Chapter 2

TOWARD A SCIENCE OF DIGITAL FORENSIC EVIDENCE EXAMINATION

Fred Cohen

Abstract   Digital forensic evidence examination is not a normal science at this time. This paper discusses the important issue of moving toward a science of digital forensic evidence examination. It highlights key areas in which progress has to be made in order for digital forensic evidence examination to become a normal science.

1.   Introduction

Like almost every scientific endeavor, the examination of digital forensic evidence (DFE) started out somewhere between an art and a craft. People with special skills and knowledge leveraged their skill sets and knowledge to put forth notions about the meaning of DFE in the context of legal matters. While the court system greatly appreciates science and its role through expert testimony in providing probative information, the appreciation is substantially challenged by the lack of a scientific base. A scientific base includes a well-defined and well-understood body of knowledge, an underlying scientific methodology, an experimental basis, adequate peer-reviewed publications associated with professional societies, and all of the other things that go with normal science. As the volume and criticality of DFE has increased, there is an increasing recognition of the limitations of DFE, and more importantly, the limitations of the underlying science.

To clarify the notion of a science of DFE examination, it is instructive to examine the advancement of science in other disciplines. In most disciplines, a scientific methodology consists of four basic elements: (i) studying previous and current theories, methods and their experimen-

tal basis; (ii) identifying inconsistencies between current theories and repeatable experimental outcomes; (iii) hypothesizing new theories and performing experiments to test the new theories; and (iv) publishing the results. However, in an area where there is no pre-existing scientific infrastructure, a new theory, methodology, experimental basis, and perhaps even a new physics, have to be built from scratch. In the case of DFE examination, only one such attempt has been made so far [3], and this paper is substantially about that attempt.

2.   The Call for Science in Forensics

The U.S. Supreme Court has spoken [8] and the National Research Council has concurred [6]. A rigorous scientific approach is needed for forensic evidence to warrant its use in courts in the United States. Much of the world is likely to follow this approach.

To a substantial extent, this call for science stems from some dramatic failures of forensics. For example, in the Madrid bombing case, the FBI declared that a fingerprint from the scene demonstrated the presence of an Oregon attorney. However, the attorney, after having been arrested, was clearly demonstrated to have been in another part of the world during the time in question. The side effect is that fingerprints are now being challenged as valid scientific evidence around the world [5].

A similar situation exists in cases where forensic examiners have done poor quality work and have testified in a number of cases, typically for the prosecution. The inability to effectively challenge evidence by these supposed experts using a scientific methodology and inquiry process is extremely disconcerting, all the more so because of the limits of human integrity. In case after case, when the details are examined, forensic evidence seems to come up short under competent challenges and close scrutiny. The solution is simple. Build and apply real science, and the truth will out.

3.   Proposing a Science

This first attempt at proposing a science for DFE examination involves the creation and enumeration of elements of an epistemology and a physics of digital information, a model of the DFE examination process in the legal environment, and the interpretation of existing information, theory, and experimental results in the new model. This is the first attempt to create a scientific model for DFE examination, and a suitable name would be "the standard model." But of course, it will only become a standard to the extent that it is embraced by the community. Also, it would have to be adapted over with time as the digital forensic

community comes to decide and adapt to the realities of the scientific method.

3.1   Epistemology for Digital Forensics

Epistemology studies the nature of knowledge, its presuppositions, foundations, extent, and validity. In the case of DFE examination, some basics may be reasonably assumed for the purposes of creating a science. The following are some of the epistemological issues that have already been identified.

Digital evidence consists entirely of sequences of binary values called bits. The physics of DFE is different from that of matter and energy, and thus the normal assumptions made with respect to how the world works do not apply – or do not apply in the same way – to DFE.

Two major differences are that DFE has observation without alteration and duplication without removal. Also, computational complexity limits what can be done with resources in a given time frame – one could say that the speed of light is "different" for DFE.

Physical evidence is very often transfer evidence and is sometimes trace evidence. In contrast, DFE is always trace evidence, but essentially never transfer evidence. Also, DFE is normally latent in nature in that it can only be observed through the use of tools. This implies a multitude of requirements surrounding DFE tools and their use.

In a "scientific" approach, the theories are not casual theories, but "scientific theories." That means that:

- Constructs are testable.

