INT 601
Fall 2019

The Therac-25 Fiasco

or, Hanlon's Razor Cuts Both Ways

BENJAMIN D. HUTCHINS

In 1982,[1] the Canadian crown corporation[2] Atomic Energy of Canada Ltd. (AECL) launched the Therac-25, a radiation therapy machine touted as being capable of operating in two different modes. It could be configured to deliver either a relatively low-powered beam of beta radiation (i.e., electrons), or a more powerful X-ray beam.

What no one knew at the time, and would not discover for several years, was that the Therac-25 had an unintentional *third* mode: Under exactly the wrong circumstances, it could deliver a raw 25-million-electron-volt (MeV) electron beam, hitting the patient with 250 gray (25,000 rad) of unfiltered beta rays—*twenty-five times* the whole-body dose considered lethal to 100% of patients by the Centers for Disease Control and Prevention.[3] Between 1985 and 1987, a series of misconfiguration accidents caused six patients at various facilities to be severely overdosed by the Therac-25s they were being treated with, resulting in at least two deaths.

Background: The Therac-25

The Therac-25 was a linear particle accelerator (*Therac* is a contraction of *ther*apeutic *ac*celerator), a device which uses a magnetron to accelerate electrons into a linear stream. It was the third in its line, following on successful machines designated the Therac-6 and Therac-20, respectively.[4] Like the Therac-20 before it, the Therac-25 was a *dual-mode* accelerator, meaning that it could be configured to deliver either straight

---

[1] Nancy G. Leveson and Clark S. Turner, "An Investigation of the Therac-25 Accidents". *Computer* 26, no. 7 (1993), 20.
[2] An enterprise that is operated like a corporation, but is the property of the Crown, i.e., the government.
[3] U.S. Centers for Disease Control and Prevention, "Acute Radiation Syndrome: A Fact Sheet for Clinicians". *Centers for Disease Control and Prevention: CDC 24/7.* April 4, 2018. https://www.cdc.gov/nceh/radiation/emergencies/arsphysicianfactsheet.htm (accessed December 10, 2019).
[4] The numbers in the Therac-series machines' names specify the power levels of their radiation sources in millions of electron volts.

electrons or X-rays. The mechanism by which it did this is important to understand, since it is ultimately also the mechanism by which the machine could become a lethal weapon.

In electron mode, the Therac-25 (like the Therac-20 before it) was designed to deliver a beam of electrons, attenuated by magnets to be much less powerful than the magnetron's raw 25-MeV output, to attack tumors at or near the surface of the patient's body. When the machine was shifted to X-ray mode, the magnets were replaced in the electron beam's path by a metal plate. The full-power electron beam would strike the metal plate, which would block the stream of electrons and emit a stream of X-ray photons from its other side, in effect "converting" the beam from electrons to X-rays. These X-rays would then pass through a modulating device, which altered the beam to produce the appropriate dosage, and irradiate tumors deeper within the patient's body.[5]

Systems operating on the dual-mode, magnetron-source principle have advantages over therapy machines that use a fixed radioactive object (e.g., a pellet of a concentrated radioisotope such as cesium-137) as their X-ray source. From a mercantile standpoint, the obvious advantage is that a medical facility does not need to buy and maintain two separate pieces of equipment, one for electron beam therapy and one for X-rays. From an occupational health and safety viewpoint, magnetron-based devices like the Therac series have no parts that are radioactive when the power is off, making them safer to operate, maintain, and eventually dispose of. (Deserviced radiotherapy machines' fixed radiation sources have fallen into the wrong hands before, with unfortunate results.[6])

---

[5] This description of the Therac-25's operating principles was gleaned from several sources, most notably Leveson-Turner 1993.

[6] International Atomic Energy Agency, *The Radiological Accident in Goiânia* (Vienna: 1988), 1.

However, their operating principle does introduce a significant danger: If the machine is operated with neither the electron beam scanning magnets *nor* the X-ray plate and beam flattener in place, the patient receives the electron beam at full power with no modulation at all—a catastrophically high dose that will do tremendous damage to the tissues it hits, with potentially fatal results.

This potential flaw is an unavoidable consequence of dual-mode operation. In order for the hardware to be flexible enough to generate both types of output, its radiation source must be capable of producing far more power than should ever be applied directly to the patient. Manufacturers of such machines were aware of this potential flaw, and AECL was no exception. It existed in the Therac-20 as well. However, the Therac-20s in service never injured a patient; several Therac-25s did. The reason for this lies in certain "improvements" that AECL made when updating the Therac-20.

