in control.
- Twenty-three other changes will be made to the software to improve its operation, including disabling of unused keys, changing the operation of the *set* and *reset* commands, preventing copying of the control program on site, changing the way various detected hardware faults are handled, eliminating errors in the software that were detected during the review process, adding several additional software interlocks, disallowing changes in the service mode while a treatment is in progress, and adding meaningful error messages.
- The known software problems associated with the Tyler and Yakima accidents will be fixed.
- The manuals will be fixed to reflect the changes.

Figure A.5 shows a typical Therac-25 installation after the CAP changes were made.

Ed Miller, the director of the Division of Standards Enforcement, Center for Devices and Radiological Health at the FDA, wrote in 1987:

FDA has performed extensive review of the Therac-25 software and hardware safety systems. We cannot say with absolute certainty that all software problems that might result in improper dose have been found and eliminated. However, we are confident that the hardware and software safety features recently added will prevent future catastrophic consequences of failure.

No Therac-25 accidents have been reported since the final corrective action plan was implemented.

A.4 Causal Factors

Many lessons can be learned from this series of accidents. A few are considered here.




FIGURE A.5
A typical Therac-25 facility after the final CAP.

Overconfidence in Software. A common mistake in engineering, in this case and in many others, is to put too much confidence in software. There seems to be a feeling among nonsoftware professionals that software will not or cannot fail, which leads to complacency and overreliance on computer functions.

A related tendency among engineers is to ignore software. The first safety analysis on the Therac-25 did not include software—although nearly full responsibility for safety rested on it. When problems started occurring, it was assumed that hardware had caused them, and the investigation looked only at the hardware.

Confusing Reliability with Safety. This software was highly reliable. It worked tens of thousands of times before overdosing anyone, and occurrences of erroneous behavior were few and far between. AECL assumed that their software was safe because it was reliable, and this led to complacency.

Lack of Defensive Design. The software did not contain self-checks or other error-detection and error-handling features that would have detected the inconsistencies and coding errors. Audit trails were limited because of a lack of memory. However, today larger memories are available and audit trails and other design techniques must be given high priority in making tradeoff decisions.

Patient reactions were the only real indications of the seriousness of the problems with the Therac-25; there were no independent checks that the machine and its software were operating correctly. Such verification cannot be assigned to operators without providing them with some means of detecting errors: The Therac-25 software “lied” to the operators, and the machine itself was not capable of detecting that a massive overdose had occurred. The ion chambers on the Therac-25 could not handle the high density of ionization from the unscanned electron beam at high beam current; they thus became saturated and gave an indication of a low dosage. Engineers need to design for the worst case.

Failure to Eliminate Root Causes. One of the lessons to be learned from the Therac-25 experiences is that focusing on particular software design errors is not the way to make a system safe. Virtually all complex software can be made to behave in an unexpected fashion under some conditions: There will always be another software bug. Just as engineers would not rely on a design with a hardware single point of failure that could lead to catastrophe, they should not do so if that single point of failure is software.

The Therac-20 contained the same software error implicated in the Tyler deaths, but this machine included hardware interlocks that mitigated the consequences of the error. Protection against software errors can and should be built into both the system and the software itself. We cannot eliminate all software errors, but we can often protect against their worst effects, and we can recognize their likelihood in our decision making.

One of the serious mistakes that led to the multiple Therac-25 accidents was the tendency to believe that the cause of an accident had been determined (e.g., a microswitch failure in the case of Hamilton) without adequate evidence to come to this conclusion and without looking at all possible contributing factors. Without a thorough investigation, it is not possible to determine whether a sensor provided the wrong information, the software provided an incorrect command, or the actuator had a transient failure and did the wrong thing on its own. In the case of the Hamilton accident, a transient microswitch failure was assumed to be the cause even though the engineers were unable to reproduce the failure or to find anything wrong with the microswitch.

In general, it is a mistake to patch just one causal factor (such as the software) and assume that future accidents will be eliminated. Accidents are unlikely

to occur in exactly the same way again. If we patch only the symptoms and ignore the deeper underlying causes, or if we fix only the specific cause of one accident, we are unlikely to have much effect on future accidents. The series of accidents involving the Therac-25 is a good example of exactly this problem: Fixing each individual software flaw as it was found did not solve the safety problems of the device.