- Refutation can destroy a theory. Finite confirmations do not prove a theory; they can only confirm it.

- Scientific theories change slowly. Once theories are accepted, they only change – in very rare cases – because of dramatic changes in the understanding of the underlying physics.

The "theories" of DFE lead to a physics of digital information. Many of them are based on existing and widely-accepted mathematical knowledge. However, some are still conjectures from computer engineering, computer science, discrete mathematics, and related areas.

3.2   Information Physics

The physics of digital information is significantly different from that of the physical world. The differences are described in more detail elsewhere [3]. However, to get a sense of the differences, many of the under-

lying assumptions of the physical world, such as smoothness, continuous space, the notion of transfer, continuous time, and even the speed of light, are very different in the digital world. Indeed, the assumptions simply do not hold and the implications are, in some sense, profound.

Input sequences to digital systems produce outputs and state changes as a function of the previous state. To the extent that the state and outputs produce stored and/or captured bit sequences, these create traces of the event sequences that caused them. Thus, a trace may be defined as a set of bit sequences produced by the execution of a finite state machine.

We generally think of the physical world as a space that diverges with time, with any given initial conditions in history producing a wide variety of possible future outcomes. As a result, when looking at a physical trace, at least theoretically, it is possible to identify a unique historical event sequence that produced such a trace. But the digital space converges with time, so instead of the one-to-many relations seen in the physical world, there are many-to-one relations in the digital world. This means that a very large number of potentially different input sequences and initial states may produce identical traces (i.e., from subsequent states and sequences). Almost any digital trace could be the result of a large number of different historical event sequences, and the number of these sequences increases dramatically with the passage of time (i.e., execution of finite state machines). Thus, the traces from digital devices are not, in general, unique as to the input sequences that produced them.

Another less mathematical problem is the relationship between the unlimited granularity of the physical world in time and space and the finite granularity of the digital world in time and space. Because of this difference, a discontinuity exists at the interface between the physical and digital world. Minor differences are exaggerated near the discontinuity while major differences are ignored away from the discontinuity. The limited sensor and actuator capacity of devices that convert between the digital and physical world prevents the exchange of a variety of information that is potentially probative, and makes it much easier to create a variety of forgeries at the interface. This implies that input sequences may not directly demonstrate which non-digital events sequences may have produced them. As a result, additional effort is required to attribute traces to real-world causes and forgery is much easier in the digital space than in the physical space.

Viewing DFE as the result of processing by finite state machines inherently limits the potential utility of DFE examination for providing probative information about real-world events. DFE examiners must

take the limitations into account when performing their examinations and when testifying about the results of the examinations. These limitations are directly due to the limits of DFE and the methodologies used to understand and work with it.

3.3   DFE Examination Model

The model of DFE examination is related to an overarching model of digital forensics [4]. It can be codified in mathematical terms as follows:

- **Laws:** $L : \{l_1, \ldots, l_n\}, \; R : \{r_1, \ldots, r_m\}, \; L \times R \rightarrow [F|T]$

- **Violations:** $L \times R \Rightarrow V$

- **Claims:** $E : [\{E_1, \ldots, E_o\}]$

- **Events:** $\forall e, e \in E^*$ that demonstrate claims
  $[\forall Ex \in E, \; Ex : (ex_1 \in E^*, \ldots, ex_p \in E^*)]$

- **Traces:** $T : (t_1, \ldots, t_q)$

- **Internal Consistency:** $C : T \times T \rightarrow [-1, 1]$

- **Demonstration Consistency:** $D : T \times E^* \rightarrow [-1, 1]$

- **Forensic Procedures:** $P : \{p_1, \ldots, p_n\}$,
  $\forall p \in P, \; p \rightarrow c \subset C, \; p \rightarrow d \subset D, \; p \rightarrow c \not\subset C, \; p \rightarrow d \not\subset D$

- **Resources:** $R : (T, \$, C, E)$

- **Schedule Sequence:** $S : (s1, s2, \ldots), \forall s \in S, \; s : (l \subset L, r \subset R,$
  $h \subset H, e \subset E, t \subset T, c \subset C, d \subset D, p \subset P, r \subset R, t, t')$