Like its predecessors, the Therac-25 was controlled digitally, using a DEC PDP-11 minicomputer running a proprietary software image. Digital controls were convenient for the machines' operators and were considered more precise and less susceptible to error. The Therac-20, however, had mechanical safety interlocks built into its hardware, which physically prevented the machine from being operated in X-ray mode with the target plate out of position. In the Therac-25's design, these hardware safety mechanisms were *deliberately removed*, leaving the machine entirely dependent on its operating software for preventing potentially dangerous misconfigurations, under the mistaken impression that software safeties were more reliable.

By 1993, this practice was commonplace enough for Leveson and Turner to remark on it in their landmark article for *Computer*, writing, "This approach is becoming more

common as companies decide that hardware interlocks and backups are not worth the expense, or they put more faith (perhaps misplaced) on software than on hardware reliability."

Perhaps misplaced indeed.

Malfunction 54

Sales of the fully-digitally-controlled Therac-25 began in 1982, and the first of them became operational the following year. By 1985, 11 of the machines were in service, six at treatment facilities in Canada and five in the United States,[7] all of which had been operating without incident since going online. Then, between June 3, 1985, and January 17, 1987, six patients at four different clinics were seriously overdosed during treatment with Therac-25 machines: one in Marietta, Georgia; one in Hamilton, Ontario; one in Yakima, Washington; two in Tyler, Texas; and finally, a second patient at the clinic in Yakima.[8]

The exact number of these patients who died as a result of Therac-25 radiation exposure is variously reported. The patient in Marietta suffered permanent disability, but survived, eventually to die in an automobile accident in 1990. The Hamilton patient died from the cancer that was being treated, but was found at autopsy to have suffered a severe radiation injury as well. The first Yakima victim survived with "minor disability and some scarring".[9] Both patients in Tyler died from their overdoses, one within five months of the accident, the other (who had been receiving treatment on his face and consequently suffered massive irradiation of his brain) in only three weeks.[10] The final victim, the

[7] Leveson-Turner 1993, 21.
[8] Leveson-Turner 1993, 22.
[9] Leveson-Turner 1993, 27.
[10] Leveson-Turner 1993, 28.

second at Yakima, died three months after the exposure, and was suffering complications from the overdose at the time of his death, but was also another potentially terminal cancer patient;[11] his death may be the hardest to apportion.[12]

The fact that the two least ambiguous deaths took place in the same clinic (the East Texas Cancer Center in Tyler, Texas), within a few weeks of each other, and involving the same personnel, eventually led to the discovery of the Therac-25's fatal flaw. The clinic's health physicist, Fritz Hager, and the machine's operator (who had been at the console for both accidents) were eventually able to re-create the conditions that inadvertently weaponized the device—a combination of software faults that simultaneously allowed the machine to deliver an unfiltered 25-MeV electron beam *and* provided no warning to the operator that it had done so. (In fact, the error message the misconfiguration produced, "Malfunction 54", was so uninformative that the operator, assuming the machine had done nothing, retried twice with the first patient, thus exposing him to the full-power beam *three times*.)[13]

Hager's report to AECL about his findings caused the company to take immediate, but somewhat lukewarm, action. Upon being informed that a sufficiently fast typist trying to correct a simple typographical error could unwittingly tie the Therac-25's software into knots and cause it to enter a catastrophically undesirable operating mode, the company issued a notice to all Therac-25 operators not to use the edit function of the user interface.

---

[11] Barbara Wade Rose, "Fatal Dose". *Saturday Night*, June 1994: 24+.
[12] The ambiguity of some of the deaths seems to have clouded later reportage about the affair. For example, a 2005 *Wired* article by Simson Garfinkel, "History's Worst Software Bugs", claims that "at least five patients die[d]," which is not borne out by any other source reviewed for this paper, and directly contradicts any that were at all specific about the incidents.
[13] Steven Casey, *Set Phasers on Stun: And Other True Tales of Design, Technology, and Human Error, Second Edition* (Santa Barbara: Aegean Publishing Co, 1998), 19.

This letter advised the machines' owners to *remove the keycap from the "up" arrow key* on the system's console to prevent its use.[14] As an unidentified agent of the U.S. Food and Drug Administration (which was already investigating the Therac-25 based on reports of the first Tyler accident, and declared the machine defective two months after it happened) noted in an acerbic response, "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 that the inconvenience to operators outweighs the need to disable the key." AECL was ordered to try again.[15]

As it happened, that AECL took *any* action after Hager's report was a departure from the company's pattern of behavior up to that point in the Therac-25 affair. The Marietta clinic's health physicist, Tim Still, made contact with AECL after the first accident to inquire whether the company knew of any way the Therac-25 could overdose a patient. He was told that such a thing was impossible. AECL personnel later used the fact that Still had asked a question rather than explicitly reporting a problem as a justification for claiming, when staff at Yakima Valley Memorial Hospital inquired after the first of their hospital's two accidents (the third overall), that there had "apparently been no other instances of similar damage to… patients."[16]

