Column: The Physics of Digital Information[1]

Fred Cohen

One of the interesting questions underlying creating a science of digital forensics is what form that science might take. At the end of the IEEE Oakland Conference in 2011, I presented some of the underlying questions and identified my approach and why I favored it.[2] My basic position is that science is about causality and a scientific theory should require that cause(C) produces effect (E) via mechanism M (written $C \rightarrow^{M} E$). The scientific method then identifies the criteria for rejecting (refuting) or accepting (for a time) a scientific theory.

The method mandates that someone proposing a theory hypothesize a specific $C \rightarrow^{M} E$ by identifying specific sets of causes, mechanisms, and effects. The term specific is intended to indicate that C,M, and E are identifiable by independent third parties and properly differentiated to a defined degree of precision and with a defined degree of accuracy in terms of things that can be measured. Experiments are then defined so as to seek to refute the hypothesis by performing experiments that will produce outcomes outside of the defined results of the theory. As and if these hypotheses are found unrefuted time and again, they become part of the theory of the day,. Always subject to refinement and potential refutation.

Of course a classic example of a scientific theory was the theory that the world is flat. The experiment; keep sailing west - if you come back from the east, the world is round. This experiment was repeated many times, confirmed again and again as lots of ships didn't come back. So many didn't come back, that this became a valid scientific theory – until the first one made it around. Then it was refuted and replaced with another theory – the world is round. But even today, the flat-world theory use useful for special cases. For example, buildings are not designed based on the curvature.

Interestingly, sciences like physics and biology came to be in different ways than the current science underlying digital systems. In particular, physics and biology both come about based on observations of nature, theories (e.g., Newton, microorganisms), testing and generalization, the ability to build reliable systems within defined limits, problems resulting in refutation, new theories (e.g., quantum, genetics), more testing and generalization, and the

1 This editorial piece is extracted and modified from F. Cohen, "Digital Forensic Evidence Examination", ASP Press, 2011

2 F. Cohen, "The need for and progress in science for information protection and digital forensics", IEEE Oakland Conference – May 25, 2011. Available online at http://all.net/Talks/2011-05-25-IEEE-SecurityScience.pdf

ability to again build reliably in more circumstances and to tighter limits. Digital science has proceeded differently. It has gone from mathematical theories to complexity issues, some models were built, but they were problematic, and thus we still cannot build such systems very reliably (at the level of systems we build today). There has been no higher level theory similar to Newtonian physics or the theory of cells, and here we are.

Other approaches to protection science include archeology, which is the NSF's current thinking on digital forensics, engineering, which has continued to fail over a long time, and social science approaches, which hold promise for many but not all aspects of the area, largely because the statistical approach is problematic in the digital domain.

My approach has been to try to build out a middle level theory for digital forensics – a physics of digital information. Full details of the current version of this physics are available online[3], and a short summary of a few key points of difference between the digital space and the physical space is all that can be provided here to give a sense of where it is going.

In keeping with previous results suggesting that consensus around common definitions for the field of digital forensics do not exist,[4] each item discussed will start with a loose definition and example.

Definition: Finite granularity means that, at the end of the day, you can only get so small, and no smaller. Finite granularity in space means that there is a minimum size (i.e., the bit) and finite granularity in time means that there is a minimum time unit (i.e., the number of bits that represent a time unit, or alternatively, the maximum clock speed of the mechanisms generating the bits).

In terms of digital evidence (not the physical media that may contain it), no matter how much of it you have, there is no unit of that evidence smaller than a bit, and no matter has much precision you make measurements with, the times associated with those bits cannot be infinitesimally small. Thus:

Digital evidence is finite in granularity in both space and time.

Compare this to the physical world where, according to the current widely accepted theories of physics, we don't have a defined limit on the smallest unit of space or time. This is because we generally consider space and time to be continuous.

Definition: Observation of digital information means to be able to determine

3 http://infophys.com/

4 F. Cohen, "Toward a Science of Digital Forensic Evidence Examination", IFIP TC11.8 International Conference on Digital Forensics", Hong Kong, China, Jan 4, 2010 (Keynote), Published in Advances in Digital Forensics VI, pp17-36, Springer, ISBN#3-642-15505-7, 2010.

what binary value each bit of the information has, while alteration is the changing of a bit value to the other bit value (digital information by definition only has two different values for bits).

At the level of digital evidence, when reading bits from a disk, the bits on the disk are observed by the mechanisms of the disk drive. When writing to a disk, some of the bits on the disk may be altered. Of course the specific technology may physically alter the stored representation of the bits on every read, for example by reading and then rewriting the original bits in a different physical location or rearranging the coding of the bits as stored. But this doesn't change the digital value represented. Thus:

Normally, we can observe digital information without altering it.

In the current widely accepted theory of physicals, this is not true at the lowest level of granularity. Every observation alters the thing being observed.[5] Of course looking at paint doesn't change it from blue to red, and this "observer effect" is not normally noticeable. In the case of bits, since they are engineered and stored and/or coded so that altering a small portion of their physical state does not alter the observed value at the level of the bit, we don't have the same problem. If this were not true we would have a very different field.

Definition: Duplication of digital information means to make an exact copy of the bit sequences comprising the digital information. Removal of digital information means making alterations so that the original information (i.e., the original bit sequence) is no longer present where it originally was present before.

For example, a duplicate of the bit sequence 1 0 1 will also be the bit sequence 1 0 1. The removal of the bit sequence 1 0 1 from a location would mean that the bit sequence at that location was no longer 1 0 1. Thus:

You can usually duplicate digital information without removing it.

