Jon Binney
803-106-721

Response to Therac-25 Paper

The design process and testing procedures for the Therac-25 were deeply flawed, and produced a system with many faults. The design and testing of each subsystem assumed the correct operation of the other subsystems. For example, the software assumed that all of the microswitches worked properly, and the hardware assumed that the software verified the turntable position correctly before starting the treatment.

The company, AECL, apparently had experience producing machines for the same purpose. They based much of the design for the Therac-25 on an earlier model, the Therac-6. The most notable difference is that the Therac-6 had merely added computer control to an existing manually operated system, whereas the Therac-25 was designed to "take control of computer control from the outset; AECL did not build on a stand-alone machine." (Leveson, 20) For the Therac-25, many of the safety features were moved into software. Another model produced by AECL, the Therac-20, was also based on the Therac-6, but kept most of the hardware safeties. Both the Therac-25 and the Therac-20 had software bugs, but "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." (Leveson, 29)

Even when assessed separately from the hardware system, the software design process was quite poor. The software was developed by a single person in assembly, and there was a lack of documentation. The user interface was particularly bad; it provided ways for the operator to quickly resume treatment after errors, without restarting the entire process. This combined with problems with the data input routines led to treatments with improper settings.

The problems with the design process meant that a robust and accurate system was probably beyond hope, but thorough testing could still have kept faulty machines from being sent to consumers. The testing was blatantly insufficient, and in particular, "unit and software testing was minimal, with most effort directed at the integrated system test" (Leveson, 20) Because of the complexity of the system, testing it as one unit could not possibly reveal all of the problems. The number of permutations of conditions which could occur, including operator actions, patient actions, and interactions with the environment were huge, and could not all be tested. Because of this, it was essential for each part of the system to be tested thoroughly on its own, including testing the response of the module when other parts of the system acted incorrectly. Assuming that the other modules will always work as specified creates a fragile system which could fail catastrophically whenever one module fails.

Additionally, when testing the system after problems had already occurred, the underlying cause of the problem was never verified. A possible source was found and fixed, but the possibility of other failure points was not considered. No unified software test plan was created and carried out to ensure the correct operation of the device. Instead varieties of poorly documented tests, as well as past experience, were relied upon. Combined with the disorganized design process, this led to a dangerous and error prone system.