By that point, an AECL investigation into the *second* accident, at the clinic in Hamilton, Ontario, had blamed the fault on a "transient failure in the microswitch used to determine turntable position"—i.e., a hardware failure—and AECL had issued a letter

---

[14] Leveson-Turner 1993, 31.
[15] Ibid.
[16] Emmanuelle Fauchart, "Moral Hazard and the Role of Users in Learning from Accidents", *Journal of Contingencies and Crisis Management* 14, no. 2 (2006): 99.

advising users to double-check the machine's positioning visually. When asked by the Canadian Radiation Protection Bureau to retrofit the hardware for greater reliability, AECL took no action.[17] No mention of this accident, its investigation, or the result thereof was made to the Yakima staff.

Even after the first Tyler accident, when AECL personnel visited the East Texas Cancer Center in person to inspect the machine, they insisted that the machine could not be at fault, and that they knew of no other incidents in which a Therac-25 had injured a patient—even though, as Emmanuelle Fauchart has pointed out in the *Journal of Contingencies and Crisis Management*, "AECL was certainly aware of the Hamilton accident that had occurred seven months before,"[18] and was then in the process of being sued by the victim of the Marietta accident. Not until the FDA got involved did the company take any action other than to deny any knowledge of previous accidents and suggest that the operators cease using a documented feature of the control software (i.e., the input editing function).

Both Fauchart and Douglas Birsch, writing for the journal *Ethics and Information Technology*, have analyzed AECL's response to the accidents from the standpoint of moral responsibility. In Fauchart's view, the critical catalyst of AECL's at best half-hearted, at worst obstructive reaction to users' reports of trouble before the FDA got involved was the asymmetry of information that existed between the company on one side, and the users of its product on the other.

---

[17] Ibid.
[18] Fauchart 2006, 100.

In the mid-1980s, the Internet was not yet a significant factor in communications outside some rarefied academic circles, which did not significantly include the medical profession. There was no formal Therac-25 users' group. The individual users of Therac-25 machines were thus initially unaware of each other's problems with the machines.

As a result, according to Fauchart (for some reason referring to AECL as a singular male), "The manufacturer was the only one to possess information on all the accidents… He thus used the information asymmetry to pretend that each accident was a one-off fluke." She argues for the establishment of formal users' groups to counteract what she describes as "opportunistic behavior" on the part of manufacturers—in essence, to keep them honest.[19]

Birsch, by contrast, ascribes the most definite moral responsibility for the Therac-25 accidents to the programmer who developed the machine's flawed operating software, and thereby hangs the oddest part of this odd technological tale.

Who Is Keyser Söze?

In Birsch's analysis, "The Programmer meets all three conditions for some degree of moral responsibility for the harm caused by the Therac-25 system failures because he or she was a significant causal link in the joint action [i.e., the production and sale of the machine], was negligent in regard to providing crucial information connected to the joint action [that the software could be dangerously faulty], and knew the dangers of being negligent." But "the Programmer", as it turns out, is the most mysterious figure in the entire affair.

[19] Fauchart 2006, 101.

In their 1993 report for *Computer*, Leveson and Turner described the genesis of the Therac-25's operating software thusly in a sidebar: "We know that the software for the Therac-25 was developed by a single person, using PDP-11 assembly language, over a period of several years." They then noted, a bit astonishingly: "The programmer left AECL in 1986 [i.e., right in the midst of the series of accidents]. In a lawsuit connected with one of the accidents, the lawyers were unable to obtain information about the programmer from AECL… none of the AECL employees questioned could provide any information about his educational background or experience… *We have been unable to learn anything about his background.*"[20] [Emphasis mine.]

This mysterious lone programmer had recycled the operating image from the Therac-6 and Therac-20 without discovering or correcting dangerous flaws. Subsequent investigation showed that the Therac-20 had the same "edit mode" problem as the Therac-25, but had been prevented from harming anyone by its hardware interlocks—the hardware interlocks that were deleted from the Therac-25 as unnecessary. He also introduced new problems, including a separate bug that caused the second Yakima accident and is believed to have caused the one in Hamilton. AECL admitted during the FDA recall process that the system had never been properly tested. And then, in the midst of the developing crisis over the machine's lethal shortcomings, he seems to have quietly decamped. Lawyers involved in ongoing liability lawsuits concerning his work were unable to find him. His own former co-workers professed to know nothing about him. In the historical record, he becomes a strange, shadowy, vaguely sinister figure who may or

[20] Leveson-Turner 1993, 20.

may not even exist, a sort of PDP-11-programming Keyser Söze (the widely-feared but possibly mythical crime lord in Bryan Singer's 1995 motion picture *The Usual Suspects*).

