Radiation-induced Cryptographic Failures and How to Defend Against Them

Peter Gutmann

University of Auckland

Crypto Fault Attacks

If you get a fault during a crypto computation, an attacker may be able to recover your encryption key(s) from the faulty output

- First (publicly) acknowledged in the late 1990s
- Studied to death since then

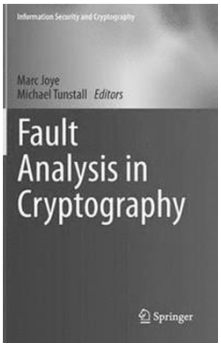

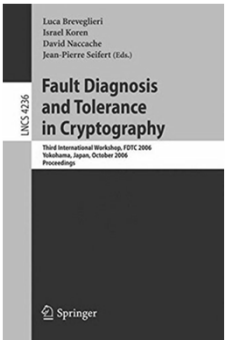

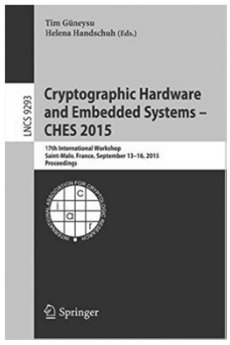

Faults in Cryptosystems

ECC is particularly susceptible to faults

- Fault with the in-memory key: Leak the private key
- Fault with the ECC computation: Leak the private key
- Fault with the RNG: Leak the private key
- You get the picture

General idea is to move the computation from the secure curve to another, inevitably weaker, one or to produce a faulty point on the original curve

> Faults can be injected in a variety of ways and almost all parts of the system can be targeted, e.g. the base point, system parameters, intermediate results, dummy operations and validation tests
> — "Fault Attacks on Elliptic Curve Cryptosystems"

Faults in Cryptosystems (ctd)

RSA also has issues, but nowhere near as bad as ECC

RSA has a fault problem in the RSA-CRT computation if you sign the exact same message twice

- This essentially never occurs
- Not in IPsec, SSH, TLS, CMS, S/MIME, PGP, SCEP, TSP, OCSP, CMP, …
- Any protocol that allowed this would also trivially allow replay attacks

ECC in contrast has entire catalogues of fault problems

- These don't require duplicate signatures

Faults in Cryptosystems (ctd)

SRP, PSK, etc have no issues

- Authentication doesn't require the use of signatures
  - Or certificates, or CAs, which is why there's close to zero support for it in browsers
- Built around MACs/PRFs (hash-based)
- Little research published on the issue, but probably because there's no obvious attack

Faults in Cryptosystems (ctd)

Symmetric crypto (e.g. AES) doesn't have radiation-related fault issues

Attacks require injection of specific attacker-controlled faults, not random faults in random locations

- Example: Create 1-byte differentials in input to AES MixColumns
- Example: Create 255 different byte faults in the AES middle rounds
- Example: Create 1-bit fault in 128 bits of SubBytes input to AES last round

Faults in Cryptosystems (ctd)

Similar to RSA, attacks require encryption of the same data two or more times

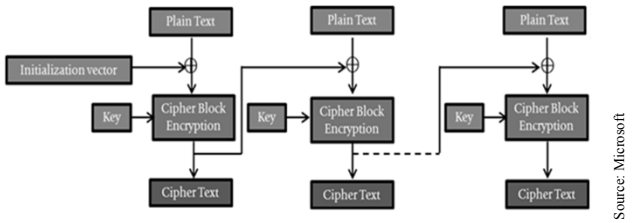

- Won't happen for the common CBC or CTR/GCM modes

CTR/GCM mode, however, fails catastrophically on an IV fault

- Both confidentiality and integrity protection collapse

Faults in Cryptosystems (ctd)

The two trendiest encryption mechanisms, ECC and AES-GCM, are also the most brittle in the presence of faults

- Worst case, a fault in the RNG, and you lose everything in one go
- ECC private key
- AES-CTR confidentiality
- AES-GCM integrity-protection

The most robust mechanisms are probably RSA and AES-CBC + HMAC

- They're not fashionable ☹

Causes of Faults

Electrical glitches

- Overvoltage
- Undervoltage
- Spikes
- Clock glitches
- Noise

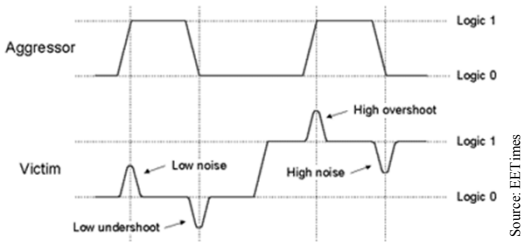

Thermal issues

Radiation

- Often induces electrical glitches
- Can also change circuit operation, temporarily or permanently

Characteristics of Radiation-induced Faults

Possible: Random bit(s) 0 → 1

Possible: Random bit(s) 1 → 0

Unlikely: Random bit fault during computation

- Most CPUs have at least error detection on the CPU core
- Some have full ECC and more, e.g. Cortex A, Cortex R, IBM Power, Intel, MIPS
- See later slides for extreme cases, e.g. Intel

Not present: Non-random, attacker-controlled faults

- In any case if an attacker can disassemble your device and sit there injecting controlled hardware faults at will, it's probably game over anyway

Characteristics of Rad-induced Faults (ctd)

Research results are often difficult to apply…

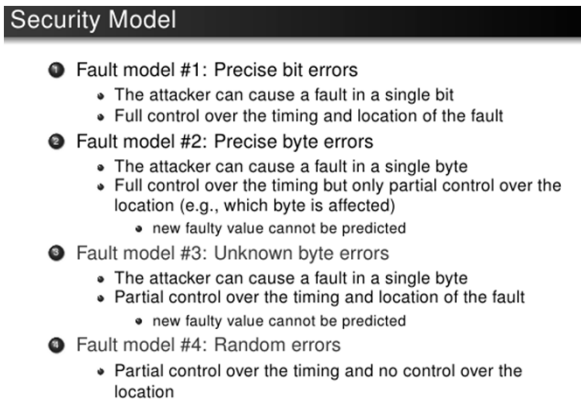

Source: Gemplus

Fault model #5: No control over the timing or location, no duplicate data to act on

Characteristics of Rad-induced Faults (ctd)

Lack of understanding by cryptographers…

> release even _one_ any way faulty signature computed using RSA-CRT and your private key walks
> — CFRG list comment

- Garbled sound byte from a 20-year-old research paper

… or appreciation that TLS crypto exists outside the web…

> […] non-starter as web browsers […] fix the reasons why web browsers […] the web browser vendors […]
> — CFRG list comment, responding to a message that talked specifically about non-browser TLS use

Characteristics of Rad-induced Faults (ctd)

… or just plain denial

I'm aware of invalid curve attacks, which can be completely mitigated by using a twist-secure curve and point compression — CFRG list comment

- "The mathematician looked at the fire extinguisher and the fire, said 'a solution exists', and went back to bed"

Cryptographers and SCADA/embedded implementers don't talk to each other

- Cryptographers

  They're not using our fine theoretical design!

- Implementers

  This stuff doesn't do what we need, we'll have to come up with our own way of doing it

When are there Radiation-induced Faults?