In essence, the legal claims constitute a set of runs through the elements of laws that produce violations. This can be conceptualized as a partially-ordered set (poset). The events and traces are entities that are evaluated to determine the outcome of the legal matter, and they form the basis for the claims that are demonstrated using the violation poset. If the events and traces are consistent with an unbroken path through the poset, a violation is indicated; if not, inadequate indications for a violation are present. If T and E are inconsistent with the poset, then they may act to sever all of the paths forming violations, in which case adequate basis may be present to definitively demonstrate that no such violation is justified. To the extent that T and E are internally or demonstrably inconsistent, C and D may be used to show that the evidence or the claims are less probative, or potentially even prevent the admission of elements of T and/or E into the matter.

The fundamental theorem of DFE examination in this model may be stated in relatively simple terms:

> *What is not consistent is not true.*

DFE examination then consists largely of testing hypotheses related to the poset that form V as demonstrated by T and E in order to attempt to refute them by showing that they produce inconsistencies. This also implies some things about language and usage.

Consistency and inconsistency are demonstrated by logic and the theories associated with the physics of digital information. So, for example, given that a claim is based on an event e_1 causing a trace t_1, events and/or traces showing that t_1 happened before e_1 would be inconsistent with the claim of causality because information physics demands that cause precedes effect.

There are several consequences of this model related to: (i) the sizes of the model components; (ii) available computing power and its impact on thoroughness; (iii) limitations due to resources and schedules; (iv) limitations of available procedures; (v) legal limitations on what can be used and how; and (vi) probative versus prejudicial value and its relationship to consistency and related matters. In the example above, refutation is based on traces and events that may themselves be problematic. Thus, C and D are defined over the range [-1, 1].

In many cases, because of the limitations of DFE examinations as described here and elsewhere, more certainty is desired. Two general classes of methods exist to provide higher surety of DFE examination results: (i) identifying additional traces or procedures to gain additional demonstrations of consistency or inconsistency; and (ii) identifying redundant paths to prove hypotheses so that even if some paths are less certain or are eliminated, the overall hypotheses remain intact. These issues are covered by the model presented here.

3.4   Use of Defined Terms

No matter how many tests are performed, except for special cases, DFE examination results cannot prove a broad claim to be true [7]. The best that can be done is to show that the tests undertaken fail to refute the hypotheses and to show the extent to which the tests were thorough. This leads to the notion of what can reasonably be asserted as the most authoritative claim in [opposition] support of a hypothesis regarding DFE:

> *"The results of {the tests I did} were [in]consistent with {the hypotheses}."*

To the extent that some of these statements are combined by logical reasoning, an overarching statement may be made with regard to the claims. This could be of the form:

> *"Based on {the basis}, the {traces and events} are [in]consistent with {identify the claim(s)}."*

Or in some cases, when this is true:

> *"In my examinations of {traces and events}, everything I found was consistent with {claims} and nothing I found was inconsistent with {claims}."*

On the other hand, a single refutation disproves a hypothesis. The least that can be reasonably said if such a refutation is identified is:

> *"The {procedures I performed} demonstrate that {traces and events} are [inconsistent with/refute] {the hypothesis}."*

Thus, the methodology underlying the science of DFE involves:

- Devising testable hypotheses (E)

- Testing the hypotheses against the evidence (T and E) using forensic procedures (P) and logic to determine Type C and D consistency by attempting to refute the hypotheses.

- Making careful and limited statements about the results of these tests, typically using wording such as that identified above.

The following are some wordings that may apply in other circumstances. Some of the more commonly misused ones are also identified, along with definitions appropriate for use by DFE examiners.

- **Suggest:** Imply as a possibility ("The evidence suggests ...") – calls to mind – propose a hypothesis or possible explanation.

- **Indicate:** A summary of a statement or statements or other content codified ("His statement indicates that ...") – a defined set of "indicators" are present and have, through some predefined methodology, and identified as such ("The presence of [...] smoke indicates ...").

- **Demonstrate:** Exemplify – show – establish the validity of – provide evidence for ("The reconstruction demonstrates that ...").

- **Correlate:** A statistical relation between two or more variables such that systematic changes in the value of one variable are accompanied by systematic changes in the other as shown by statistical studies ("Based on this statistical analysis, the use of the KKJ account is correlated (p=95%) with ...").

