For Good Measure
Paradigm

DAN GEER

Dan Geer is the CISO for In-Q-Tel and a security researcher with a quantitative bent. He has a long history with the USENIX Association, including officer positions, program committees, etc. dan@geer.org

Whether cybersecurity is, or could become, a science is no easy question, and any seemingly easy answer is simplistic. We need a science of security, that much is sure. Yes, we are making, and have made, significant advances in technique, but there is no doubt that something more generative needs to come onto the scene. Consider, via Figure 1, six major advances in technique and their useful lifetime [1].

But what would a science of security be? There is much to think about in revisiting T. S. Kuhn's 1962 landmark work, *The Structure of Scientific Revolutions*. Kuhn begins and ends with what is a circular idea, that a scientific community is defined by what beliefs practitioners share, and what beliefs practitioners share defines what community they are in. This is, in fact, instructive as no science begins in mature form, but rather any new science will begin in much more modest circumstances where, early on, consensus is not even a concept. As such, part of becoming a mature science is the development of a broad consensus about the core concerns of that branch of knowledge. Kuhn's word for the collections of exemplars of good science was "paradigm," a word whose meaning today is all but entirely Kuhn's, even among those who've never read a word he wrote.

But what is a "paradigm" and why do we want one? As Kuhn puts it, "[Paradigms] are the source of the methods, the problem-field, and standards of solution accepted by any mature scientific community at any given time." Kuhn's book and the two-decade-long back and forth between Kuhn and philosophers notwithstanding, the simplest version is that a paradigm is all the things that a scientist can assume that his or her colleagues will congenially understand about their common work without explicitly explaining them or arguing them from first principles again and again.

Again quoting Kuhn, "Men whose research is based on shared paradigms are committed to the same rules and standards for scientific practice. That commitment and the apparent consensus it produces are prerequisites for normal science, i.e., for the genesis and continuation of a particular research tradition.... Acquisition of a paradigm and of the more esoteric type of research it permits is a sign of maturity in the development of any given scientific field." Competing schools of thought are always present before a mature science first appears. As Kuhn said, "What is surprising, and perhaps also unique in its degree to the fields we call science, is that such initial divergences should ever largely disappear. For they do disappear to a very considerable extent and then apparently once and for all." Perhaps it is thus possible to say that topics of study that never coalesce their competing schools are either fated never to be sciences or are in some state of arrested development that may someday be cured. Those that can and will, but have not yet done so, are what Kuhn called pre-paradigmatic, meaning not yet a science. The appearance of a paradigm that all can accept is when the transition to being a science occurs, or, as Kuhn put it, "Except with the advantage of hindsight, it is hard to find another criterion that so clearly proclaims a field a science."

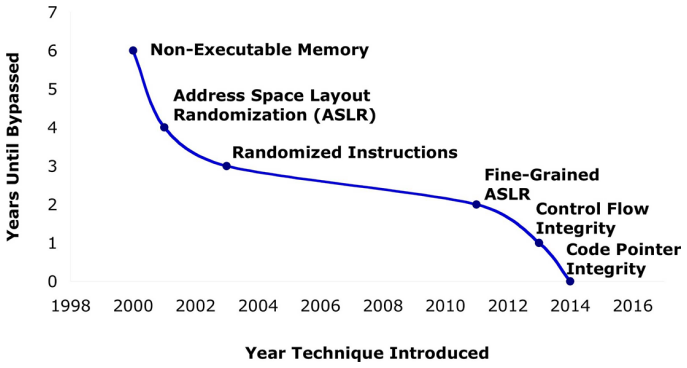

Figure 1: Major advances in technique and their useful lifetime

For Kuhn, the appearance of a paradigm transforms those who merely study first into a discipline and then into a profession. One can even say, and Kuhn does, that the paradigm itself is the last result of the science in question that can be appreciated by the lay audience—after that, all progress is in journal articles not readable by non-specialists, to the extent that "[t]he scientist who writes [for the lay reader] is more likely to find his professional reputation impaired than advanced" for having done so.

