

“Cycles, Cells and Platters: An Empirical Analysis of Hardware Failures on a Million Consumer PCs”

Edmund B. Nightingale, John R. Douceur
Microsoft Research

Vince Orgovan
Microsoft Corporation

Presentation by Rafał Rawicki
rafal@rawicki.org

Introduction

- This is the first large-scale analysis of hardware failures on consumer PCs
- Two data sets:
 - RAC - from Windows' Experience Improvement Program
(collected from approx. 950 000 machines)
 - ATLAS - from reports sent when Windows boots after crash

Data limitations

- Only Windows crashes were reported. There is no data about unrecoverable failures or application crashes.
- Opt-in participation in both programmes.

Terminology

- TACT - Total Accumulated CPU Time
- Failures divided by type of hardware:
 - CPU and associated components
 - DRAM
 - disk subsystem

Failures are recurring

Failure	min TACT	Pr[1st failure]	Pr[2nd fail 1 fail]	Pr[3rd fail 2 fails]
CPU subsystem	5 days	1 in 330	1 in 3.3	1 in 1.8
CPU subsystem	30 days	1 in 190	1 in 2.9	1 in 1.7
DRAM one bit flip	5 days	1 in 2700	1 in 9.0	1 in 2.2
DRAM one bit flip	30 days	1 in 1700	1 in 12	1 in 2.0
Disk subsystem	5 days	1 in 470	1 in 3.4	1 in 1.9
Disk subsystem	30 days	1 in 270	1 in 3.5	1 in 1.7

Underclocking vs. overclocking

	Vendor A		Vendor B	
	No OC	OC	No OC	OC
Pr[1st]	1 in 400	1 in 21	1 in 390	1 in 86
Pr[2nd 1]	1 in 3.9	1 in 2.4	1 in 2.9	1 in 3.5
Pr[3rd 2]	1 in 1.9	1 in 2.1	1 in 1.5	1 in 1.3

	Underclocked	Rated
CPU subsystem	1 in 460	1 in 330
DRAM one-bit flip	1 in 3600	1 in 2000
Disk subsystem	1 in 560	1 in 380

Desktops vs. laptops

	Desktops	Laptops
CPU subsystem	1 in 120	1 in 310
DRAM one-bit flip	1 in 2700	1 in 3700
Disk subsystem	1 in 180	1 in 280

Interdependence of failure types

	DRAM failures	no DRAM failures
CPU failures	5 (0.549)	2091 (2100)
no CPU failures	250 (254)	971,191 (971,000)

	Disk failures	no Disk failures
CPU failures	13 (3.15)	2083 (2090)
no CPU failures	1452 (1460)	969,989 (970,000)

	Disk failures	no Disk failures
DRAM failures	1 (0.384)	254 (255)
no DRAM failures	1464 (1460)	971,818 (972,000)

Summary

System	Topic	Finding
CPU	initial failure rate	1 in 190
DRAM	initial failure rate	1 in 1700
Disk subsystem	initial failure rate	1 in 270
CPU	rate after first failure	2 order-of-magnitude increase
DRAM	rate after first failure	2 order-of-magnitude increase
Disk subsystem	rate after first failure	2 order-of-magnitude increase
DRAM	physical address locality	almost 80% machines had a recurrence at the same address
all	failure memorylessness	failures are not Poisson
all	overclocking	failure rate increase 11% to 19%
all	underclocking	failure rate decrease 39% to 80%
all	brand name / white box	brand name up to 3x more reliable
all	laptop / desktop	laptops 25% to 60% more reliable

Summary

System	Topic	Finding
cross	CPU / DRAM	dependent
cross	CPU / Disk	dependent
cross	DRAM / Disk	independent
CPU	increasing CPU speed	fail. incr. per time, const per cycle
DRAM	increasing CPU speed	failures increase per time & cycle
Disk subsystem	increasing CPU speed	fails incr. per time, decr. per cycle
CPU	increasing DRAM size	failure rate increase
DRAM	increasing DRAM size	failure rate increase (weak)
Disk subsystem	increasing DRAM size	failure rate decrease
CPU	calendar age	rates higher on young machines
Disk subsystem	calendar age	rates higher on old machines
all	intermittent faults	15%-39% faulty machines