When you're using the crypto to monitor nuclear materials

Used to check compliance with nonproliferation treaties

Crypto in High-radiation Environments

Monitoring of fuel storage ponds

Source: World Nuclear News

Ensure fuel rods don't go missing (particularly in breeder reactors)

Crypto in High-radiation Environments (ctd)

Monitoring of reactor refueling

Source: World Nuclear News

Check what goes in and out

Crypto in High-radiation Environments (ctd)

Monitoring of waste management

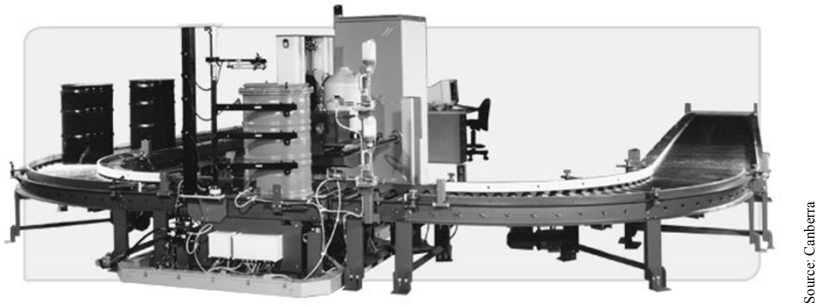

Source: Canberra

Check what's leaving the facility via nondestructive assay (NDA)

Crypto in High-radiation Environments (ctd)

Most of those aren't truly high-radiation environments

- Humans have to work there

Higher-than-normal radiation, but not classed as high-radiation

- Other equipment is deployed to high-radiation areas

Leads to an interesting definition of tamper-*discouraging* crypto

> It would take you three days to put up the scaffolding and disassemble the monitoring gear. The radiation would kill you in one day

- Who needs "tamper-resistant" when you've got that…

See later slides for environmental requirements

Crypto in Harsh Environments

Not specific to reactors though…

Devices can experience faults in harsh environments in general

- Covered by numerous standards
- EN 50128 – Railway applications – Communication, signalling and processing systems
- EN 50129 – Railway applications – Safety related electronic systems for signalling
- EN 50402 – Requirements on the functional safety of fixed gas detection systems
- IEC 60601 – Medical electrical equipment safety

[Continues]

Crypto in Harsh Environments (ctd)

[Continued]

- IEC 60880 – Nuclear power plants – Instrumentation and control systems important to safety
- IEC 61508 – Functional Safety
- IEC 61511 – Safety instrumented systems for the process industry sector (also ANSI S84)
- IEC 61513 – Nuclear power plants – Instrumentation and control important to safety
- IEC 62061 – Functional safety of electrical, electronic and programmable electronic control systems (also ISO 13849)
- ISO 26262 – Road vehicles – Functional safety

Many, many more

Notable Failures due to Ionising Radiation

Advanced Simulation and Computing Program (ASC) Q Supercomputer at Los Alamos

Source: LANL

- Built with DEC Alpha 21264 CPUs

Notable Failures due to Ionising Radiation

Error detection but not correction on level 3 cache tag (BTAG) RAM

- Too much slowdown in this speed-critical case
- Standard data RAM does have ECC

Faults were detected via parity checks, but not corrected, node crashed

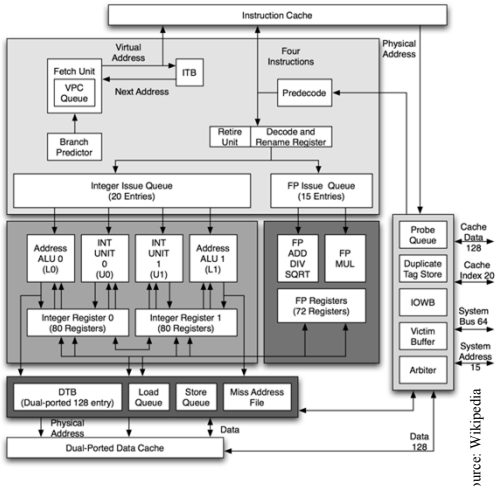

Source: Wikipedia

Notable Failures due to Ionising Radiation

Suspected cosmic rays

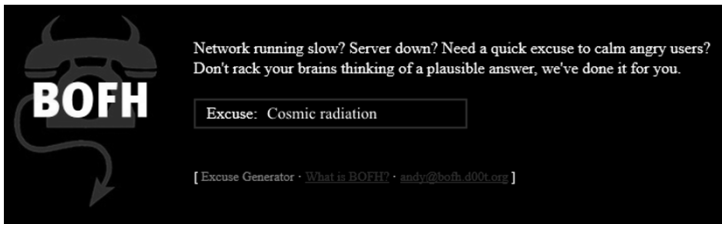

Source: d00t.org

ASC Q is at LANL, elevation 7,500ft (2,300m)

- Cosmic radiation is 6x as intense as at sea level

For comparison, avionics computers are at 30,000ft

- Radiation is 150x as intense as at sea level

Notable Failures due to Ionising Radiation

Single node at sea level experiences fatal soft error once in 50 years

- 500-node cluster at elevation experiences one every 1½ hours

Los Alamos just happens to have the Los Alamos Neutron Science Centre (LANSCE)

- Confirmed that it's radiation-induced

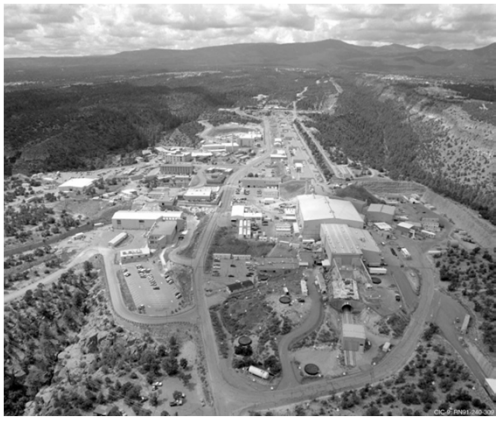

Source: LANL

Notable Failures due to Ionising Radiation

Dealt with by

- Scrubbing cache RAM before program runs
- Checkpointing during runs to allow recovery
- Leaving spare nodes available to restart failed jobs on
- etc

(NB: Often-repeated 2016 IEEE Spectrum article mentions more examples, but these contain multiple factual errors and/or are unverifiable. Don't believe what Google will turn up).

Radiation-induced Effects

Total Ionising Dose, TID

- Long-term, permanent damage

Single Event Effects, SEE

- Transient damage

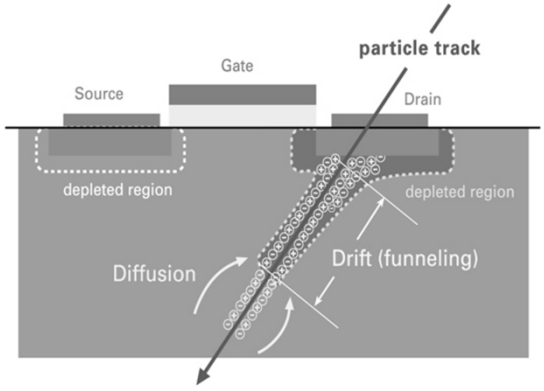

Source: Silvaco

TID

Things tunnel into gate oxides, shifting device characteristics

- Some can be detrapped by thermal energy (annealing), others can't

Affects the operation of circuits like RNGs

- See ECC issues earlier

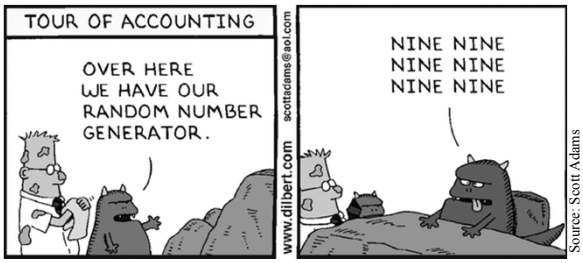

Source: Scott Adams

Testing for this is hard

- Long-term accumulation of low dose != short-term high dose

SEE

Usually soft errors

- SEU, Single Event Upset, the most common
- SET, Single Event Transient

Can be hard errors rather than soft

- SEB, Single Event Burnout
- SEGR, Single Event Gate Rupture
- Permanent
- Experienced in power MOSFETs in train engines in Europe and Japan
  - Caused by atmospheric neutrons

SEE (ctd)

Also persistent errors

- SEL, Single Event Latchup
- SEFI, Single Event Functional Interrupt

Reset by power cycling

Stuck bits aren't uncommon

- Bit(s) become stuck at some value for several seconds
- Eventually relax back to a r/w state
- RNGs and ECC again (or AES-CTR)

SEE (ctd)