Now as everyone here knows, from time to time a science may undergo a revolution, which in Kuhn's terms is precisely the laying down of one paradigm in preference for another. The title of his book is to be understood as precisely that, that scientific revolutions share aspects of structure that we can now describe as there have been enough of them in the last 400 years to discern that structure. If you consider physics to be the paragon of a hard science, then the transition from Newtonian mechanics to Einsteinian relativity demonstrates exactly the point Kuhn was making, that there comes a moment when research has reached a kind of impasse where the nature of what now look to be puzzles needing further study cannot be profitably investigated within the paradigm that now holds.

Kuhn referred to these impasses as the appearance of an anomaly, one that the existing paradigm cannot evaluate by way of further research consistent with the paradigm then in place. His review of past revolutions centered in each case on the appearance of irreconcilable anomalies that made a given field ripe for revolution. That roasting metals caused them to gain weight thus indicating that they had absorbed some fraction of the air around them, a fraction that could be exhausted, led to the idea that air might not be the one and only gas but rather a combination of gases. Perhaps more significantly to the very idea of revolution is that even though Lavoisier had discovered oxygen, others in the field, notably Priestly, never accepted the existence of oxygen and held to the phlogiston theory to the end of their careers. I say "more significantly" as the trite version of "What is a scientific revolution?" is that it is a time when newcomers to the field adopt the new paradigm while those already in the field slowly die off. It is a generational change.

Kuhn's idea of crisis is the dual of his idea of paradigm. Where a science's paradigm suggests puzzles that further research will solve, in a crisis this is no longer so. Yet the occasional crisis is itself necessary for advancement as any paradigm whose theories completely explain all observable fact ceases to be science and becomes engineering. In other words, a crisis is not the end of research but the substitution of a new paradigm for an old and a new set of research puzzles awaiting solution.

Just as "a scientific theory is declared invalid only if an alternate candidate is available to take its place," to reject one paradigm without simultaneously substituting another is "to reject science itself." In short, when an anomaly appears, there are only three resolutions available: (1) solve the problem, (2) leave the problem for future scientists, or (3) use the crisis to force a new paradigm on the field.

When there is a shift of paradigm, that is to say a scientific revolution, it may serve to redirect a field so completely that some parts of it fall away entirely, the separation of astronomy from astrology or the separation of chemistry from physics being two examples. As Kuhn put it, the choice between paradigms is a choice between incompatible modes of community life (and, yes, he does mean "scientific community" in an altogether social sense). "Since no two paradigms leave all the same problems unsolved, paradigm debates always involve the question: Which problems is it more significant to have solved?"

One of the first questions we might ask is whether cybersecurity is a science or, if not, whether it ever will be. I am one of several expert reviewers for the National Security Agency's annual "Science of Security" competition and award [2]. Quoting its rationale, "The competition was established to recognize the current security paper that best reflects the conduct of good science in the work described. [Science of Security] is a broad enterprise, involving both theoretical and empirical work. While there can only be one best paper, any one paper cannot span that full breadth. Nonetheless, the field is broad and work in all facets is encouraged and needed. The common denominator across the variety of approaches is solid methodology and effective communication, so those aspects of the papers [are] strong factors in our decision" [3].

Papers are nominated for consideration, and I encourage you to do so, but I am also here to report that among the reviewers our views of what constitutes a, or the, Science of Security vary rather a lot. Some of us would prioritize purpose, agreeing with Charles Darwin that "all observation must be for or against some view if it is to be of any service" [4]. Some of us view aspects of methodology as paramount, especially reproducibility and the

clarity of communication on which it depends. Some of us are ever on the lookout for what a physicist would call a unifying field theory. Some of us insist on the classic process of hypothesis generation followed by designed experiments. We vary, and I take that to be a vote of sorts on whether cybersecurity is yet a science.

Whether cybersecurity is yet a science is a hard question. Let's consider candidate paradigms of cybersecurity; if they exist and have turned over from time to time then, yes, cybersecurity is now a science. Take one of the most basic tools we employ, that of authentication. Authentication is the solution to the puzzle of identity establishment, a puzzle that derived from the paradigm of perimeter control.