Complacency. Often it takes an accident to alert people to the dangers involved in technology. A medical physicist wrote about the Therac-25 accidents:

In the past decade or two, the medical accelerator “industry” has become perhaps a little complacent about safety. We have assumed that the manufacturers have all kinds of safety design experience since they’ve been in the business a long time. We know that there are many safety codes, guides, and regulations to guide them and we have been reassured by the hitherto excellent record of these machines. Except for a few incidents in the 1960’s (e.g., at Hammersmith, Hamburg) the use of medical accelerators has been remarkably free of serious radiation accidents until now. Perhaps, though we have been spoiled by this success [284].

This problem seems to be common in all fields.

Unrealistic Risk Assessments. The first hazard analyses initially ignored software, and then they treated it superficially by assuming that all software errors were equally likely. The probabilistic risk assessments generated undue confidence in the machine and in the results of the risk assessment themselves. When the first Yakima accident was reported to AECL, the company did not investigate. Their evidence for their belief that the radiation burn could not have been caused by their machine included a probabilistic risk assessment showing that safety had increased by five orders of magnitude as a result of the microswitch fix.

The belief that safety had been increased by such a large amount seems hard to justify. Perhaps it was based on the probability of failure of the microswitch (typically 10^{-5}) AND-ed with the other interlocks. The problem with all such analyses is that they typically make many independence assumptions and exclude aspects of the problem—in this case, software—that are difficult to quantify but which may have a larger impact on safety than the quantifiable factors that are included.

Inadequate Investigation or Followup on Accident Reports. Every company building safety-critical systems should have audit trails and incident analysis procedures that are applied whenever any hint of a problem is found that might lead to an accident. The first phone call by Tim Still should have led to an extensive investigation of the events at Kennestone. Certainly, learning about the first lawsuit should have triggered an immediate response.

Inadequate Software Engineering Practices. Some basic software engineering principles that apparently were violated in the case of the Therac-25 include the following:

- Software specifications and documentation should not be an afterthought.
- Rigorous software quality assurance practices and standards should be established.
- Designs should be kept simple and dangerous coding practices avoided.
- Ways to detect errors and get information about them, such as software audit trails, should be designed into the software from the beginning.
- The software should be subjected to extensive testing and formal analysis at the module and software level; system testing alone is not adequate. Regression testing should be performed on all software changes.
- Computer displays and the presentation of information to the operators, such as error messages, along with user manuals and other documentation need to be carefully designed.

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 of the testing was done as a system. It appears that unit and software testing was minimal, with most of the effort directed at the integrated system test. At a Therac-25 user’s meeting, the same man stated that the Therac-25 software was tested for 2,700 hours. Under questioning by the users, he clarified this as meaning “2700 hours of use.” The FDA difficulty in getting an adequate test plan out of the company and the lack of regression testing are evidence that testing was not done well.

The design is unnecessarily complex for such critical software. It is untestable in the sense that the design ensured that the known errors (there may very well be more that have just not been found) would most likely not have been found using standard testing and verification techniques. This does not mean that software testing is not important, only that software must be designed to be testable and that simple designs may prevent errors in the first place.

Software Reuse. Important lessons about software reuse can be found in these accidents. A naive assumption is often made that reusing software or using commercial off-the-shelf software will increase safety because the software will have been exercised extensively. Reusing software modules does not guarantee safety in the new system to which they are transferred and sometimes leads to awkward and dangerous designs. Safety is a quality of the system in which the software is used; it is not a quality of the software itself. Rewriting the entire software in order to get a clean and simple design may be safer in many cases.

Safe versus Friendly User Interfaces. Making the machine as easy as possible to use may conflict with safety goals. Certainly, the user interface design left much to be desired, but eliminating multiple data entry and assuming that operators would check the values carefully before pressing the return key was unrealistic.

User and Government Oversight and Standards. Once the FDA got involved in the Therac-25, their response was impressive, especially considering how little experience they had with similar problems in computer-controlled medical devices. Since the Therac-25 events, the FDA has moved to improve the reporting system and to augment their procedures and guidelines to include software. The input and pressure from the user group was also important in getting the machine fixed and provides an important lesson to users in other industries.

SAFEWARE

System Safety and Computers

Nancy G. Leveson
University of Washington

QA
76.76
.R44
L48
1995



SAD

ADDISON-WESLEY PUBLISHING COMPANY

Reading, Massachusetts • Menlo Park, California • New York • Don Mills, Ontario
Wokingham, England • Amsterdam • Bonn • Sydney • Singapore
Tokyo • Madrid • San Juan • Milan • Paris

**INDIANA UNIVERSITY
SOUTH BEND**