Other interesting works

- Bitsquatting - DNS Hijacking without exploitation
Artem Dinaburg, July 2011, Raytheon Company
- DRAM Errors in the Wild: A Large-Scale Field Study, June 2009, Google

Bitsquatting

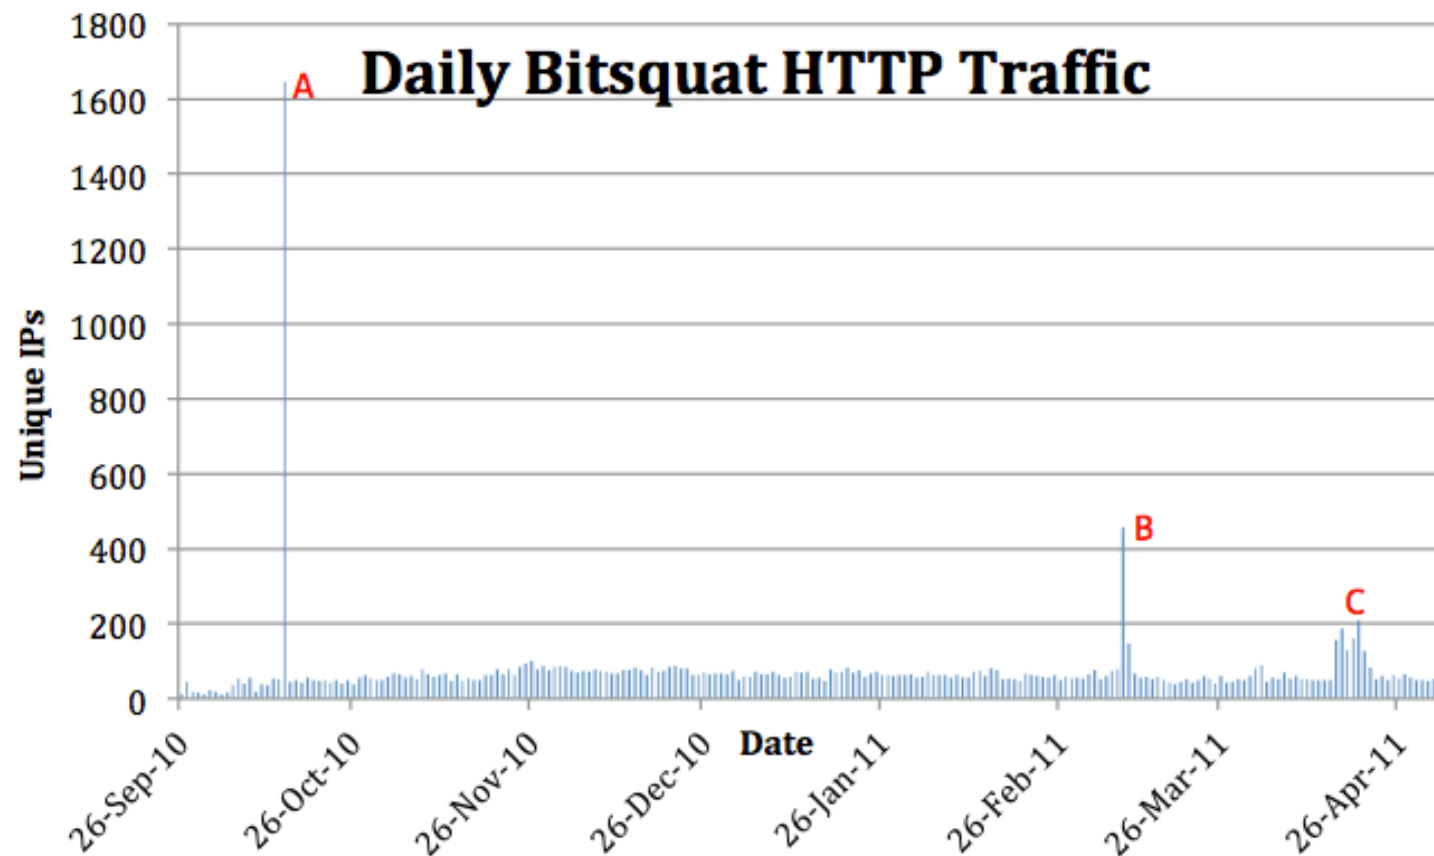
- Some domains differing by one bit from popular ones were aquired

