

Re: [BugTraq] Peter Gutmann data deletion theory?

Source: <http://www.derkeiler.com/Mailing-Lists/securityfocus/bugtraq/2005-07/0394.html>

From: Robin Whittle (rw_at_firstpr.com.au)

Date: 07/22/05

Date: Fri, 22 Jul 2005 13:03:18 +1000

To: bugtraq@securityfocus.com

Peter Gutmann's 1996 paper is at:

Secure Deletion of Data from Magnetic and Solid-State Memory

http://www.treachery.net/~jdyson/infosec/secure_del.html

I will discuss four types of memory: Magnetic media (tape, hard-drive, floppy etc.), Static RAM, Dynamic RAM and FLASH (or EPROM or EEPROM).

Magnetic media typically involves a flat recording surface, with a read head, write head and sometimes an erase head before the write head. The width of the read head may be less than that of the write head, so that slight misalignments in reading or writing still cause the read signal to be picked up from within the wider written band.

Magnetic media probably has some kind of depth effect too – it is a three-dimensional object which is magnetised in one step and read in another. Unless all trace of previous writes can be removed, some remnant of the original signal will remain. To what extent this can be read – either by the standard read system or by fancy forensic techniques – would vary enormously.

Any recovered original signal will be only a small fraction of the actual read signal. The rest will be noise, the last recorded signal, and remnants of other recordings in the past. There needs to be a certain signal-to-noise ratio before data can be recovered in any meaningful manner.

I can't imagine how these forensic techniques work with modern hard drives, where data is packed so densely and recorded with highly complex encoding systems.

However, to ensure the destruction of previously written data, I think its easier to physically destroy the media than do enough research to prove beyond doubt that there is no possible recovery technique.

SecurityFocus Bugtraq: Re: [BugTraq] Peter Gutmann data deletion theory?

Probably what I have written above also applies in principle to optical media as well.

Semiconductor memory (and core memory) is completely different. There is no three-dimensional recording medium – so there is no misalignment between conventional read and write operations. There is a single quantity which is used to encode a bit of information – the voltage of a capacitor (Dynamic and Flash RAM) or of one side of a flip-flop (Static RAM).

A Static RAM cell consists of a four transistor flip-flop with two additional components, either transistors or resistors, as loads for each side. Each side is an inverting amplifier, usually made with an N- and a P-channel Field Effect Transistor. Reading involves switching one or both of the sides of the flip-flop (using additional transistors) to bit lines, and then reading these lines. Writing involves a similar connection to the bit lines, but then forcing the state of the flip-flop to one state or the other.

At the end of a write operation, the flip-flop is constantly powered and is free from external intervention. Each inverting amplifier constantly inverts its input and forms the input to the other. This constant amplification would quickly (fractions of nanoseconds) move the voltage of one side towards ground and the other towards the power supply voltage. I can't imagine any detectable difference in the physical state of the flip-flop due to its previous state surviving more than tiny fractions of a nanosecond. Since all these flip-flops are operational as soon as power is applied, any charge states after the chip was turned off would soon be over-ridden by the amplifying nature of the flip-flops. The question is whether the chip could be powered up in a way which would amplify slight charge differences remaining in the chip since it was last operating. Such differences would need to be significant compared to the natural bias in each cell towards powering up one way or another, due to atom-scale differences in dimensions of channels, gates and doping. This seems unlikely to me, but I am not a semiconductor engineer.

I recall reading, via some third-hand account, of some very old memory chips from the 1970s which could have their internal structure changed by the sub-micron voltage gradients over long periods of time spent with a particular data pattern, and this could be perceived in the power-up state of each bit, as the usual randomness was skewed by these particular changes.

Peter's paper discusses changes to the device due to repeated writes of 1 and 0. Only a semiconductor engineer would be able to advise how valid this theory is, and whether any such changes could possibly be detected by powering-up the chip in a special way, or by some other external test mode or invasive microscopical technique.

Dynamic RAM involves a single transistor connecting a single capacitor to a single bit line, to which many other capacitors could also be connected – but only one at a time. Reading involves measuring the voltage of the bit line to discern whether it is high or low. In practice, this involves a non-inverting amplifier connected to the bit line, with its output connected to its input, and therefore to the bit line. There is a pre-charge system to bias the amplifier in the middle of the range of voltages before connecting the bit line and turning it on, so that, for instance, in a 5 volt DRAM cell (now they are all 3.3 volts or less) any bit-line voltage above 2.5 volts will be amplified up to 5 volts (in a nanosecond or so) and any voltage below this will likewise be amplified to 0 volts.

Thus a read involves a "refresh" – forcing memory cells which are deemed to be "low" to (or strongly towards and usually very close to) 0 volts and those which are deemed to be "high" to 5 volts.

