Destroying Flash Memory-Based Storage Devices (draft v0.9)

Dr. Steven Swanson
Director, Non-volatile Systems Laboratory
Department of Computer Science and Engineering
University of California, San Diego
swanson@cs.ucsd.edu

We examine the problem of mechanically destroying NAND flash-based storage devices (as commonly found in USB "thumb" drives, commercially available solid-state disks, and portable electronics such as cell phones, tablet computers, and media players) with the goal of making it prohibitively difficult to recover data from the device after destruction. We present an analysis of the minimum particle sized required to achieve this goal.

In this report, the size of a particle is measured as the length of the particle's longest linear dimension in any direction.

The necessary particle size depends on the sophistication of the "adversary" that is expected to attempt to recover data from the destroyed device. We present analysis for the particle sizes needed for three different adversaries:

1. The *typical* adversary is assumed to have access to a well-equipped electronics lab.
2. The *sophisticated* adversary has substantial but limited funds and technical expertise.
3. The *worst-case* adversary has almost unlimited funds, technical expertise, and time.

Protecting Against a Typical Adversary

The least sophisticated way to extract data from a damaged SSD is to remove an intact flash device from the SSDs printed circuit board and extract data from it using a flash reader. We have demonstrated this technique in our lab using equipment costing less than $1000. A competent electrical engineer (e.g., with a 4-year degree) could easily implement it.

This technique requires an *intact* flash device, so guarding against it requires that at least the packages for all the flash chips in the SSD be substantially damaged.

Currently available flash devices come in range of packages. The table below summarizes the dimensions of each (Data from Open NAND Flash Interface standard version 2.3 and measurements in our lab):

Package Type	width (mm)	height (mm)
LGA-52	14	18
BGA-63	10	10
zBGA-100	12	18
TSOP-48	18	12

Table 1: Flash device package sizes.

The smallest package (BGA-63) is 10mm on a side. This means that particle sizes less than 10 mm will ensure that no flash packages remains intact. To be conservative, we recommend 75% of this value or 7.5mm.

However, the BGA-63 package is not very common. The other packages are more common, and would allow larger particle sizes without compromising safety. The difficulty is in knowing what package a particular SSD uses without opening it. The conservative option is to assume it is BGA-63.

Since these package dimensions are standard, they will not change over time, so the minimum particle size will remain constant so long as new, smaller packaging standards do not enter wide use. To ensure that this guidance remains up to date, users should be in communication with flash manufacturers and standards bodies to keep abreast of any new packages that become common.

Protecting Against a Sophisticated Adversary

We assume that the sophisticated adversary has access to an extremely well equipped laboratory and is capable of "deprocessing" flash devices to remove the ceramic or plastic package that protects the chip and supports the electrical connection via the package's pins. We further assume that the adversary can attach new bonding wires to the chips pads, effectively replacing the chips package and pins. This capability has been demonstrated by other researchers in other contexts, and the "deprocessing" step is available for hire commercially.

With this capability, it would be possible to remove a flash die from a damaged package and effectively repairing or replacing the damaged package. Once this was complete, extracting data from the device would be relatively easy. Since the flash packages are often larger than the chips inside, a smaller particle size may be required.

To prevent recovery by this type of adversary, the particle size must be small enough to ensure that the flash chip itself (rather than just the package) is broken

into one or more pieces. A survey of flash devices produced in the last 5 years shows flash die sizes are relatively constant and range from 116 to 230 mm^2 with the shortest dimension ranging from 10 to 15mm. To be conservative, we would suggest targeting a particle size 1/2 of this size so reducing SSDs to particles between 5 and 7.5mm in size is sufficient.

Protecting Against the Worst-Case Adversary

To protect data from the worst-case adversary (i.e., one with nearly unbounded amount of time, money, and expertise) we assume that we must prevent data recovery via techniques that are known to work as well as those that may (eventually) be possible.

Flash arrays are composed of many pages. Depending on the manufacturer, the pages contain between 2048 and 16384 bytes. The bytes in a single page are typically consecutive bytes from a single file. Because of the complex remapping algorithms that SSDs utilize the data on consecutive *pages* is not necessarily related, although they may be.

Our analysis assumes that it is sufficient to guarantee that no single page remains intact. That is, that after destruction, the contents of each page should be spread across at least two fragments of the flash chip. This means the acceptable particle size after destruction should be less than the maximum dimension of a single page. In this analysis, we aim for particles that will contain no more than one half a page worth of data. Therefore, the minimum particle size is one half the maximum dimension of a single page.