- **Match:** An exact duplicate ("These two documents have matching publication dates, page counts, ...").

- **Similar:** A correspondence or resemblance as defined by specified and measured quantities or qualities ("The 28 files were similar in that they all had syntax consistent with HTML, sizes under 1,000 bytes, ...").

- **Relate:** A defined and specified link ("The file system is related to FAT32 in that FAT32 was derived from ...").

- **Associate:** Make a logical or causal connection with basis provided ("I associate these bit sequences with program crashes because ...").

Through the careful use and consistent application of these terms, the field of DFE examination may move forward more quickly, and peer reviews could create a body of work that is meaningful across endeavors and time. However, if, the DFE community is inconsistent or if its peer review process fails to force compliance with the terminology, then DFE examination is unlikely to proceed as a normal science.

3.5   Tools of the Trade

As an area of science, DFE examination has a relatively small number of peer-reviewed and repeatable scientific experiments. Most experiments are of limited applicability and are not focused on building a fundamental understanding. They do not meet the standards of scientific rigor expected in other fields. They are oriented toward confirmation rather than refutation, which makes them dubious as science.

Furthermore, there is a methodological challenge associated with experiments for several reasons. DFE is latent and, therefore, experiments require tools. Of course, this means that the experiments are limited by the tools and, like any other area of science, the examiner must understand the limits of the tools in order to understand the limits of the experiments. This, in turn, leads to the need to have a methodology to evaluate tools. Without such a methodology, regardless of what the tools may indicate, the interpretation of the results is open to question.

A methodology for understanding tools might start with the development of an error model. Classical error models for digital systems [1, 2] may well be applicable, but their utility will not be known until they are applied in DFE examinations.

It is also important to understand how to calibrate and test tools, and to create systematic approaches for doing so. Calibration processes

typically involve validation with known samples, which is readily done in most cases. The testing process typically involves verification of some sort, which, in the case of software, normally involves mathematical proofs or tests that verify the results against error models. Again, this is an area where DFE examination is lacking. Redundancy via the independent verification of results may provide an alternative in cases where no well-defined testing methodologies and practices are available.

Regardless of how "good" a tool is, it must be properly used, the results must be meaningfully interpreted, and the limits of the tool must be understood. This implies that the examiner must have knowledge, skills, education, training, and experience suited to the use of the tools they apply. DFE examination has few advanced students and teachers and, as a result, produces "niche experts," who are of limited utility. It has many niche experts who can potentially speak to very narrow domains and it has expert claimants who profess expertise beyond their actual knowledge, skills, education, training, and experience. Indeed, DFE examination as a field has too few "real experts" at this time.

3.6   Presentation

Another major issue with tools today is how they present results, both in support of the examination process and when the results of examinations are presented in reports or before judges and juries.

Presentation is intimately tied to, but not directly part of, examination. Because DFE is latent, presentation will always be an issue. For the examiner, the results of the experiments must be presented using tools. For the jury, presentation is again fundamental to understanding the evidence and the examination results. For the judge, the same is true to evaluate admissibility. For the opposition, presentation is just as critical to evaluating expert reports,. Today, however, there is no standard for even presenting the most common representations of DFE. Even something as simple as presenting a text file is fraught with potential errors.

Different ways of presenting the same information can lead to different interpretations and outcomes. Consider this simple example:

- Plaintiff's sworn statements are inconsistent with the evidence.

- If Plaintiff's sworn statements are to be believed, then the evidence refutes Plaintiff's claims.

- If the evidence is to be taken at face value, then it refutes Plaintiff's sworn statements.

The first of these statements encompasses the other two. The second statement appears to say that one can assume that the plaintiff is telling the truth, but the evidence does not support the claim. The third statement appears to say that the plaintiff is lying.

Technical presentation errors are also problematic. For example, the digit "O" and the letter "0" are almost indistinguishable, as are the digit "l" and the letter "1." Spaces at the ends of lines, and the differences between a leading tab, a leading space followed by a tab, and leading spaces cannot be discerned in normal printouts. As an aside, the nature of the problem becomes very clear if the reader failed to notice that the numbers "0" and "1" and the letters "O" and "l" are interchanged above.