Things can fail in unexpected ways

- Expose PIII and K7 to gamma source

What failed wasn't the CPU but the CPU fan

- A PWM fan-control chip in the fan motor died long before the CPU did

SEE (ctd)

During the test, all components except the CPU were heavily shielded

- CPU was raised up above the shield by a riser board

Scattering caused faults in the shielded components

> Multiple motherboards, memory modules and video cards have lost functionality in the pursuit of the total dose limit of the DUT processors
> — "Total Ionizing Dose Testing of the Intel Pentium III and AMD K7 Microprocessor"

SEE (ctd)

PIII took 100x the dose of the K7

> Intel processors verge on radiation-hardened devices
> — "Terrestrial-Based Radiation Upsets: A Cautionary Tale"

- Possibly due to Intel's bad experience with alpha-particle induced faults in 16K DRAMs in the late 1970s

Use e.g. ECC and bit-interleaving (to fortify SECDED) in caches (c.f. 21264 BTAG issues)

No visible SEUs in cache RAM

- SEUs in registers and SETs in the logic are the main issue

No known SELs

- These are rare in general

SEE (ctd)

It's clear that [Intel] are addressing an issue with cosmic rays, since they have become progressively more rad-hard over the years
— "Terrestrial-Based Radiation Upsets: A Cautionary Tale"

- Another reason why you see Intel in data centres and not ARM

IBM Power is just as careful

- ECC, parity, residue checking, …
- Tested by proton-beam irradiation

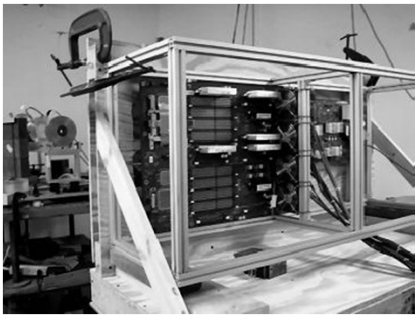

Source: IBM

SEE (ctd)

Rate of SEUs is measured in Mean Time To Upset, MTTU

Intel claim MTTU of 25 years for server-grade CPUs

- Presumably at sea level
- Possibly in a lead vault?

IBM claim "dramatic improvements in soft-error recovery"

- Compared to what?

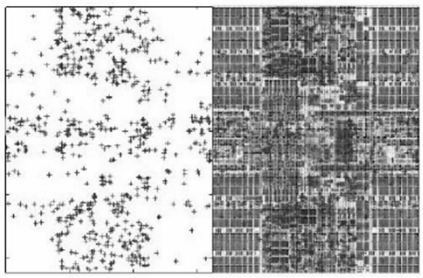

As part of the process to verify the coverage model, the latch flip distribution (left) was overlaid on a POWER6™ die photo (right).

Source: IBM

Other Effects on Computers

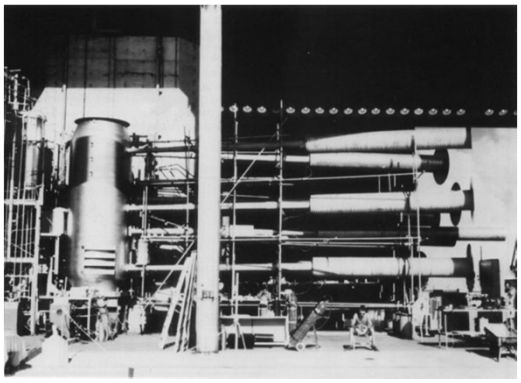

Source: NuclearSciences

Original wet hydrogen bomb employed fusion of deuterium

- Completely impractical
- Enormous size and weight (82 tons), cryogenic cooling with liquid hydrogen to -250°C, …

Other Effects on Computers (ctd)

Weaponised in Mk17 dry bomb using LiD to generate T

- $^{6}Li + n \rightarrow {}^{3}T (+ {}^{4}He + 4.7MeV)$

Other Effects on Computers (ctd)

What's the near-universal high-energy-density battery made of?

- Not that serious since quantities involved are tiny
- Resulting minute quantities of tritium and helium are trapped within the battery
- Lithium battery cases are built to deal with gas formation
- Incidentally, lithium was the first atom split in a lab in 1932 by Rutherford's colleagues Cockroft and Walton in (p, 2α) reaction

Does create a (low-energy) beta source though

- Unusually, no gamma

Other Effects on Computers (ctd)

Helium generation from alphas can be a problem for very old sealed radium sources

- 1 atom ^{226}Ra generates 5 helium atoms from alphas before the decay chain stops at ^{206}Pb
- Overpressure of 0.2 atmospheres/year for 1 gram ^{226}Ra in 1cm^3 volume

Glass-encapsulated, would rupture when enough pressure built up

- Other encapsulation was tubes sealed with a soft solder bead, which could also rupture

Other Effects on Computers (ctd)

Lots of these are still floating around in random places (orphaned leftovers from hospitals, doctor's surgeries, labs, …)

- ~65% of the US stock of ^{226}Ra is unaccounted for

We buried 10 grams of radium bromide on the 14th hole of the officer's golf course at Barksdale Airbase…

"Obvious" Solutions to the Problem

Shielding

Rad-hard components

Fault-tolerant hardware

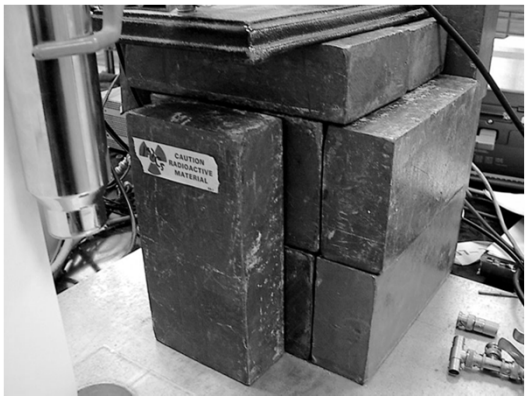

Source: Universal Medical

None of these are actually practical

- The explanation is rather long…

"Obvious" Answers (ctd)

Short version:

- Try protecting against that!

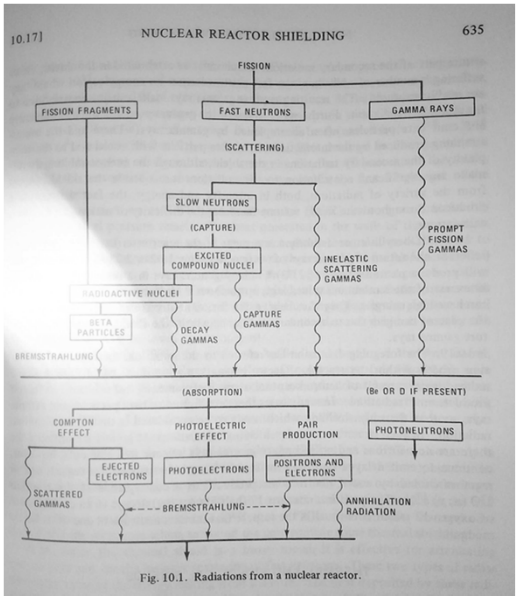

Fig. 10.1. Radiations from a nuclear reactor.

Shielding

The higher the energy and/or radiation intensity, the thicker the shielding

For photons (x-rays/gammas) we want high Z

- High atomic weight materials
- The usual suspect, lead

For neutrons we want low Z

- High hydrogen content
- Neutron shielding is actually far more complex than this, but lots of hydrogen helps

A Note on Terminology

Lots of ways of measuring dose, activity, dose equivalence, and fluence

- rad, rem, Gray, Sievert, Curie, Becquerel, $x/m^2/s$, and more
- Needs a full book to explain (e.g. "Radiation Protection", Shapiro)

For the purposes of this talk, just treat the values as magic numbers

- I'll try and give comparative values where possible

Alphas

α, $^{4}_{2}He^{2+}$, helium nucleus

- Depicted as helium with a 2+ charge due to missing electrons
- Can gain electrons from the environment and become helium

Large size and charge, high probability of interacting with something

- Anything will stop them

Still a problem with neutron activation of internal materials, which produces alphas in close proximity to the item of interest

- Also a health problem if you breathe ^{222}Rn, which decays into a pile of alpha-emitting solids that cook your lungs

Betas

β, fast electron (or positron)

Widely-available beta source: bananas

- ^{40}K produces 1.3 MeV betas

Body regulates potassium levels

- Can't overdose on bananas, it's eliminated as waste

Also found in other foods that contain potassium

- Beware garbled information on the Internet!