What few mentions of this individual exist in later literature do nothing to clarify the point. Even his name seems to be unknown. The most recent reference to him I could find, in a 2016 article by Patricia A. McQuaid in *Journal of Software: Evolution and Process*, only deepens the mystery, noting, "To this day, not much is known about the sole programmer who ultimately created the software, other than that he had minimal formal training in writing software."[21] Tantalizingly, the article contains no indication of where that last data point comes from.

To be sure, the mysterious programmer's professional qualifications (or lack thereof) were far from the only ethical problem facing the Therac-25. We have seen already that, as McQuaid summarized, "It was a common practice for [AECL's] engineering and other departments to dismiss claims of machine malfunction as user error, medical problems with the patient beyond AECL's control, and other circumstances wherein the blame would not fall on AECL." Beyond having to be forced to admit that anything was happening, let alone take action, the company also manifestly lacked a quality culture, particularly in its software development and (lack of) testing policies.

Conclusion

In the end, no more Therac-25s were sold, although the existing machines were eventually retrofitted (after a great deal of wrangling with the FDA) and put back into service—except the one at the East Texas Cancer Center, which was returned to the

[21] Patricia A. McQuaid, "Software disasters—understanding the past, to improve the future", *Journal of Software: Evolution and Process* 2012, no. 24: 465.

manufacturer after staff refused to use it.[22] AECL spun its medical equipment arm off into a separate corporation, rebranded "Theratronics", the year after the last of the Therac-25 accidents. No one was ever criminally prosecuted for the production and sale of the lethally defective machine, nor for the dilatory and evasive response of its manufacturers to reports of the initial accidents (when a more timely and meaningful reaction could conceivably have prevented later ones).

Its manufacturers' sense of urgency or lack thereof notwithstanding, lessons *were* learned from the Therac-25. In the software engineering community, the Therac-25 fiasco is a benchmark case study in many different fields, from human factors to user-experience design to how *not* to react to user reports of serious problems. As a result of the lengthy investigation and contentious rectification process that ensued, the FDA began taking a harder line on the review of software used in medical equipment (which had previously been seen as something of an afterthought to certification of the hardware).[23] We may be forgiven for hoping that, as a result, patients are now safer than they were in 1985.

[22] Rose 1994.
[23] Ibid.

Bibliography

Birsch, Douglas. "Moral responsibility for harm caused by computer system failures." *Ethics and Information Technology*, no. 6 (2004): 233-245.

Casey, Steven. *Set Phasers on Stun: And Other True Tales of Design, Technology, and Human Error, Second Edition.* Santa Barbara, California: Aegean Publishing Co, 1998.

Fauchart, Emmanuelle. "Moral Hazard and the Role of Users in Learning from Accidents." *Journal of Contingencies and Crisis Management* 14, no. 2 (2006): 97-106.

Garfinkel, Simson. "History's Worst Software Bugs." *wired.com.* November 8, 2005. https://www.wired.com/2005/11/historys-worst-software-bugs/ (accessed December 10, 2019).

International Atomic Energy Agency. *The Radiological Accident in Goiânia.* Vienna: IAEA, 1988.

Israelski, Edmond W., and William H. Muto. "Use-Error Focused Risk Analysis for Medical Devices: A Case Study of the Therac-25 Radiation Therapy System." *Proceedings of the Human Factors and Ergonomics Society 47th Annual Meeting*, 2003: 1564-1568.

Leveson, Nancy G. "The Therac-25: 30 Years Later." *Computer* 50, no. 11 (November 2017): 8-11.

Leveson, Nancy G., and Clark S. Turner. "An Investigation of the Therac-25 Accidents." *Computer* 26, no. 7 (1993): 18-41.

McDaniel, James G. "Improving system quality through software evaluation." *Computers in Biology and Medicine* 32 (2002): 127-140.

McQuaid, Patricia A. "Software disasters—understanding the past, to improve the future." *Journal of Software: Evolution and Process* 24 (2012): 459-470.

Neumann, Peter G. "Risks to the Public in Computers and Related Systems." *ACM SIGSOFT Software Engineering Notes* 12, no. 3 (1987): 2-17.

Parnas, David L., John van Schouwen, and Shu Po Kwan. "Evaluation of Safety-Critical Software." *Communications of the ACM* 33, no. 6 (1990): 636-648.

Rose, Barbara Wade. "Fatal Dose." *Saturday Night*, June 1994: 24+.

United States Centers for Disease Control and Prevention. "Acute Radiation Syndrome: A Fact Sheet for Clinicians." *Centers for Disease Control and Prevention: CDC 24/7.* April 4, 2018. https://www.cdc.gov/nceh/radiation/emergencies/arsphysicianfactsheet.htm (accessed December 11, 2019).