When examining the output from widely-used and trusted tools, the presentation produced by the tools often fails to aid the examiner in seeing these sorts of differences. In case after case and in tool after tool, differences that might allow the examiner to detect inconsistencies go unseen and, thus, the inconsistencies are commonly missed. Even something as simple as a forensic font would largely alleviate these problems, but this notion was only first introduced in late 2009.

Clearly, presentation is fundamental to the advancement of a science of DFE examination. Also, presentation is critical to the effective use of tools upon which essentially all of digital forensics depends.

4. State of the Science

The principal elements of DFE examination are analysis, interpretation, attribution and reconstruction. While this categorization of the discipline may be somewhat arbitrary, it is sensible in that each of these elements, while not entirely distinct from one another, encompasses different techniques that may require different expertise and tools.

4.1 Analysis

Analysis engages techniques that provide definitive answers to specific questions. Existing analytical techniques are limited in number and in the nature of the definitive results they produce. The techniques are computational in nature with defined complexity and defined input and output limitations, and their accuracy can be verified (at least theoretically). Methods that do not meet these criteria should not be called analysis.

Analysis starts with a bag of bits. Redundancy in the bag of bits confirms or refutes hypotheses about what the bag of bits is. Through analysis, features and characteristics are detected, symbol sets are iden-

tified, trace typing is undertaken, content is parsed, normalized and elucidated, and indicators are analyzed. Any of these may return the examiner to a bag of bits, especially if it produces inconsistencies that destroy the chain of analytical results supporting the current hypotheses about the bag of bits.

During analysis, characteristics and features are analyzed for consistency, traces are ordered and out-of-order entries are detected, sourcing and travel patterns are identified, consistency is checked across related records, anchor events are used for external validation, time differentials and jitter are considered, and all of these are compared with hypotheses in order to identify consistency or inconsistency. Also, sieves are constructed and items are counted, derived traces are formed for analytical convenience, counts are made of various features and characteristics of interest to the examiner, mechanisms are combined and the resulting errors identified and mitigated, and results are verified by independent means where feasible. Additionally, intentionally-hidden items of interest are sought, content placed in hard-to-find locations is found, steganographic and other transformed content are identified and inverted, recursive embedded languages are parsed, and indicators are identified and sought relative to the issues at hand.

In these analytical processes, automated and human cognitive methods are combined to identify potential Type C and D consistencies and inconsistencies either automatically or through DFE examiner interaction. All of these processes are limited by the available forensic procedures (P) and their computational complexity, which impacts and is impacted by resources.

4.2   Interpretation

Traces, events, claims, and analytical results must be interpreted in the context of the legal matter in order to be meaningfully examined. Interpretation involves selecting among alternative explanations (hypotheses). Different interpretation methods are used for different circumstances, with one major difference being the treatment of structured and unstructured traces.

Over-interpretation by examiners is common. For example, terminology is misused, conclusions are drawn in excess of what is supported by the data, the terms "match" and "correlate" are overused, and the details of the basis are often left unspecified.

Special care must be taken in making statistical claims. In addition to simple interpretation errors (e.g., percentages that add up to more than 100%), claims often ignore data that is not present, conceal assumptions

related to random stochastic processes and other similar things, and assert precision exceeding accuracy. This is particularly problematic in digital systems, where 33.333% results from an experiment with only 9 samples. At best, there is a little less than one digit of accuracy, so this precision is misleading. The examiner must interpret the output of analyses and present information in clear terms in order to properly interpret results for presentation – 3 of 9, 1/3 or 33% are all valid, but 33.333% is not.

Tools commonly interpret traces as well, and to the extent that they do so, these interpretations make assumptions, often without basis. The interpretations often present false results based on the assumptions and result in presentations and other depictions that may mislead the examiner as well as the judge and jury. No methodology currently exists for evaluating interpretation by tools, although a wide range of cognitive errors are known to cause incorrect or imprecise interpretation.

Interpretation includes the identification and explanation of missing traces and their implications to the matter at hand. While some experts may know the traces that are commonly present in some situations, there is neither a widely available library of what constitutes "normal" behavior in an operating environment nor a basis for comparison in the existing literature.