Betas (ctd)

Stopped by 5-10mm of shielding

However, all that energy has to go somewhere

- Electron passing near the nucleus is attracted
- Curves around the nucleus
- Loses kinetic energy which is emitted as a photon (bremsstrahlung)

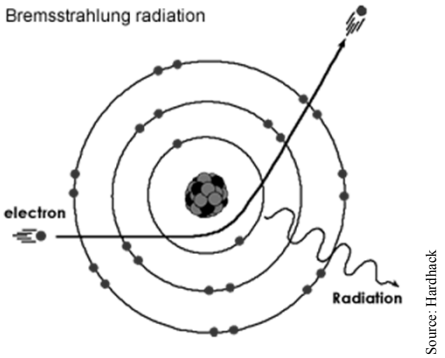

Source: Hardhack

Betas (ctd)

Industrial production of X-rays

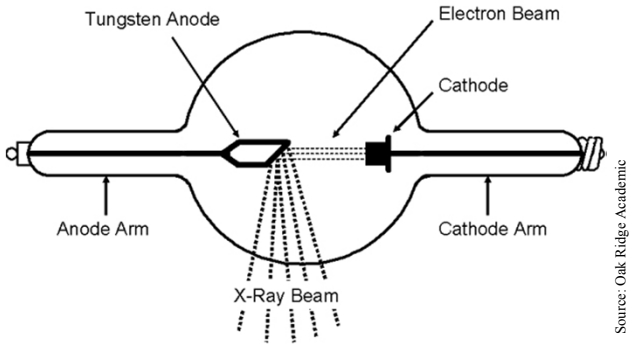

Accelerate an electron beam at a tungsten anode

- Produces X-rays of equivalent intensity

Betas (ctd)

Betas are fast electrons, lead == tungsten (an approximation)

Lead shielding of betas dumps energy as bremsstrahlung

- Same principle as X-ray tubes
- Converts betas into roughly equal-energy X-rays

Use HDPE (low-Z) to shield betas to prevent the creation of bremsstrahlung X-rays

- Back it with lead to catch anything that gets through

Gamma Rays

γ, high-energy photons (like X-rays)

Almost any interaction seems to produce gammas somewhere

- Calculus: "… plus a constant"
- Nuclear physics: "… plus a gamma"

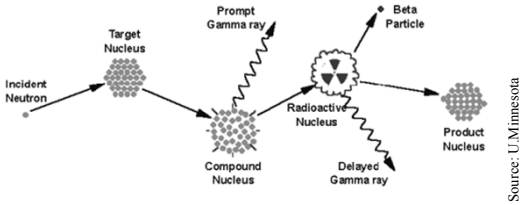

Source: U.Minnesota

Gamma Rays

Remember this diagram?

- Count the gammas

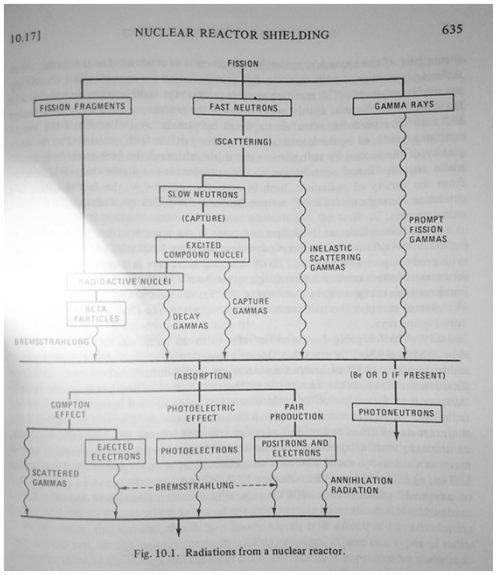

Fig. 10.1. Radiations from a nuclear reactor.

Gamma Rays

Can occur at multiple energy levels

- ^{60}Co has two gammas at 1.2 MeV and 1.3 MeV

For reference

- UVA / UVB = 4eV
  - UV penetrates poorly, so dumps all its destructive energy at once at the surface
  - Don't be fooled by the low energy level
- Medical X-ray = 75 – 120 keV

Low-activity/flux ^{137}Cs at 1.2 MeV vs. high-activity/flux ^{241}Am at 60 keV

- Calculations get complicated fast

Gamma Rays (ctd)

Half Value Layer (HVL)

- Thickness of material needed to reduce radiation by half

Energy (MeV)	**Approx. HVL in cm**			
	Lead	Iron	Concrete	Water
0.5	0.5	1.0	3.3	7.6
1.0	0.8	1.5	4.6	10.0
1.5	1.3	1.8	5.8	12.2
2.0	1.5	2.0	6.6	14.0

Example: ^{226}Ra

- 7.7 MeV α, 2.8 MeV β, 2.4 MeV γ
- HVL = 14mm lead, 70mm concrete

Gamma Rays (ctd)

Tenth Value Layer (TVL) follows from HVL

- Operation is non-linear, since radiation loses energy the further it penetrates
- Typically 3-5x HVL

Actual level depends on how much you can safely have behind the shielding

- Even TVL may not be enough

Gamma Rays (ctd)

Generally just use lots of lead

- Also good as a universal backstop to stop byproducts (X-rays, gammas) of other radiation
- Particularly necessary for neutron shielding

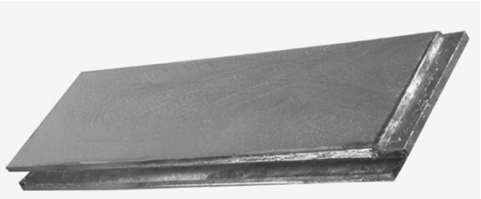

Source: Masrsmetal

Scattered photons can produce X-rays which contribute to the total dose even if the gammas don't directly cause damage ("dose enhancement")

Neutrons

What makes a reactor work

- You can't really remove them from the picture

Horrible things

- Penetrate everything
- Leave a trail of charged particles and destruction in their wake

No charge themselves, but can induce SEEs through secondary-particle production

Neutrons (ctd)

Thermal neutrons

- In thermal equilibrium with their environment

Fast Neutrons

- 1 – 10 MeV
- Low probability of interaction due to high speed
- In a reactor, too fast to cause further fission
- Slowed with a moderator, e.g. water, graphite

Neutron has no charge, so can't push through an atom's electron shield to the nucleus

- Slow (thermal) neutrons spend more time in the vicinity of the nucleus, so have a greater chance of being captured

Neutron Shielding

1. Slow the neutrons
2. Absorb the neutrons
3. Absorb the resulting gammas
   - There's always *&#^"*ing gammas…

Slowing Very Fast Neutrons

Use iron

- Inelastic scattering can reduce energy to 1MeV (standard fast neutron) in a single collision

Reactor vessels are made of 316 stainless or equivalent

- "Marine grade stainless"
- Heavily studied
- Long experience with use in reactors
- Properties well known

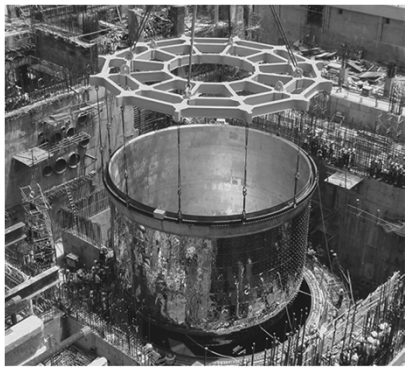

Source: Larson & Tuobro

Slowing Very Fast Neutrons (ctd)

