

## The Therac 25

A case study in safety failure

- Radiation therapy machine
- “The most serious computer-related accidents to date”
- People were killed
- References:


Nancy Leveson and Clark Turner, “The Investigation of the Therac-25 Accidents”, Computer, 26, 7 (July 1993) pp 18-41.

Nancy Leveson, “Medical Devices: The Therac-25” appendix in *Software: System Safety and Computers*, Addison-Wesley, 1995

## Therac 25 Background

- Medical linear accelerator developed by Atomic Energy of Canada, Ltd. in mid-1970s
- Delivers 25 MeV photons or electrons of various energies
- Controlled by PDP-11 minicomputer
- Software responsible for safety
- Software adapted from earlier Therac-6 & Therac 20 systems, which had hardware interlocks for safety

## The Therac 25



## Therac 25 Turntable




Figure B. Upper turntable assembly.

## Therac 25 Turntable

- Electron mode
  - 5-25 MEV
  - Magnets spread beam
  - Ion chamber monitor
- X-ray mode
  - 25 MEV electrons hit target
  - “Beam flattener” attenuates
  - 100x beam current
  - Ion chamber monitor
- Field-light mode
  - No current
  - Mirror & light used to check alignment
  - No ion chamber (since not treating)




Figure B. Upper turntable assembly.

## Therac 25 Turntable

- Computer adjusts turntable position
- Microswitches detect turntable setting
- 3-bit binary code used to encode turntable setting
- Software checks replace hardware interlocks




Figure B. Upper turntable assembly.

## Therac 25 Software Development

- Evolved from Therac 6 system (1972-1976)
- Incorporated some Therac 20 code, as well
- Written in PDP-11 assembler
- Custom operating system
- Little documentation during development
- Minimal unit and software testing
- Q/A testing was 2700 hours of use as integrated system

## Therac 25 Software Development

- Evolved from Therac 6 system (1972-1976)
- Incorporated some Therac 20 code, as well
- Written in PDP-11 assembler
- Custom operating system
- Little documentation during development
- Minimal unit and software testing
- Q/A testing was 2700 hours of use as integrated system
- **Programmer left AECL in 1986, little information available about his background**



*"I know this may be an awkward time,  
but do you recall him ever mentioning source code."*

## Therac 25 Software Functions

- Monitors machine status
- Sets up machine for treatment
- Turns beam on and off in response to operator
- Monitors interlocks
- If fault, either prevents treatment start or causes a pause/suspend

## Therac 25 Software Structure

- Critical tasks:
  - Treatment monitor
  - Servo
  - Housekeeping
- Non-critical tasks:
  - Checksum
  - Keyboard
  - Calibration
  - etc.
- Concurrent access to shared memory, “test” and “set” of variables not indivisible, race conditions

## Operator Procedures

- Position patient on table
- Manually set treatment field size and gantry rotation; attach accessories
- Leave room
- Use VT-100 console to enter patient data, dose data, etc.
- (System compares manual settings with system values)
- If “verified”, operator can start machine
- Else must re-enter data

## Operator Screen Layout

PATIENT NAME : TEST	BEAM TYPE: X	ENERGY (MeV): 25	
TREATMENT MODE : FIX			
	ACTUAL	PRESCRIBED	
UNIT RATE/MINUTE	0	200	
MONITOR UNITS	50 50	200	
TIME (MIN)	0.27	1.00	
GANTRY ROTATION (DEG)	0.0	0	VERIFIED
COLLIMATOR ROTATION (DEG)	359.2	359	VERIFIED
COLLIMATOR X (CM)	14.2	14.3	VERIFIED
COLLIMATOR Y (CM)	27.2	27.3	VERIFIED
WEDGE NUMBER	1	1	VERIFIED
ACCESSORY NUMBER	0	0	VERIFIED
DATE : 84-OCT-26	SYSTEM : BEAM READY	OP. MODE : TREAT	AUTO
TIME : 12:55: 8	TREAT : TREAT PAUSE	X-RAY	173777
OPR ID : T25V02-R03	REASON : OPERATOR	COMMAND:	

## Operator Procedures

- **Complaint**
  - Re-entering all that data manually is very tedious
- **Response**
  - Set things up so that “carriage return” copies previous data for entry
  - Series of carriage returns effectively permits fast re-entry of unchanged parts of data

## Operator Procedures

- **Error Conditions**
  - “Treatment suspend” requires complete machine reset
  - “Treatment pause” can be resumed if operator types “P” at console
  - Machine insists on reset after 5 “P”s
  - Malfunction messages fairly common & usually do not affect safety
- **Error Messages**
  - Cryptic
  - Some were of the form “Malfunction NN”

## FDA Comment on Manual

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.

from *Nancy Leveson, “Medical Devices: The Therac-25”* appendix in *Software: System Safety and Computers*, Addison-Wesley, 1995



## Hazard Analysis by AECL

**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.

from *Nancy Leveson, “Medical Devices: The Therac-25”* appendix in *Software: System Safety and Computers*, Addison-Wesley, 1995