While redundancy may be used to mitigate interpretation errors, interpretation is highly subject to individual variations. Proper interpretation is limited and couched in terms of its accuracy and applicability, including the things that the examiner does not know. It is important to enforce careful word usage for technical terms in professional publications and in legal venues. How can one claim to be an expert when one does not even use the published and accepted terms of the scientific and technical community?

There is also a tendency to create casual theories and embrace them as if they were something more. To be clear, the theories of DFE examination are the results from information physics, the assumptions of the model, and the foundational logic and results of computer science and engineering. Any other theory, unless and until it passes peer review and is reconciled with the existing theory, is only a hypothesis at best.

Examiners, like all people, make cognitive errors during the interpretation process. The examiner would be well served to examine the literature in this area and apply it to maintain clarity regarding his/her own as well as others' potential for making mistakes. Indeed, it is unwise to believe what one thinks is right unless one really knows. What some reviewers call "put a stake in the ground" should be shunned in the DFE community in favor of properly identifying the limitations and

not overstating a case. Peer review of examination results, particularly interpretations, should be sought in every case and for every report.

As a working assumption, the DFE examiner should assume that each system and situation is in fact different from others. Experience is a great teacher, but it can be misleading. Care must be exercised in not over-interpreting the data. And when the interpretation is unclear, reconstruction is a viable and worthwhile approach for gaining additional data. All interpretation should also be made within the context of information physics, and with the working assumption that every interpretation will be tested against information physics and refuted if it is inconsistent.

Events are also subject to interpretation because words are interpreted by examiners. Cognitive errors in interpretation are common and, since examiners assume the context of their knowledge and words mean different things to different people, the careful interpretation and close scrutiny undertaken by an examiner may be viewed as being "picky" by others; but this is the nature of interpretation in the legal environment. Furthermore, events are often difficult to interpret and leave many possible interpretations, particularly in light of the lack of specific technical knowledge by the individuals who offer the information related to the events.

Events should presumably be viewed through the lens of information physics. Claims are often inconsistent with and are, thus, refuted by the careful use of physics. As an example, causality is tricky because of the requirement for time ordering and the fact that complexity limits actions in the digital space; this results in a poset of causal relationships. Reviewing information physics for events may be helpful, but being thorough in this regard is problematic because of resource limitations and the lack of apparent progress when the examiner is thinking about the issues in light of other considerations. At this time, it is not possible to automate much of this examination and most DFE examiners lack the time or knowledge to do a thorough (or even thoughtful review) against the full set of theories.

Resources limit interpretation and tight schedules may prohibit contemplation of the issues in many cases. Computational resources and costs may be prohibitive for certain procedures, especially when a large volume of traces are involved. Available traces may change with time and legal actions as well.

Statements made and documents created by examiners are often unnecessarily interpretive. Careful wording is highly desirable in written and verbal communications. There are few "standard wordings" for DFE examination, and examiners often make statements that end up

being wrong, but that could have been accurate had they been stated differently. For example:

"*I found X B's in File Q*"

is true even if there are more than X B's present, because, presumably, the examiner did not find the other ones. On the other hand, the interpretation:

"*There were X B's in File Q*"

is potentially problematic given the same circumstance. It is a good idea to minimize interpretation and favor statements of fact wherever possible.

Similarity is almost always interpretive, in large part because DFE examination lacks adequate similarity metrics or criteria, an experimental basis for many similarity claims, and techniques for detecting similarity in many cases. Still, examiners have the tendency to use the term "similar" when it cannot be backed up with facts. Similar problems are often indicated and commonly associated with the suggestion of matches relating associated correlations. Clearly these uses are unclear and imprecise, and the lack of clarity and precision during use, while perhaps suggestive to the untrained reader, should be readily identified as indicative of a lack of knowledge, education, and training in DFE examination.

Assumptions underlying interpretations are critical, but they are often not detailed in statements and reports. Thus, the statements and reports lack the basis for the interpretations. It may be hard to identify all the assumptions, but confirming/refuting assumptions and hypotheses is the process that should be relied upon in interpretation; therefore, identifying and presenting them is important to obtaining the right answer.

