Therac-25

The Therac-25 machine was a state-of-the-art linear accelerator developed by the company Atomic Energy Canada Limited (AECL) and a French company CGR to provide radiation treatment to cancer patients. The Therac-25 was the most computerized and sophisticated radiation therapy machine of its time. With the aid of an onboard computer, the device could select multiple treatment table positions and select the type/strength of the energy selected by the operating technician. AECL sold eleven Therac-25 machines that were used in the United States and Canada beginning in 1982.

Unfortunately, six accidents involving significant overdoses of radiation to patients resulting in death occurred between 1985 and 1987 (Leveson & Turner 1993). Patients reported being "burned by the machine" which some technicians reported, but the company thought was impossible. The machine was recalled in 1987 for an extensive redesign of safety features, software, and mechanical interlocks. Reports to the manufacturer resulted in inadequate repairs to the system and assurances that the machines were safe. Lawsuits were filed, and no investigations took place. The Food and Drug Administration (FDA) later found that there was an inadequate reporting structure in the company, to follow up with reported accidents.

There were two earlier versions of the Therac-25 unit: the Therac-6 and the Therac-20, which were built from the CGR company's other radiation units--Neptune and Sagittaire. The Therac-6 and Therac-20 units were built with a microcomputer that made the patient data entry more accessible, but the units were operational without an onboard computer. These units had built-in safety interlocks and positioning guides, and mechanical features that prevented radiation exposure if there was a positioning problem with the patient or with the components of the machine. There was some "base duplication" of the software used from the Therac-20 that carried over to the Therac-25. The Therac-6 and Therac-20 were clinically tested machines with an excellent safety record. They relied primarily on hardware for safety controls, whereas the Therac-25 relied primarily on software.

On February 6, 1987, the FDA placed a shutdown on all machines until permanent repairs could be made. Although the AECL was quick to state that a "fix" was in place, and the machines were now safer, that was not the case. After this incident, Leveson and Turner (1993) compiled public information from AECL, the FDA, and various regulatory agencies and concluded that there was inadequate record keeping when the software was designed. The software was inadequately tested, and "patches" were used from earlier versions of the machine. The premature assumption that the problem(s) was detected and corrected was unproven. Furthermore, AECL had great difficulty reproducing the conditions under which the issues were experienced in the clinics. The FDA restructured its reporting requirements for radiation equipment after these incidents.

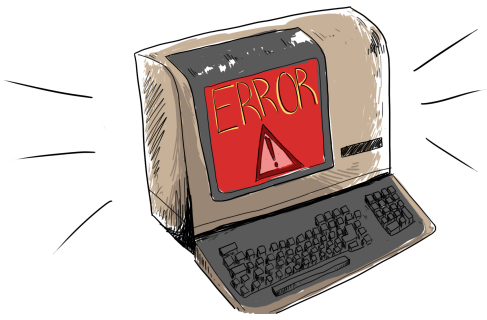

As computers become more and more ubiquitous and control increasingly significant and complex systems, people are exposed to increasing harms and risks. The issue of accountability arises when a community expects its agents to stand up for the quality of their work. Nissenbaum (1994) argues that responsibility in our computerized society is systematically undermined, and this is a disservice to the community. This concern has grown with the number of critical life services controlled by computer systems in the governmental, airline, and medical arenas.

According to Nissenbaum, there are four barriers to accountability: the problem of many hands, "bugs" in the system, the computer as a scapegoat, and ownership without liability. The problem of too many hands relates to the fact that many groups of people (programmers, engineers, etc.) at various levels of a company are typically involved in creation of a computer program and have input into the final product. When something goes wrong, there is no one individual who can be clearly held responsible. It is easy for each person involved to rationalize that he or she is not responsible for the final outcome, because of the small role played. This occurred with the Therac-25 that had two prominent software errors, a failed microswitch, and a reduced number of safety features compared to earlier versions of the device. The problem of bugs in the software system causing errors in machines under certain conditions has been used as a cover for careless programming, lack of testing, and lack of safety features built into the system in the Therac-25 accident. The fact that computers "always have problems with their programming" cannot be used as an excuse for overconfidence in a product, unclear/ambiguous error messages, or improper testing of individual components of the system. Another potential obstacle is ownership of proprietary software and an unwillingness to share "trade secrets" with investigators whose job it is to protect the public (Nissenbaum 1994).

The Therac-25 incident involved what has been called one of the worst computer bugs in history (Lynch 2017), though it was largely a matter of overall design issues rather than a specific coding error. Therac-25 is a glaring example of what can go wrong in a society that is heavily dependent on technology.

Discussion Questions:

1. Who should be responsible for the errors in a medical device?

2. What moral responsibility do creators of software have for the adverse consequences that flow from flaws in that software?

3. What steps are creators of software morally required to take to minimize the risk that they will sell flawed software with dangerous consequences?

4. What should constitute FDA approval of a medical device? Should the benefit outweigh the harm? Should the device be 100% safe prior to approval? Should FDA approval guidelines take into consideration novel therapies for protected populations such as children or patients with rare conditions?

5. Should updated medical devices be reviewed by the FDA as a new device or as an improvement in an older design? If reviewed as an improvement, at what point can/should a device be subject to a full review process? If reviewed as a novel device, how might this effect the production of modified/ improved devices and the overall companies that produce medical devices?

Resources:

Gotterbarn, Donald, "Software Engineering Ethics," Encyclopedia of Software Engineering (2002), https://onlinelibrary.wiley.com/doi/abs/10.1002/0471028959.sof314.

Leveson, Nancy & Turner, Clark, "An Investigation of the Therac-25 Accidents," Computer 26:7, p. 18 (July 19993), https://web.stanford.edu/class/cs240/old/sp2014/readings/therac-25.pdf
Leveson, Nancy, Medical Devices: The Therac—25 (1995), http://sunnyday.mit.edu/papers/therac.pdf.

Lynch, Jamie, "Therac-25 Causes Radiation Overdoses," Bugsnag Blog (2017)
https://www.bugsnag.com/blog/bug-day-race-condition-therac-25
Nissenbaum, Helen, "Computing and Accountability," Communications of the ACM, 37:1, p. 73 (1994).
http://delivery.acm.org/10.1145/180000/175228/p72-nissenbaum.pdf?ip=128.62.211.38&id=175228&acc=ACTIVE%20SERVICE&key=603D2E7028CD4EF5%2E4D4702B0C3E38B35%2E4D4702B0C3E38B35%2E4D4702B0C3E38B35&__acm__=1576603038_0292b9d0b31643bd7b06dde8efa509f7

Author:
Chris Apgar and Robert Prentice
The University of Texas at Austin