The paradigm of perimeter control has been in an evident crisis for some time now. The crisis is not merely because the definition of perimeter may have been poorly applied in practice, but because some combination of always-on and universal addressability collectively make the paradigm of a defensible perimeter less and less a paradigm where research is itself likely to patch up the mess and retain the core and guiding paradigm of perimeter control. The parallel with Ptolemaic astronomy is pretty fair. On the one hand, every improvement in observational accuracy made the motions of the planets more complicated to describe with epicycles upon epicycles. On the other hand, our hand, the threat to systems from always-on universal addressability has become too rich to be just a new set of puzzles solely within the paradigm of perimeter control—the defensible perimeter began to have its own version of epicycles within epicycles by a shrinking of what a perimeter could, or should, control [5].

A second crisis for the paradigm of perimeter control is upon us now as exemplified with a commercial example. Table 1 counts cores in the Qualcomm Snapdragon 801.

Central CPU	4
Adreno 330 GPU	4
Video out	1
Hexagon QDSP	3
Modem	2–4
Bluetooth	1
USB controller	1
GPS	1
Wifi	1–2
Charging	?
Power	?
Display	?

Table 1: Identifiable cores in Qualcomm Snapdragon 801

That is somewhere between 18 and 21 cores. In the vocabulary of the Internet of Things, is that one "thing" or the better part of two dozen "things"? Is the perimeter to be defended the physical artifact in the user's pocket or is it the execution space of each of those cores? The compound annual growth rate of deployed "things" approximates 35%—meaning a 27-month doubling time. That's a lot of new perimeter.

If the paradigm of perimeter control is no longer producing puzzles that can be solved by further scientific research, then what? Noting, as Kuhn does, that "to reject one paradigm without simultaneously substituting another is to reject science itself," what might be a substitution? If everything we are or do is unique if examined closely enough, then the idea of authentication as verifying an assertion like "My name is Dan" can easily morph into an observable like "Sensors say that this is Dan." In other words, our paradigm of an authentication transaction before any other perimeter-piercing transaction is itself showing its age.

The paradigm that is the obvious alternative to perimeter control, and thus authentication as a gating function, is accountability based on one single unspoofable identity per person. If I am right—that real-soon-now identity is simply an observable that needs no assertions—then that single identity which the individual has but does not need to prove may be fast upon us. The National Strategy for Trusted Identities in Cyberspace is not worded in that way although that is how I read it. That means there is a crisis in privacy research, too. Privacy's paradigm has long been, "Privacy is the power to selectively reveal oneself to the world" [6], but all uses of virtual reality come with total surveillance.

Nevertheless, if being part of the modern world in no more robust way than appearing unmasked on a public street is the same as submitting to unitary identification observable at a distance by things you never heard of, then that means either submission or withdrawal for the individual.

Note that Kuhn never said that a switch to a new paradigm would be delightful or comforting, he merely said that it would better explain the way the world works while suggesting new puzzles for scientists who share that paradigm to pursue. Authentication transactions as a prodrome to authorization transactions in service to a paradigm of perimeter control may soon be behind us, including in the peer-to-peer world. If, in fact, being authenticated as yourself is unavoidable, then there is no proving that this is the Dan for whom there is a book entry allowing him into some robot-protected building, but rather an accountability regime based on whether that Dan did or did not enter a building for which he might later be penalized. His crime would not be masquerading as some identity other than his own so as to get in, but rather that of he was observed to have gone in even though he was forbidden to do so.

What I am suggesting as the crisis around the paradigm of selective revelation is that, as with metadata, there is so much redundancy in what is observable that prohibiting one or another form of collection has no meaningful effect whatsoever on those agencies, whether intelligence or advertising, who would build a model of you from metadata alone. As but one example, with current technology I can read the unique radio signature of your beating heart at five meters. As with anything that has an electromagnetic output, the only technological question is the quality of the antenna. If I can take your picture on the public street without your permission or notice, why can't I record your heart? Or your iris? Or your gait? Or the difference in temperature between your face and your hands? That list is long and getting longer. It is a crisis for which the paradigm of selective revelation can scarce put up puzzles fast enough, and scientific solving of those puzzles can, at best, trail the curve.

The crisis is simply that what heretofore we have known as confidentiality is becoming quaint and irrelevant. Perhaps science will have to reposition confidentiality within some new paradigm that prioritizes integrity, not confidentiality. Perhaps a world in which data can and will be collected irrespective of selective permission granting is a world in which the data had better be right. If more and more intelligent actors are to be out there doing our implicit bidding long after we've forgotten their configuration interface, then data integrity had better be as absolute as we can make it, and that is then where the research puzzles will have to be found.