Presentation is also an interpretive function that drives every aspect of DFE examination. It acts within the process of examination to skew the approaches and procedures used by the examiner. The presentation provided by a tool is the only thing the examiner has to rely on to understand the latent traces. Examiners must understand that they are interpreting the output of tools and that the outputs and interpretations are fraught with the potential (often realized) for producing cognitive faults, leading to failures that may significantly impact examination results and process.

4.3   Attribution

Correlation is not causality; before is not because. On the other hand, a lack of correlation certainly throws considerable doubt on the notion of causality, and the cause had better be before the effect.

Attribution of actions to actors is centered on the notion of causality, to the point where they are, in essence, inseparable. The fundamental assumption of causation in the digital world is:

> *Traces come about by the execution of finite state automata that follow the physics of the digital world.*

The physics of the digital world is useful in assessing attribution claims. For example, A caused B implies that A came before B. How long before is determined by the digital version of the speed of light. Computational complexity, the performance levels of devices, the physical speeds associated with the devices and their components, and the mechanisms available for storage, transport and processing, all place limits on the interval of time between the cause and the effect.

Statistics do not apply in most cases relating to attribution because the past is what it is; there is no "might have been" or likelihood that some particular thing happened. Either it happened or it did not happen.

Establishing a causal chain is non-trivial as well. It often involves redundant records and almost never eliminates all other possibilities. But it may provide a seemingly overwhelming body of information that is consistent with the hypothesis of attribution, and no information available may lead to inconsistencies with a hypothesis. This becomes a compelling argument for a judge or jury, even if the examiner never claims it to be definitive.

Finite state machines are highly predictable because driving state and input to output leads to the same answer every time in almost all cases. However, finite state machines converge with time (while the real world diverges), so where simulation may produce identical outputs, reversing time does not give a unique answer. Therefore, in the digital world, convergence implies that many paths lead to the same traces. In addition, because sensors that change physical inputs to digital outputs are highly nonlinear, small differences are expanded near a nonlinearity while large differences are reduced far from a nonlinearity. Thus, the interface tends to break the digital perfection of even forward-driven causal chains.

Attributing actions to human actors is even more problematic. Authentication methods are of limited value. Most biometrics are not good for identification, but rather only for selecting a known individual out of a group of a few thousand known individuals when deception is not in use. Something the user has can be taken or exploited; something the user knows can be known by others; something the user can do can be done by others.

DFE can almost never put a person at a keyboard at a time, although other events may be able to help. And even if the user is at the keyboard, it does not necessarily mean that he/she was in control of the computer. Behavioral attribution may use words and word sequences, commands that are executed and usage patterns, keyboard use and timing patterns, etc., but all of these fail under deception. Also, they are limited to picking known individuals out of small known groups and they have significant error rates even under non-stringent test conditions.

Device authentication and attribution may be accomplished using various indicators (e.g., operating environment identification data, device identifiers, and known behaviors of programs and mechanisms). While most of these are readily subverted by deception, some have properties that make them difficult to forge. Redundant traces may be applied to reduce the impact of deception.

Attribution of damage to parties is also important, e.g., to establish threshold requirements for criminal charges. Actual damages are typically divided into: (i) physical damage; (ii) conversion to use by another party; (iii) deprivation of utility to the owner, which is often the key DFE issue; and (iv) lost value or lost rights (e.g., disclosure of trade secrets or release of pre-patent data). If attribution can be done of the cause to effect, computation methods are available to identify the extent of the deprivation. This may include things such as the cost per usage of electricity, cost in reduced lifetime of equipment, cost in demonstrable lost business, and cost in reduced life of equipment.

4.4   Reconstruction

Driving time backwards is one approach to reconstruction, but information physics shows that this is problematic. The lack of adequate traces over time leads to very large envelopes of possible histories, which make it almost impossible to tell what was "original." In addition, theft may not be identifiable, travel time and jitter produce ordering uncertainties, and reversing time in a unique manner through homing sequences is impossible. Error accumulation also leads to large expansions when reversing time. The list goes on and on, which leads to the alternative of experimental demonstration of operation in the forward direction.

Experimental demonstration of operation in the forward direction is a form of reconstruction that can be used to test hypotheses. As such, it can confirm, refute, or be unrevealing. The basic methodology deals with constructed traces (C-traces) and original traces (O-traces). Similarity measures are used to define, in advance, the criteria for identifying sets of

C-traces mapped to O-traces, and the implications of outcome class sets. The examiner creates reconstruction(s) based on hypotheses, generates C-traces, and compares the C-traces to the O-traces to confirm or refute hypotheses.