316 stainless contains nickel, ^{58}Ni

- Under neutron bombardment, bad things happen

  $^{58}Ni + n \rightarrow {}^{59}Ni + \gamma$

- Neutron enters the nucleus and agitates the nucleons
- Excess energy is emitted as the obligatory gamma ray
- ^{59}Ni has a huge neutron-capture cross-section (13 barns)

  $^{59}Ni + n \rightarrow {}^{56}Fe + \alpha$

  $\alpha \rightarrow He$

- Also generates H by various other processes

Slowing Very Fast Neutrons (ctd)

Neutrons will also physically displace matter

- Direct collisions create primary knock-on (PKA) atoms
- These then displace further atoms, secondary knock-on atoms

Creates a displacement cascade

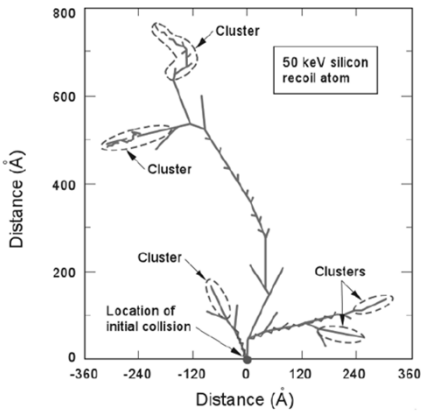

Source: JPL

Slowing Very Fast Neutrons (ctd)

Causes displacement damage in which atoms are moved into new locations

- Neutron interactions with matter require only 25 eV
- 1 MeV neutron can produce many defects due to multiple interactions before energy is lost

Stainless steel tends to fracture along grain boundaries

- These then fill up with helium from the alphas, which accelerates void swelling
- Swelling can change dimensions up to 15%

Slowing Very Fast Neutrons (ctd)

Radiation embrittlement

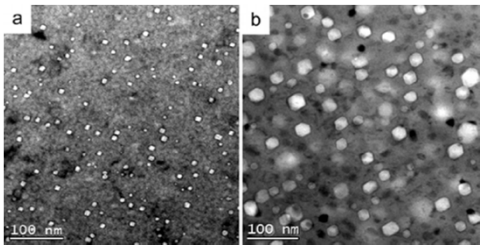

Source: Journal of Nuclear Mat.

Mitigating factors:

- Cold-working the steel gives a 4x improvement in swelling suppression
- During initial operation while the metal is still ductile it can deform if required to minimise stresses

Slowing Fast Neutrons

Use hydrogen (or something containing lots of hydrogen)

- Almost the same mass as a neutron
- Maximum energy transfer from collision with the hydrogen nucleus
- Reduces/moderates neutrons to < 1MeV

Using purely high-Z elements is no good, first collision is inelastic but the rest are elastic for which heavy elements are poor moderators

Neutron removal cross-section

- Iron = 2
- Lead = 4
- Paraffin (wax), HDPE = 80

Slowing Fast Neutrons (ctd)

Water

- ✓ Cheap
- ✓ Heat transfer medium (in reactors, e.g. BWR, PWR)
- ▪ Evaporates or drains away (LOCA)
- ▪ Heat transfer system also increases reactivity due to neutron moderation

Paraffin (wax)

- ✓ Relatively cheap
- ✓ Can be cast into any form
- ▪ Burns a bit too well

Slowing Fast Neutrons (ctd)

HDPE

- ✓ Easy to machine
- ✓ Doesn't evaporate or burn as readily as paraffin
- ▪ Expensive

Grad students

The human body is better at absorbing neutrons than a slab of lead

— "Strange Glow"

- ✓ Cheap
- ✓ Easily replaced
- ✓ Self-positioning
- ▪ Irregular density

Absorbing Neutrons

Again, hydrogenous matter works well

Cadmium works really well too

- 10,000 times the neutron capture cross-section of lead
- However, in both cases you get gammas again

Incorporate boron, which only emits a low-energy capture gamma ray

- 1600 times the capture cross-section of lead
- Much cheaper than cadmium
- Borated concrete, borated paraffin

Source: United Nuclear

Absorbing Neutrons (ctd)

Finally

- Yet again, use a lead backstop
- Stops capture gammas and various decay products from neutron activation of the preceding material

Lead for Shielding

Lead has its own problems though

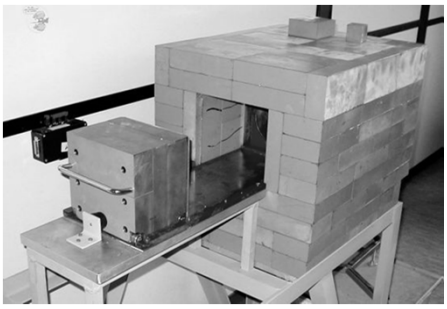

- Heavy
  - It's astounding how heavy even a moderate amount of shielding is
- Expensive
- Not ROHS compliant
  - Wouldn't want to shield your plutonium with an officially-declared Hazardous Substance
- Need to check your lead for radioactivity when you get it
  - Some lead may be, um, "recycled"

Neutron Activation (ctd)

Chinese-made solder, for example

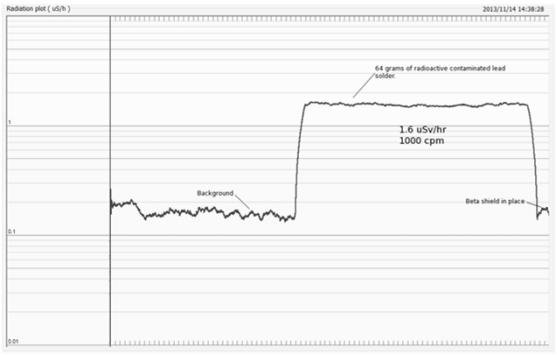

Looks like some enterprising vendor recycled some lead that should have been disposed of…

- 1000 cpm ~ avionics-level radiation

Neutron Activation (ctd)

Exactly how this came about is unclear

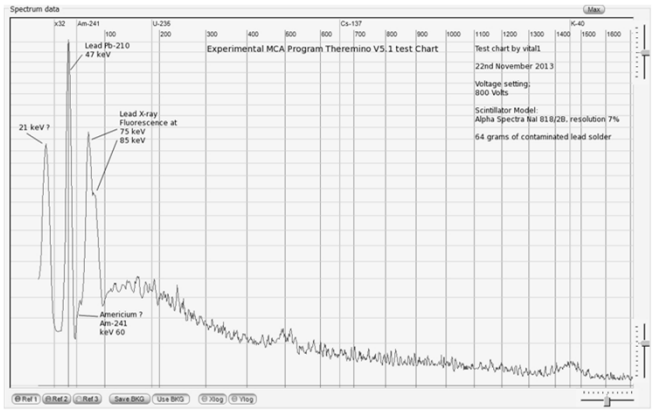

^{210}Pb in the decay chain of ^{238}U via ^{226}Ra and then ^{222}Rn, but not found in these quantities in nature

Streaming

Radiation has a remarkable ability to find its way out through gaps and cracks

- Need to get power and signals in and out of a device

Use a labyrinth to prevent scattered radiation from getting out

- Rule of thumb: Radiation must scatter at least twice to reach the entrance

For wiring, need complex zig-zag paths through the shielding

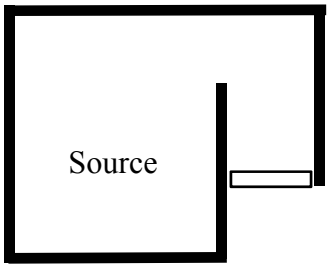

The Brute-force Solution

Concrete, and lots of it

Source: Berliner Unterwelten

Neutron Activation

Neutrons tend to make many materials they hit radioactive

- Almost all caused by thermal neutrons, can ignore for 2 MeV fast neutrons

Affect anything they come into contact with

- Plastics: Carbon, hydrogen, chlorine, flourine, oxygen
- Stainless steel: Iron, chromium, nickel, molybdenum, manganese
- Electronic circuitry: You name it
- (Chinese solder?)