Since capacitors leak, DRAM systems repeatedly read and refresh all the capacitors, such as every 16 milliseconds. (Dipping at random into a 1990 Hitachi memory data book. Nonetheless, I recall interrupting refresh in a Z80 CPM system and having the data, or at least enough of it to run the system, survive for 10 seconds or so.) Each such refresh or read effectively erases any detail of the prior voltage of the capacitor. There might be some extremely small difference in the final voltage of the capacitor in the following situations:

- 1 – Initially 5 volts, then written to "low", so almost 0 volts, with just a little of the residual charge, due to limited resistance of the bit line and the limited time of the write cycle.
- 2 – Initially 0 volts then written or refreshed to "low". This would be effectively 0 volts – there is no other charge for there to be a residual.
- 3 – Initially 2.4 volts, so deemed to be low, but refreshed or written to low. (There's no reason in ordinary operation for a capacitor to every be at this voltage, unless it was left to leak due to there being no refreshes for a long time).

So I can't see how any trace of previous states could possibly be measured, including by directly probing the capacitor by some means – rather than by using the sense amplifiers, which are only making a decision about whether the voltage is above or below 2.5 volts – after one or more read, refresh or write cycles.

Perhaps keeping a DRAM chip for weeks, months or years with a particular state of data, as may well happen, could result in long-term changes in the exact physical nature of the semiconductor material, insulators etc. of the capacitor and its immediate surrounds. However, unlike Static RAM, I can't see how this could be measured (except by some direct probing of the chip, which seems highly impractical, since such probes

SecurityFocus Bugtraq: Re: [BugTraq] Peter Gutmann data deletion theory?

would be noisy compared to the slight changes) because the base state of the capacitors at power up is 0 volts or close to it, and the sense amplifiers can only discern fine differences in voltage around the 2.5 volt range.

So I think that short term storage of data in Static or Dynamic RAM leaves absolutely no detectable trace after it has been over-written. I can't imagine how long-term storage patterns could be detected in DRAM and I think it is probably impossible, by even the most heroic means, to discern any such long-term patterns in SRAM.

There are two classes of data recovery here:

- 1 – Turning the chip on in a special way to discern its state when it was turned off.
- 2 – Trying to sense long-term changes in the device to find out what data it has stored for long periods of time.

The first approach might involve removing the chip from its system by desoldering it (or maybe unplugging a SIMM) or by cutting the PCB and making low-temperature bond connections to the chip to it in a forensic test rig. Alternatively, the device may be left in the system and powered up with different signals being forced onto the memory chip to stop the usual system operations, which usually include erasing DRAM or SRAM (except for battery backed up SRAM, in which case the previous data is easily readable).

So maybe someone could build a special system to plug a DRAM SIMM into in order to detect tiny charges remaining in the capacitors from the state the chip was in when turned off. There could be many reasons why this is impossible, not least the highly complex nature of modern DRAM chips, with their fancy protocols, compared to the much more direct nature of chips from the 70s and 80s.

Flash, EEPROM and EPROM memories also use a single capacitor, but this is not connected or sensed directly by using a transistor to connect it to a bit line. The aim of data recovery in these cases is to find what data was stored in the chip prior to the data which was most recently written. Flash memory has block erase modes, so it is possible that an area has just been erased, without new data being written.

The capacitor of each memory cell is a small island, typically of silicon, entirely surrounded by quartz or some other insulator. Its voltage (Flash and EEPROM) is raised and lowered by high voltage temporary breakdown (quantum mechanical tunnelling of electrons) through the insulator. The capacitor's voltage is sensed by it part of a gate of a MOSFET – so it can be read without any current entering or leaving.

EPROM capacitors are charged and read as described above, but they can only be discharged by illumination with short wavelength UV light, which gives individual electrons enough energy to tunnel through the

capacitors insulator.

I expect that in all these cases programming and erasure is not an exactly complete process – that some trace would remain of the previous state. To what extent this could be used it hard to say, but perhaps there is a way in some cases, for instance with a FLASH or EEPROM memory. Lets say the chip was programmed with some data we wish to recover. Then lets say it was all erased to "1". Now we power up the chip and try to discern individual voltages of each cell. Maybe there's a way of increasing the power supply voltage so that the sense amplifiers are discerning fine voltage difference in the range of the normal "1" voltage. (Flash, EEPROM and EPROM typically use internal or external power supplies to create voltages much higher than the standard power supply rails used for reading, so the exact voltage of the capacitor, may be rather high. What counts is the much higher power supply voltage of the cell or read amplifier which is needed to read a "0" when the cell is actually programmed as a "1". Then by changing the power supply voltage slightly, multiple reads can be used to finely measure the supply voltage at which this cells reads as a 0 instead of a 1. Maybe we could discern some variations which would have something to do with the original data. I can't see how this could work if the capacitors were written to 0 volts.

Unlike DRAM and SRAM, Flash etc. has no continual or repetitive amplification function. Traces of previous charges may remain, but I doubt they would be recoverable, except perhaps by the most drastic forensic techniques – and even then, each write or erase operation would reduce the remnants further still.

– Robin <http://www.firstpr.com.au>