## Accident History

- 11 Therac 25s installed (5 US, 6 Canada)
- Six accidents involving massive overdoses between 1985 and 1987
- Machines recalled in 1987
- Related problems in Therac 20 discovered later but hardware interlocks prevented injuries

## Accident History

- June 3, 1985
  - Kennestone Regional Oncology Center, Marietta, Ga.
  - Never really investigated
- July 26, 1985
  - Hamilton, Ontario
  - AECL decides failing microswitch was cause
  - Independent consultant recommended adding a potentiometer
- September 1985
  - AECL makes first round of changes and notifies customers

## Yakima Valley, December 1985

- Therac 25 modified in September 1985 in response to earlier overdose problems in Hamilton, Ontario.
- Woman treated in December 1985
- Developed parallel-striped red pattern on right hip
- Treatments continued until January 6, 1986 because reaction was not determined to be abnormal
- Hospital staff investigated various causes such as heating pad patient slept on. But were puzzled because nothing seemed to fit.
- Eventually described problem as “cause unknown”

## Yakima Valley, 1987

- Second overdose in Feb. 1987 led hospital staff to suspect that first incident was also an overdose.
- Further investigation showed signs of tissue damage in first patient, which was repaired surgically. Patient survived.
- Staff concluded that first overdose must have been less severe than second, since damage only developed some time after the exposure.

## Yakima Valley, 1985

- In report written after second overdose, medical physicist said:



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 percent) we did not believe that there could have been any accelerator malfunction. These modifications to the accelerator were completed on 5,6-Sep-85.

### E.g., East Texas, March 1986

- History of 500 patients treated successfully
- Prescribed: 22MeV electrons, 180 rads
- Operator selected x-rays by mistake, used cursor keys to change to electrons
- Machine tripped with “Malfunction 54”
  - Documentation explains this is “dose input 2” error
- Operator proceeded; machine tripped again

### E.g., East Texas, March 1986

- Patient felt something wrong on first jolt, tried to get up
- Video/audio links to room not functioning
- Patient felt jolt on arm while getting up, pounded on door
- Treatment cancelled for day
- Calibration checks seemed normal
- Later found patient had gotten 16,500-25,000 rads over 1 cm square
- Patient eventually died after 5 months

### E.g., East Texas, March 1986

- AECL engineers could not replicate a Malfunction 54
- AECL home office engineer said machine could not overdose patient
- AECL suggested patient got an electric shock
- No grounding problems found
- Machine returned to service April 7, 1986

### East Texas/ April 11, 1986

- Prescription 10 MeV, area 7 x 10 cm
- Operator used cursor keys to change x-rays to electrons, saw “beam ready”, and turned machine on
- Loud noise, shutdown, malfunction 54
- Patient in great pain
- Patient died three weeks later

## East Texas/ April 11,1986

- Machine taken out of service
- ETCC eventually reproduced malfunction 54
  - Data entry speed critical factor
  - Observed 4000 rad dose
- AECL later measured 25,000 rads
- In lawsuit, earlier “cursor up” problems reported, which AECL believed to have been fixed

## Tyler Accident Race Condition

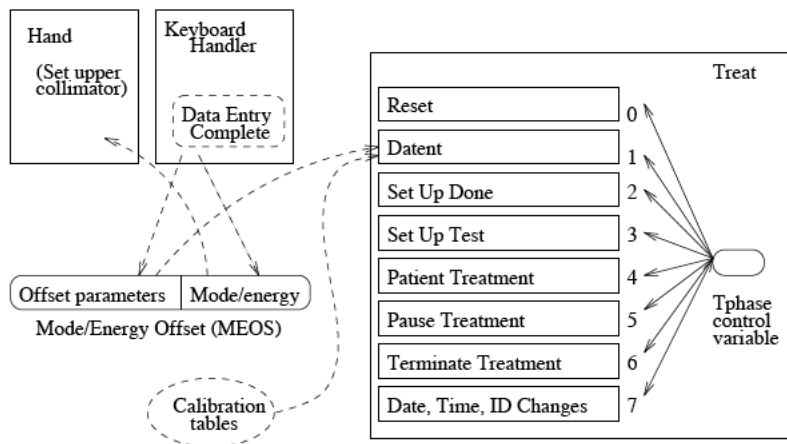


Figure 3: Tasks and subroutines in the code blamed for the Tyler accidents.

Nancy Leveson and Clark Turner, "The Investigation of the Therac-25 Accidents", *Computer*, 26, 7 (July 1993) pp 18-41.

## Race Condition

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 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.

Nancy Leveson and Clark Turner, "The Investigation of the Therac-25 Accidents", Computer, 26, 7 (July 1993) pp 18-41.

## Datent Subroutine

```
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
```

Nancy Leveson and Clark Turner, "The Investigation of the Therac-25 Accidents", Computer, 26, 7 (July 1993) pp 18-41.

## Magnet Subroutine

```
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
```

Takes about 8 secs and invoked multiple times

Nancy Leveson and Clark Turner, "The Investigation of the Therac-25 Accidents", *Computer*, 26, 7 (July 1993) pp 18-41.