The number of bits in a page and the "feature size" of the chip determine the physical size of a page. The feature size is measured in nanometers (nm) and is the length of the smallest feature that the chip's manufacturer can create in a given lithographic chip manufacturing process. Manufacturers use the feature size to differentiate between different manufacturing process generations. For instance, a "22nm flash chip" is manufactured in a process with a 22nm minimum feature size. We will refer to the minimum feature size as *F*.

Flash organizes the bits within pages by interleaving the bits from two pages in a single row. A physical structure that measures 2F on a side stores a single bit from a page. The figure below illustrates:

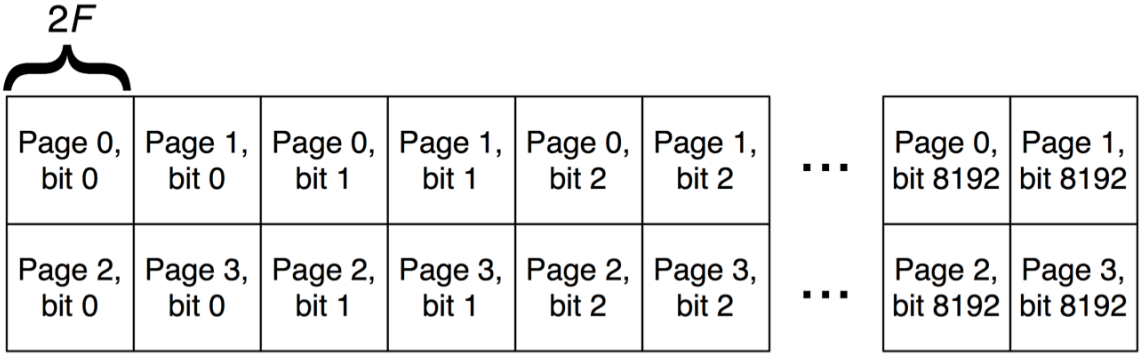

As a result the width of a single page is given by twice the page size in bits multiplied by $2F$ divided by 2 (since we want particles that contain no more than one half a page worth of data).

Using this relation, we can calculate the maximum allowable particle size for a particular page size and process generation. The table gives the values for current and projected several generations of flash memories and several common page sizes (projects from the International Technology Roadmap for Semiconductors 2010 update). Flash manufacturers are currently use 19nm feature sizes, which is ahead of the ITRS projections.

Allowable Particle Diameter by Process Generation and Page Size (in mm)

	Process Generation						
Year of Introduction		2011	2013	2015	2018	2020	2022
Feature Size	65nm	24nm	20nm	18nm	13nm	11nm	8nm
2 KB/page	2.1	0.7	0.6	0.5	0.4	0.3	0.2
4 KB/page	4.2	1.5	1.3	1.1	0.8	0.7	0.5
8 KB/page	8.5	3.1	2.6	2.3	1.7	1.4	1.0
16 KB/page	17.0	6.2	5.2	4.7	3.4	2.8	2.0
32 KB/page	34.0	12.5	10.4	9.4	6.8	5.7	4.1

Table 2: Maximum allowable particle sizes for the worst-case adversary.

As process technology improves, the allowable particle size drops. However, at smaller feature sizes, manufacturers tend to move toward larger pages, potentially increases the allowable particle size. However, there is not a consistent "rule of thumb" that relates the two.

Conclusions and Limitations

Our analysis shows that for all but the most well-funded, skillful, and determined adversary a particle size of 5mm will ensure that data is not recoverable from the flash chips inside an SSD. If more information is available about the particular flash device or packaging standard the SSD uses larger particle sizes may be acceptable as well. However, reliably determining that information on a per-SSD basis is probably impractical in practice.

For the "worst case" adversaries, much smaller particles are required to prevent recovery and the particle sizes decreases with advanced in flash manufacturing technology. Currently available SSD will require reduction to particles with maximum diameters of between 0.5 and 2.5 mm, and future SSDs may require particles as small as 0.2mm.

The techniques we considered in this analysis necessarily limit its scope. There may be other techniques for recovering data from flash devices that we have not considered, and protecting against these attacks may require smaller particle sizes. However, the above analysis for the worst case adversary, in particular, makes relatively few assumptions about the types of attacks that adversaries can mount and relies instead on the physical characteristics of the flash chips themselves and how they organize data. As such, we think it is unlikely that techniques exist or could be developed that would be effective at recovering data from SSD reduced the smallest particle sizes we suggest.