Neutron Activation (ctd)

So now you have alphas, betas, gammas…

- The radiation is coming from inside the shielding

Some facilities accommodate this with sacrificial shielding, e.g. two layers of concrete where the inner layer can be removed when the facility is decommissioned

Shielding Summary

Too complicated/impractical

- Expensive, bulky, heavy, …

Needs to be managed on a case-by-case basis

- OK for fixed experiments with known parameters in a lab
- Not OK for equipment that needs to be applied/used anywhere

Rad-hardened Electronics

Military use, mostly high-altitude aircraft or outer space

- Some terrestrial as well, “harsh environments”

Combat rarely occurs in environmentally-controlled data centres

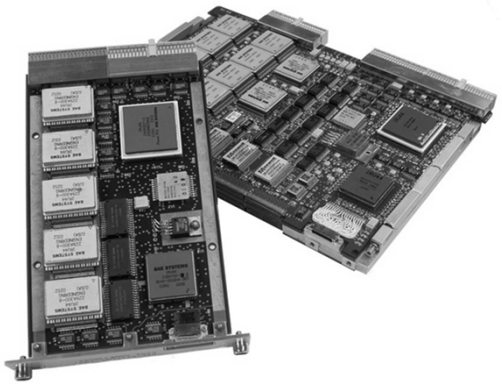

Source: Mil-Embedded

Rad-hardened Electronics (ctd)

Concern there is protons + electrons ("cosmic radiation"), not neutrons

- Strictly speaking, solar energetic particles (SEP)
- Galactic cosmic radiation (GCR) also has high-Z-high-energy (HZE) ions, which are even nastier, not considered here

1 – 10 MeV protons are stopped by just over 1mm of aluminium

- Neutrons don't even notice 1mm of aluminium

Fast neutrons are very highly penetrating […] you can look through lead shielding
— "Neutrons Provide Unique Penetrating Radiation"

Rad-hardened Electronics (ctd)

Example: Linear Tech RH1056 rad-hard op-amp

- Guaranteed by manufacturer to 100 krad ($Si0_2$), tested via ^{60}Co gamma source
- In practice typically works up to 1 Mrad ($Si0_2$)

Source: Linear

Typical COTS device = 1 krad ($Si0_2$)

- 1 krad is also the whole-body fatal dose for humans
  - See earlier note about complexities of dosage measurement
- Local doses of 4-8 krad are OK in radiotherapy

Rad-hardened Electronics (ctd)

However, device that's OK with 100 krad – 1Mrad gammas fails catastrophically with < 100 krad 200 MeV protons

Oh yeah, and rad-hard = export-controlled

Also

- Expensive
- Just a standalone op-amp when you actually need an entire ASIC
- …

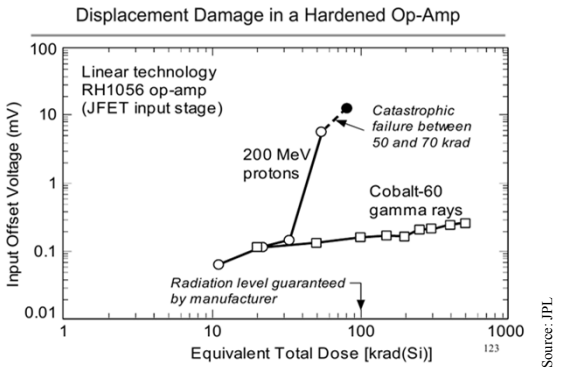

Fault-tolerant Systems

Numerous architectures designed for functional safety

- General terminology is XooY
- X out of Y units must fail for a system failure

Standard systems are 1oo1

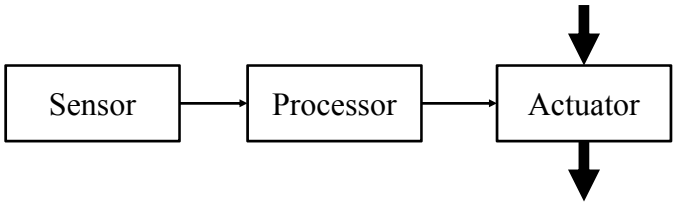

- Failure of any part causes system failure

Fault-tolerant Systems (ctd)

1oo2 for fail-safe systems

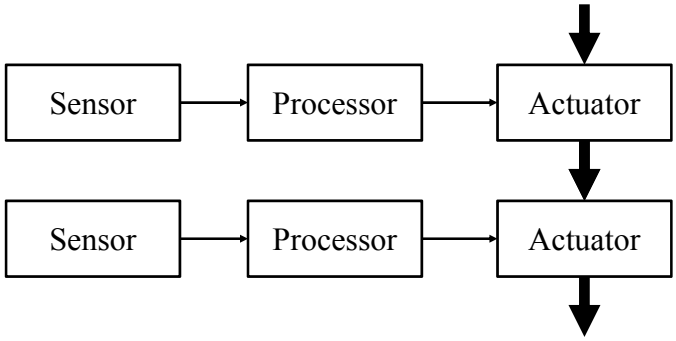

- Failure in both systems is required for an inadvertent activation to occur

Fault-tolerant Systems (ctd)

2oo2 for high-availability systems

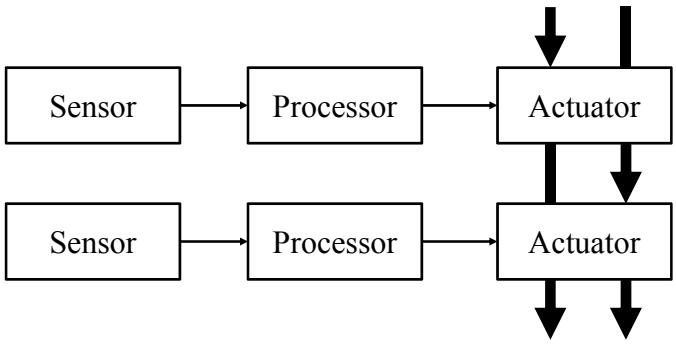

- Both systems must fail for overall failure

Fault-tolerant Systems (ctd)

Even more complex systems are possible

- 2oo3 with voting circuits

All of these (except 1oo1) require custom hardware designs

- Not practical to require this
- Can't demand completely new hardware just to accommodate an obscure crypto issue

Countermeasures Summary

None are really practical

- May be feasible, but not really practical

Need to suck it up and deal with it

Terrestrial Radiation Environments

Typical sample mix (used for rad-testing purposes)

- 50% gamma
- 25% fast neutrons (2 or 14 MeV)
  - Fission neutrons from U235 are 0.1…3 MeV, take 2 as representative
  - Fusion neutrons from D-T fusion are 14 MeV, don't really need to worry about these ones
- 25% thermal neutrons (< 10keV)

Terrestrial Radiation Environments (ctd)

Low

- 10-70uSv/hour
- 1 SEU-induced reset per week (MTTU 1 week)
- Total dose over equipment lifetime 6 Sv
- Equipment should operate for 150 months MTBF
  - That's *twelve years* continuous operation at 1 reset-inducing fault a week

Medium

- Up to 500uSv/hour
- 1 SEU-induced reset per hour (MTTU 1 hour)
- Total dose over equipment lifetime 300 Sv

Terrestrial Radiation Environments (ctd)

High/Special

- Over 500uSv/hour
- Equipment can be expected to fail within one hour unless special measures are taken
- Special-case use, e.g. short-term spikes in radiation
  - Can use shielding or other countermeasures

Total dose over equipment lifetime on a case-by-case basis, e.g. 3000 Sv

- For comparison, full-body human fatal dose is about 4000-5000 Sv

Terrestrial Radiation Environments (ctd)

One option for intermittent high-radiation operation is to temporarily shut down the system during periods of high flux ("circumvention")