## Race Condition

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.

from Nancy Leveson, "Medical Devices: The Therac-25" appendix in *Software: System Safety and Computers*, Addison-Wesley, 1995



## East Texas Govt & User Response

- Report to FDA on April 15, 1986
- Sent letter recommending temporary fix to all users

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 [an] "R" reset command must be used and the whole prescription reentered.

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

## Response, continued

- FDA comment

We have reviewed Mr. Downs' 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.

## Response, continued

- First fix plan – June 13, 1986
  - Fixed software to eliminate specific bug
  - Modified software sample-and hold circuits to detect pulse above a threshold. Shut down if have one pulse exceeding threshold, rather than 3.
  - Malfunctions 1-64 now suspend treatment, not pause it
  - Added circuit to turn off beam independent of software
  - Modify editing software to limit cursor up, etc.
  - Modify manuals
- FDA had numerous internal concerns
- FDA Letter of 7/23 agreed conceptually, but complained about lack of specific information to evaluate plan. Requested detailed description of software development procedures.

## Response, continued

- FDA Internal Memo of October 20

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 QA [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 [accidental radiation overdoses] in Texas. AECL had received official notification of a lawsuit 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 QA program currently exists. For example, a response has not been provided with respect to the software part of the CAP to the CDRH [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.

## Response, continued

- Second revised plan December 22, 1986
  - Included meaningful messages, software modifications, expanded test plan, etc.
- Sent “Component and Installation Test Plan” on Jan 26, 1987.
  - Company explained that delays were due to investigation of a new accident on Jan 17, in Yakima, California.

## Yakima Valley, January 1987

- Plan: 2 film verification exposures (3 & 4 rads) + 79 rad photon treatment
- Performed two film exposures
- Operator used hand controls to rotate table to field-light position & check alignment
- Operator set machine but forgot to remove film
- Operator turned beam on, machine showed no dose & displayed fleeting message
- Operator proceeded from pause

## Yakima Valley, January 1987

- After another machine pause, operator reentered room.
- Patient complained of burning sensation
- Patient developed severe striped burns
- Patient died in April
- Hospital obtained similar pattern on film by running machine with turntable in field light position

## Responses

- Voluntary Class II recall 8/1/85
- AECL accident report April 15, 1986
- First version of corrective action plan 6/13/86
- Second Yakima overdose 1/17/87
- Fifth (final) corrective action plan 7/21/87
- Interim safety analysis report 1/29/88
- Final safety analysis report 11/3/88

## A bit more detail on operator procedure

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.

## Yakima Accident Race Condition

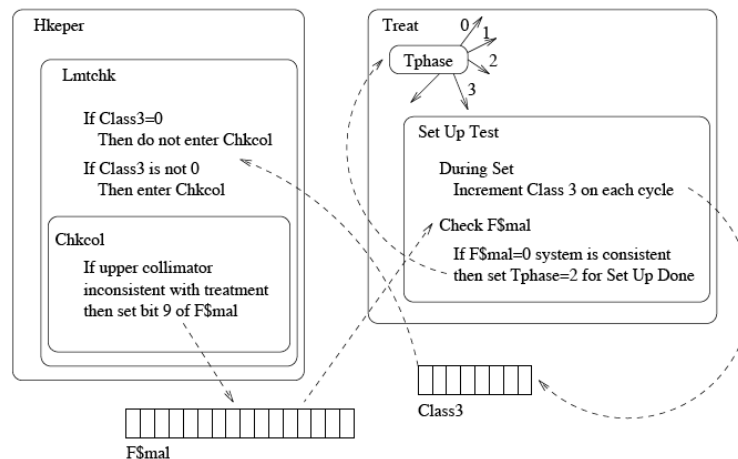


Figure 4: The Yakima software flaw.

Nancy Leveson and Clark Turner, "The Investigation of the Therac-25 Accidents", *Computer*, 26, 7 (July 1993) pp 18-41.

## The Race Condition

During machine setup, Set-Up Test will be executed several hundred times since 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 1 byte, it can only contain a maximum value of 255 decimal. Thus, on every 256th pass through the Set-Up Test code, the variable overflows and has 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 the upper collimator was still in 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.

Nancy Leveson and Clark Turner, "The Investigation of the Therac-25 Accidents", *Computer*, 26, 7 (July 1993) pp 18-41.

## Therac 25 Turntable




Figure B. Upper turntable assembly.

## Corrective Action Plan

- Numerous hardware and software changes
- All interruptions related to dosimetry not continuable
- independent hardware & software shutdowns
- potentiometer on turntable
- hardware interlocks
- “dead man switch” motion enable
- Fix documentation, messages, & user manuals
- etc

## Lessons ( Leveson & Turner)

- Complacency
- Assumption that problem was understood without adequate evidence (“the last bug” fallacy).
- Sole reliance on software for safety
- Systems engineering practices

## Lessons ( Leveson & Turner)

- Documentation key from beginning
- Use established software engineering practices
- Keep designs simple
- Build in software error logging & audit trails
- Extensive software testing and formal analysis at all levels
- Revalidate reused software
- Don' t rely only on software for safety
- Do incorporate redundancy
- Pay careful attention to human factors
- Involve users at all phases