Bitsquat Domain	Original Domain
ikamai.net	akamai.net
aeazon.com	amazon.com
a-azon.com	amazon.com
amazgn.com	amazon.com
microsmft.com	microsoft.com
micrgsoft.com	microsoft.com
miarosoft.com	microsoft.com
iicrosoft.com	microsoft.com
microsnft.com	microsoft.com
mhcrosoft.com	microsoft.com
eicrosoft.com	microsoft.com
mic2osoft.com	microsoft.com
micro3oft.com	microsoft.com
li6e.com	live.com
0mdn.net	2mdn.net
2-dn.net	2mdn.net
2edn.net	2mdn.net
2ldn.net	2mdn.net
2mfn.net	2mdn.net
2mln.net	2mdn.net
2odn.net	2mdn.net
6mdn.net	2mdn.net
fbbdn.net	fbcdn.net
fbgdn.net	fbcdn.net
gbcnd.net	fbcdn.net
fjcdn.net	fbcdn.net
dbcdn.net	fbcdn.net
roop-servers.net	root-servers.net
doublechick.net	doubleclick.net
do5bleclick.net	doubleclick.net
doubleslick.net	doubleclick.net

Table 3: Bitsquat domains registered for the experiment.

Bitsquatting

- Experiment took approx. 8 months
- “(...) a total of 52,317 bitsquat requests from 12,949 unique IP addresses.”



DRAM Errors in the Wild

Table 1: Memory errors per year:

Platf.	Tech.	Per machine				
		CE Incid. (%)	CE Rate Mean	CE Rate C.V.	CE Median Affct.	UE Incid. (%)
A	DDR1	45.4	19,509	3.5	611	0.17
B	DDR1	46.2	23,243	3.4	366	–
C	DDR1	22.3	27,500	17.7	100	2.15
D	DDR2	12.3	20,501	19.0	63	1.21
E	FBD	–	–	–	–	0.27
F	DDR2	26.9	48,621	16.1	25	4.15
Overall	–	32.2	22,696	14.0	277	1.29

Platf.	Tech.	Per DIMM				
		CE Incid. (%)	CE Rate Mean	CE Rate C.V.	CE Median Affct.	UE Incid. (%)
A	DDR1	21.2	4530	6.7	167	0.05
B	DDR1	19.6	4086	7.4	76	–
C	DDR1	3.7	3351	46.5	59	0.28
D	DDR2	2.8	3918	42.4	45	0.25
E	FBD	–	–	–	–	0.08
F	DDR2	2.9	3408	51.9	15	0.39
Overall	–	8.2	3751	36.3	64	0.22

DRAM Errors in the Wild

- ECC chips only
- Recurrence probability is consistent with “Cycles, Cells and Platters (...)”
- “A DIMM that sees a correctable error is 13–228 times more likely to see another correctable error in the same month”
- Error rate increases with age

Alpha Particles

```
1 1 1 1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1 1 1 1
1 1 1 1 0 1 1 1 1 1 1 1 1
1 1 1 1 0 0 0 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1 1 1 1
```

Figure 1

Effect of a single radioactive atom decay on a computer memory. The figure shows a readout of a portion of a 64Kb DRAM memory chip. It had been filled with all ones, and a dilute radioactive source was brought close to it. About one radioactive fragment per minute hit the chip (the source emitted alpha-particles). By observing a constant readout of the memory, it was found that a single alpha-particle could cause four memory cells to change their content from a one to a zero. From C. K. Chou, IBM Poughkeepsie, 1979 (unpublished work).

Thank you