- Not necessarily good for security purposes
- Attacker can cause the system to go offline for a period of time
- Circumvention results in some TID effects during high flux, but no SEE effects

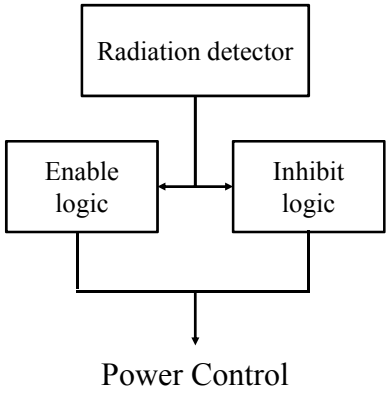

Testing with α, β, γ

α is hard to test explicitly since everything stops it

β from any natural source

γ from lab ^{60}Co source

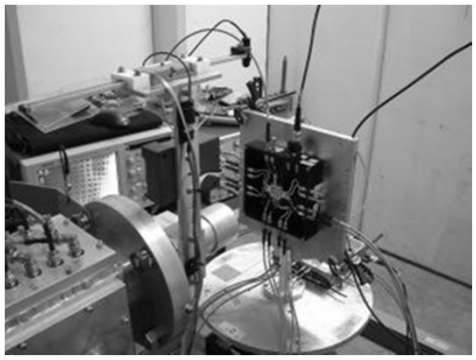

Testing with Neutrons

"Neutrons are easy, just make your own!"

- You can actually do this

OCTOBER 15, 1936 PHYSICAL REVIEW VOLUME 50

The Production and Absorption of Thermal Energy Neutrons

GEORGE A. FINK, *Pupin Physics Laboratories, Columbia University, New York, New York*

(Received August 13, 1936)

- Every school should have one!

In the small college or university laboratory, as well as in the small industrial laboratory, low-level neutron sources can be versatile tools to augment an educational or analytical program. Until

Testing with Neutrons (ctd)

We went out and got us some plutonium…

A 5-curie plutonium-beryllium source was obtained from *NUMEC*.[7] The rationale behind procurement of

[5] Kaman Nuclear Corporation, Colorado Springs, Colorado.
[6] Nuclear-Chicago Corporation, 333 E. Howard Ave., Des Plaines, Illinois.
[7] Nuclear Materials and Equipment Corporation, Appolo, Pennsylvania.

Standard disc sources today are microcuries or even fractional microcuries

- 5 Ci of plutonium is *five million times* more than this
- Oh yeah, it's ^{239}Pu, not that useless ^{240}Pu stuff

Testing with Neutrons (ctd)

Uses (α, n) reaction

- Alphas from Po, Pu, Am knock neutrons out of Be (spallation)
  - Most neutrons produced are fast neutrons
- Need alphas > 3.7 MeV, i.e. ^{210}Po, ^{226}Ra, ^{239}Pu or ^{241}Am

$^{9}_{4}\mathrm{Be} + {}^{4}_{2}\alpha \rightarrow ({}^{13}_{6}\mathrm{C})^{*} \rightarrow {}^{12}_{6}\mathrm{C} + {}^{1}_{0}\mathrm{n}$, $^{13}\mathrm{C}^{*} \equiv {}^{13}\mathrm{Be}$

Also used (with polonium as the alpha source) as the initiator for fission weapons

Testing with Neutrons (ctd)

Leave out the shielding to save cost

> Ideally, the container should be lined with cadmium sheet metal to absorb most of the neutrons which reach the outer walls, but such shielding can add up to $300 to the cost of materials. Several of the commonly used commercially produced neutron howitzers use little or no cadmium shielding.

Actually not so bad, one paper quotes 20 n/cm^2/s for both fast and thermal neutrons at 1m

- Mind you, this was based on 1960s safety levels…

Testing with Neutrons (ctd)

Still, even in 2015:

> 71864-001. The neutron source is a mixture of plutonium (^{239}Pu) and beryllium (Be) powder encased in a stainless steel capsule. It has a nominal strength of 1 Curie (Ci) (Manual, 1964). The source was made by Magna Research, and the capsule is stamped with

Oh yes, there can be gammas too

- There are always gammas

> paraffin. Because hydrogen has an appreciable cross section (0.33b) for absorbing thermal neutrons and turning into deuterium, the howitzer body glows with 2.224 MeV gammas emitted during the capture reactions (Rinard, 1991).

- 5cm of lead reduces the gammas to TVL

Testing with Neutrons (ctd)

"Neutrons are easy, just make your own!"

- Yeah, I'm going to outsource that one, thanks

What do we Want to Test

Not testing to failure, or even for failure

- Makes things much easier

Don't have to accommodate all sorts of complex effects and conditions

What do we Want to Test (ctd)

If we took this thing into the field and ran it near some kind of radiation source, what would happen to the software running on it?

- What sort of SEUs will be observed?
- Single-bit upsets?
- Multiple-bit upsets?
- $0 \rightarrow 1$, $1 \rightarrow 0$ transitions?
- Patterns of faults
- Other things

Types of SEE Faults in RAM

Lots of different fault types possible

- Expected: $0 \rightarrow 1$, $1 \rightarrow 0$
- In checkerboard patterns

  0 1 0 1 0 1 0 1 0 1
  1 0 1 0 1 0 1 0 1 0
  0 1 0 1 0 1 0 1 0 1
  1 0 1 0 1 0 1 0 1 0

  a single strike can affect adjacent cells so diagonal upsets predominate
- Entire rows can be upset if a word line takes a strike and writes a bit line value into all cells

Types of SEE Faults in RAM (ctd)

Word-line upsets can cause errors undetectable by ECC (!!!)

- Wrong memory address is read/written
- ECC is correct, but it's for the wrong data
- If multiple word lines are asserted, the bit lines will logically OR the read data

Number of bits upset depend on the strike angle

- At 90°, one- and two-bit upsets predominate
- At 45°, three-bit upsets start appearing
- At 53 and 65°, multibit/multicell upsets predominate

Types of SEE Faults in RAM (ctd)

Modern ICs (not just Intel/IBM CPUs) are surprisingly radiation-tolerant

- Expectation: Smaller feature size → smaller charge quantities → more susceptible to charged particles
- Actual: Smaller feature size → less for the charged particle to affect

Total charge trapped in oxide is proportional to thickness

- Very thin oxide
- Not much for the charge to be trapped in
- Quick recovery through detrapping (annealing)

Types of SEE Faults in RAM (ctd)

Various different reports indicating that shrinking features cause more/less problems

- Overall, it's not as bad as you'd think

The three most recent platforms exhibit lower CE [error] rates than the two older platforms […] Improvements in technology are able to keep up with adversarial trends in DIMM scaling
— "DRAM Errors in the Wild"

How to Test

Record SEUs and then inject them into a VM or binary simulator to determine effects on software

- Only need to do the physics stuff once
- Record once, use multiple times

Fault-tolerant Systems Revisited

There's a special variant that requires little to no custom work…

1oo1D

- Standard 1oo1 with diagnostic channel
- If a failure is detected by the monitoring system, halt or restart the main system

Fail-fast

- 1oo1D is pretty standard for radiation-tolerant systems
- Actually it's pretty standard for properly-designed (SCADA, not IoT) embedded in general

Fault-tolerant Systems Revisited (ctd)

1oo1D protection mechanism for code

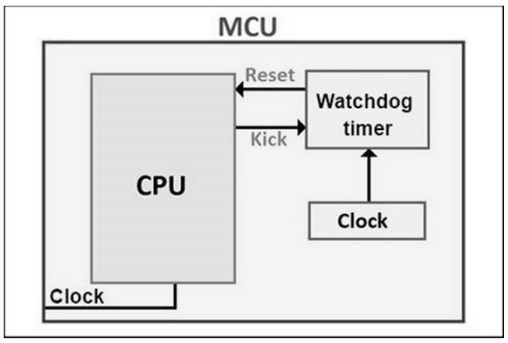

- Outright crash: Watchdog resets the system
- Non-crashing code corruption: Code (r/o) memory is scanned for corruption by D component, crash on checksum failure

Fault-tolerant Systems Revisited (ctd)

Tiny size of the D component minimises the chances of it being hit

- For hardware-based solutions, Xilinx estimates MTBF of, um, 721 years, 56,534 years, 329 years, 570,006 years, 187,256 years 75 days 3 hours 31 seconds exactly, …

Note MTBF, not MTTU

- SEUs are much more frequent
- Assumed occurring in unused cells, corrected by ECC, etc, so MTBF ≫ MTTU

Conservatively, we can multiply the MTBF by a factor of ten — "Radiation Effects and Mitigation Overview"

- "So, Bob, what do you think would be a good number?"