Compare this to the physical world where, if you "take" something, like an original painting, it is no longer there. Duplicates at the level of a painting are always differentiable from originals at some level of physicality, and thus forgery is never "perfect" in the sense of impossible to discern from the original. Of course figuring out which came first is another issue...

Definition: Transfer is a concept in which two objects coming into contact with each other each leaves something of itself with the other. This is generally

5 Heisenberg, W. (1930), Physikalische Prinzipien der Quantentheorie, Leipzig: Hirzel English translation The Physical Principles of Quantum Theory. Chicago: University of Chicago Press, 1930.

attributed to Locard,[6] even though the more precise definition of the term used here appeared later.

For example, when a shirt rubs up against a sharp object, some shards from the object may get transfered to the shirt while some fibers from the shirt may get transferred to the sharp object.

Digital evidence is normally not transfer evidence.

In physical evidence, at some level of granularity, this is almost certainly true – subatomic particles will likely transfer between objects, even though at that level, they are indistinguishable from each other. But digital evidence does not come about through such a process. While we may think of a file transfer as the same thing, it is not. It is the duplication of bit sequences through a process that produces a bit sequence elsewhere without the movement of any of the physical elements storing the original bits. This may be clearly seen if the original was stored on optical media and the result of interaction produced a bit sequence on magnetic media. None of the light and dark spots were placed on the magnetic media.

Definition: Trace evidence is evidence that is produced as the result of a process, so that the presence of the evidence is consistent with the execution of the process.

When a pen writes on paper, the indentations in the paper resulting from the writing are traces of the process of writing. Similarly, when a computer program produces bit sequences that are stored on a disk, the bit sequences are traces of the execution of the computer hardware and software that produced them. Thus:

Digital evidence is normally trace evidence.

Physical evidence is also normally trace evidence as well in that it is produced by some physical mechanism, and the presence of the evidence is consistent with some set of physical processes.

Definition: Latent evidence is evidence that cannot be directly observed by human senses.

DNA evidence, for example, is normally latent evidence in that the DNA has to be processed by a mechanism to produce a result that can be examined by human senses. So are most fingerprints, many tool marks, and many other sorts of evidence. And of course:

Digital evidence is normally latent in nature.

---

6 E. Locard, "The Analysis of Dust Traces", Revue International de Criminalistique I. #s 4-5, 1929, pp 176-249, (translated into English and reprinted in 3 parts in A, J. Police Science, 1930 in V1#3, May-Jun 1930, pp276-298, V1#4 Jul-Aug 1930, pp 401-418, and V1#5 Sep-Oct 1930, pp 496-514.)

In fact, almost no digital evidence can be presented in its original form (other than the placement of a disk drive in front of the jury). It has to be turned into something humanly observable, like a printout or a display on a video screen. And this means that the tools used to turn the evidence into something people can observe has to be validated and shown reliable in order to pass muster in court.

Definition: Computational complexity is the number of low-level computing operations required in order to perform an algorithm which processes information.

As we have all seen, computers are not infinitely fast. Different operations take different amounts of time. It takes more computer time to determine the best possible route from house to house throughout a city that to find any route that passes by all of those same houses because the former problem is more complex than the latter.

Computational complexity limits digital forensic analysis.

The time limits on examination of evidence means that certain techniques that take a great deal of computation cannot practically be applied in most cases. Even with parallel processing and cloud computing systems, this remains true. For this reason, we often use approximate algorithms or methods that are different from what might be used in other fields.

Definition: Consistency between two or more things means that each is the way you would expect it to be if the other ones are the way you observe them to be.

If you see a black box and someone else viewing the same object under the same conditions states that it is a white sphere, your observation is inconsistent with their statement. Similarly, if a sworn statement states that a particular file was created at 10AM on a particular day in a particular place, and the metadata for the file indicates that the same file was created at a different time on a different day, the sworn statement is inconsistent with the metadata. As a fundamental of digital forensics:

What is inconsistent is not true.

or in other words:

Inconsistent things cannot all be true.

If the metadata disagrees irreconcilably with the person, they can't both be right. Each actual event that takes place takes place within a bounded area of space and time, and the same event cannot take place in two mutually exclusive times and places. But this doesn't mean that, for example, the metadata associated with the file reflects the same event that the person was talking about. Thus inconsistency must be examined closely in order to make sure that there isn't some consistent explanation.

Notionally, these definitions and explanations may be viewed as the start of part of a physics of digital information. This physics may be seen as a layer above the mathematics that underlies computers.

As a physics, these rules applied in context explain and predict behaviors of digital systems. If you agree to the details of the physics above, it means that you can make forensically sound copies of digital data that, at the level of digital evidence, are every bit as good as the original. If you don't agree, then such copies cannot be forensically sound in that they cannot maintain the digital properties of the original that are critical to its use, unless you can come up with some other physics that will gain acceptance and under which they will be admitted.

Assuming you buy into this view of physics, the physics of digital systems is different than that of the normal physical world, and that means that the way we normally view the natural world is off kilter from the nature of the digital world in at least these ways. If you stole that information, how come I still see it there? If a copy of a piece of paper is not as good as the original for evidence, how come a copy of the bits on the computer, even though they are on a different sort of media, are as good as the original bits?

A final comment

The issue of consensus around these basics of digital forensics is not fully settled today. But at least for the items identified in this article, consensus is getting pretty close. At the end of the day, consensus is what the scientists, practitioners, and researchers in the field come to agree on. It's not what I say – it's what you say – and do – as a group.

I urge you to take some time and review the details of the physics of digital forensics and draw your own conclusions. Read the chapter cited above, comment on it, prove it is wrong if and where it is, show its limits, and move the field forward. Hopefully, over time, we will come to an understanding that will help us all move forward and advance the field together.

At least that's how I see it.