Perhaps I have it wrong, perhaps the topmost paradigm of the science of security is simply that of defense. Perhaps the rise of sentient opponents makes a paradigm of defense unarguable, as evidenced by rafts of paradigmatically generated puzzles of the sort of how can this or that be hardened or otherwise defended, up to and including DARPA's Grand Challenge [7].

If defense is and has been our paradigm, then that, too, is in crisis. That is in no way a failure; paradigms only change due to the success that the one paradigm has in motivating science to explore the world thoroughly enough to discover anomalies that cannot be made to fit within the paradigm that caused them to be discovered in the first place. The outgrowth of the paradigm of defense has been guidance that has allowed us, including non-scientist practitioners, to get better and better. We have discovered and then deployed better tools, we have come to understand causal chains and thus have achieved better understood practices and work with better colleagues. That's the plus side, and it is one terrific plus side. But if I am interested in the ratio of skill to challenge, then, as far as I can estimate, we are expanding the society-wide attack surface faster than our science of security is expanding our collection of tools, practices, and colleagues. The paradigm of defense is in crisis.

One embodiment of the paradigm of defense has been the movement to build security in. The successes of that movement are precisely of the sort I mentioned before when I said that we have discovered and then deployed better tools. But to remind you of the truism in Adi Shamir's 2002 Turing Award lecture, "Cryptography is typically bypassed, not penetrated." I would argue that this is true of all aspects of the cybersecurity mechanism, including those delivered by building security in; it is the possibility of bypass that ultimately matters. Our sentient opponents know that, too, and their investments in automating the discovery of methods of bypass are in a hell of a horse race with both building security in and in-static analysis of code bodies, new or old. Look (again) at Figure 1.

Kuhn takes some pains to say why it is that a paradigm shift requires a crisis, that "to an extent unparalleled in other fields, [scientists] have undergone similar educations and professional initiations." One here must ask the central question of this essay by mirroring Kuhn: are the paradigms of cybersecurity in enough of a crisis that resolution of the crisis requires a change of paradigm? The answer is by no means obvious, although to my eye there are several crises now in play. If the crises sufficient to require a reformulation of the paradigm or paradigms of cybersecurity, then a scientific revolution is upon us, what Kuhn calls "a reconstruction of group commitments." As he points out, a crisis requiring such a reconstruction might not even be in cybersecurity itself, but instead might be due to discoveries in some other field or venue, just as discoveries in physics engendered a crisis in chemistry once upon a time.

That, then, is the question before us, complicated by the changing nature of what scientists of security are studying both with respect to rapid technological change and the presence of sentient opponents, leavened, of course, with the societal demands fast upon us largely independent of what we know or say. I think I see paradigms here that are in undeniable crisis. I can, of course, be entirely wrong, and we may still be working our way up to being a science, still coalescing schools of thought into the kind of paradigm that will define us as scientists.

A fuller treatment of this topic is available at http://geer.tinho.net/geer.nsf.6i15.txt.

References

[1] Chart data courtesy of Hamed Okhravi, MIT.

[2] NSA, Science of Security: www.nsa.gov/what-we-do/research/science-of-security/index.shtml.

[3] Best Scientific Cybersecurity Paper winner: https://www.nsa.gov/news-features/press-room/press-releases/2014/best-scientific-cybersecurity-paper-competition.shtml.

[4] Letter to Henry Fawcett, 1863, as quoted in E. J. Huth and T. J. Murray, eds., *Medicine in Quotations: Views of Health and Disease through the Ages* (American College of Physicians, 2006), p. 169.

[5] IT Conversations, "The Shrinking Security Perimeter," recorded March 1, 2004: audio mirror at: geer.tinho.net/Dan_Geer_-_The_Shrinking_Security_Perimeter.mp3.

[6] E. Hughes, "A Cypherpunk's Manifesto," March 9, 1993: http://www.activism.net/cypherpunk/manifesto.html.

[7] "Cyber Grand Challenge": https://cgc.darpa.mil/.