Fault-tolerant Systems Revisited (ctd)

For software-based 1oo1D, fault in D will usually result in a crash

- Watchdog or user restarts the device

For SEU-management purposes can be retrofitted to existing systems (if they're not already present as a standard item) without requiring new hardware designs

- Moving to new hardware is ~~almost~~ never an option
- "Here's what we've got, make it work"

Example: Aircraft Control

More than usually susceptible to SEUs due to operation at high altitude

- Computers have dual units, command and monitoring
- Command unit performs the actual function
- Monitoring unit checks that the command unit operates correctly

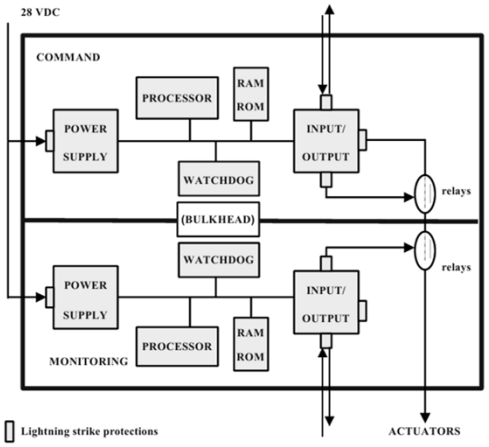

Source: Airbus

Example: Aircraft Control (ctd)

In case of discrepancy, units go into a failsafe state

- Hot-spare takes over

Total of four computers, two primaries (PRIM) and two secondaries (SEC)

- Design and software diversity: PRIM is PowerPC, SEC is Sharc, different software, compilers, etc
- Both computer teams coordinate to use different components
- SEC is simpler than PRIM
- MEL is three units operational

Example: Aircraft Control (ctd)

In standby mode, units are kept powered up to "activate potential dormant faults and isolate them"

- Opposite of practice for SEU circumvention, where units are powered off to minimise SEUs

Fault Protection

Remaining fault types (for crypto)

- Fault getting the key data into RAM
- Fault in the key data in RAM
- Fault during the crypto operation

RAM storage fault is vastly more likely then the other two

- Keyload and crypto ops take a fraction of a second
- Key sits in memory for days, months, …
- Inherent fault-resistance of CPU cores vs. RAM, see earlier slides

Fault Protection (ctd)

Example: Device outputs a signed message every six hours

Fault probabilities

- Move to RAM: 100ms (0.1s), one-off on device initialisation
- Store in RAM: 21,600s (6 hrs)
- Generate signature: 10ms (0.01s)

Time-at-risk ratios

- Signature: 1, Load: 10*, Storage: 2,100,600

Fault Protection (ctd)

Even those figures don't show the true picture

Chances of a strike on RAM vs. CPU

- Relative surface areas of the different components
- Whether the bit you care about gets hit
  - c.f. Xilinx' use of MTBF, not MTTU

CPU typically has ECC and other error management, RAM doesn't

Strike on RAM will silently corrupt, strike on CPU will often crash it

This is why you eventually need to get empirical data…

Fault Protection (ctd)

Cryptovariable ("keys") chain of custody

- Stored in encrypted form in flash
- Decrypted into RAM
- Converted from flattened stored form to internal bignum form

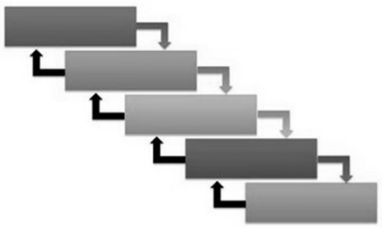

Each step needs protection

Fault Protection (ctd)

Crypto data in flash is MAC'd (cryptographic checksum)

- Decrypt + verify MAC

Conversion to bignum form is performed twice and cross-checked

- Modular redundancy

Bignum values are checksummed once loaded

Algorithm-specific validity checks on key parameters

Pairwise consistency test on loaded values

- Create signature with private key
- Verify signature with public key

Fault Protection (ctd)

Order is important

- Standard TOCTOU problem

Wrong

- Check that key form A is OK
- Convert from A to B

Right

- Convert from A to B
- Check that form A that you started with is still OK

Fault Protection (ctd)

It's really, really tricky to find every possible TOCTOU location

- Requires very careful code auditing

Example: Load and verify a key

```
loadKeyFromStorage();
if error → exit;
checkPublicComponents();
if error → exit;
checkPrivateComponents();
if error → exit;
calculateKeyChecksum();
keyOKtoUse = TRUE;
```

Fault Protection (ctd)

This isn't safe

- Can get a fault between the key validity check and the checksum calculation

```
loadKeyFromStorage();
if error → exit;
calculateKeyChecksum();    // On unvalidated data
checkPublicComponents();
if error → exit;
checkPrivateComponents();
if error → exit;
keyOKtoUse = TRUE;
```

- Order is peculiar since the data being checksummed hasn't been validated yet

Fault Protection (ctd)

Each time the private key is used

- Check key data integrity via checksum
- Generate signature
- Check key data integrity is still OK via checksum
- Verify signature with public key

Wait, verify the signature right after creating it?

Lots of fancy measures to modify RSA-CRT, and to a lesser extent ECC, signature gen. to detect corruption

- High-overhead, up to 100%, best-case is about 10%
- Very complex and difficult to get right
- May introduce additional problems, e.g. side channels

Fault Protection (ctd)

For RSA, simply verifying the signature with the public key has an overhead of 3% (Crypto++ benchmarks)

- In any case it's not really an issue, see earlier slides
- 3% overhead means you may as well do it

For ECC, the same thing has a 300% overhead (Crypto++ benchmarks)

- Makes ECC twice as slow as RSA (P256 vs. 2048, the universal standards)

Did I mention that ECC kinda sucks?

Fault Protection (ctd)

Crypto requires long-term risk planning

- What's likely to happen to this deployed system ten years from now?

For RSA, we can look at the history of fault issues

- Straight line on the graph from original attack in 1997 to now
- Break RSA-CRT, fix RSA-CRT, break RSA-CRT, …
- Doesn't really matter in any normal use case (non-dup sigs)

For ECC, the history is all over the place

- Scatter-plot of different attacks on every part of ECC
- No sign of it settling down to any particular trend

Can't really risk-plan when ECC is involved

Conclusion

We *can* make crypto function in a high-radiation environment

- Also fortifies it for use in the real world in general

Testing for effects on actual hardware is a work in progress

- Preliminary testing using purely random fault injection indicates you can make it radiation-safe

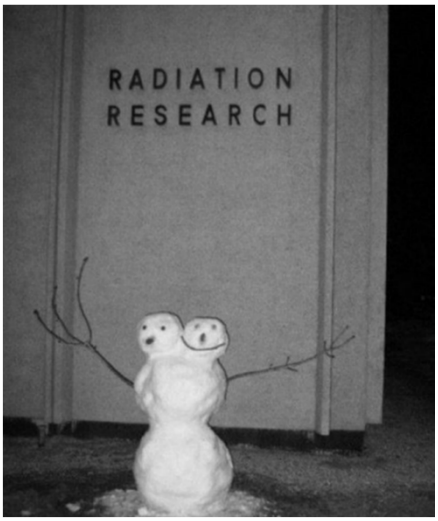

Source: The Interwebs