5. Toward a Normal Science

Based on the discussion above, the following statements can be made today based on the science of DFE examination:

"I did X and observed Y."

"I [did not find/found] X in Y."

"X is [in]consistent with the claim Y because..."

"X [suggests/indicates/demonstrates/correlates with/matches/is similar to/relates to/associates with] Y because..."

If examinations are properly undertaken, each of these can have a sound basis with the proper scientific underpinnings. Unfortunately, the current set of methodologies, processes and procedures are limited in terms of their validity, testability, reliability, calibration, and basis. Additionally, there is a lack of strong agreement within the DFE community about many aspects of the science as a whole. While most of the results are peer reviewed and accepted within individual communities, several problems exist:

- The overall collection of results (as a body of science) is not recognized as such.

- The unifying methodology expressed with regard to the application of information physics to determine consistency is gaining acceptance very slowly.

- Models that are currently in use are used for various limited purposes, and are not widely adopted.

- Procedures and their results are limited and are not formalized or standardized.

- Tools and processes are only explored to a limited extent, with notions of completeness and thoroughness just beginning to be defined.

- Error models have not been adequately applied from other mature fields. The sources and magnitudes of uncertainty are poorly defined, and confidence intervals simply do not exist.

In most cases, the honest and knowledgeable examiner is largely limited to the most basic:

"I did X and observed Y,"

with the observation being typically limited to:

"I [did not find/found] X in Y."

While these are powerful statements that are appropriately used in place of other less sound statements, they are a long way from the level of science that DFE examination has the potential to achieve.

6.   Conclusions

DFE examination is not operating as "normal science" today. While there is a scientific basis for many activities involved in DFE examinations, a widespread consensus and common methodologies are still lacking. The foundation for scientific theories exists, but little attention is paid to testing the theories and developing the science. To successfully make the transition to a normal science, the DFE community will have to ask and answer several questions:

- What well-defined and consistent terms should be used?

- What well-understood epistemology should be used?

- What theories and methodologies should be chosen?

- What strong experimental foundations should be built?

- What agreed-upon physics should be used and how should it be formulated?

- How could community consensus be built?

- Should the path outlined here be embraced?

- If not, what is the best path?

The view of this paper is that there exists at least one description of a reasonably comprehensive scientific foundation underlying DFE examination. Regardless of the problems and limits of the foundation, it is a place to start building a normal science and advancing the field. As the field matures, normal science is almost inevitable, but the normalization process is only just beginning. A community consensus is highly desired, and this paper supports and anticipates such consensus in the near future.

References

[1] A. Avizienis, J. Laprie, B. Randell and C. Landwehr, Basic concepts and taxonomy of dependable and secure computing, *IEEE Transactions on Dependable and Secure Computing*, vol. 1(1), pp. 11–33, 2004.

[2] M. Breuer and A. Friedman, *Diagnosis and Reliable Design of Digital Systems*, Computer Science Press, Rockville, Maryland, 1981.

[3] F. Cohen, *Digital Forensic Evidence Examination*, ASP Press, Livermore, California, 2009.

[4] F. Cohen, Two models of digital forensics examination, *Proceedings of the Fourth International Workshop on Systematic Approaches to Digital Forensic Engineering*, pp. 42–53, 2009.

[5] G. Fine, Statement of Glenn A. Fine, Inspector General, U.S. Department of Justice before the House Committee on the Judiciary, Subcommittee on Crime, Terrorism and Homeland Security concerning Section 1001 of the USA Patriot Act, U.S. Department of Justice, Washington, DC (www.justice.gov/oig/testimony/0505b.htm), May 10, 2005.

[6] National Research Council of the National Academies, *Strengthening Forensic Science in the United States: A Path Forward*, National Academies Press, Washington, DC, 2009.

[7] K. Popper, *The Logic of Scientific Discovery*, Hutchins, London, United Kingdom, 1959.

[8] U.S. Supreme Court, Daubert v. Merrell Dow Pharmaceuticals, Inc., *United States Reports*, vol. 509, pp. 579–601, 1983.