Cybersecurity, Nuclear Security, Alan Turing, and Illogical Logic

Martin E. Hellman
Professor Emeritus of Electrical Engineering
Stanford University
hellman@ee.stanford.edu

My work that is being recognized by the 2015 ACM A. M. Turing Award is in cybersecurity, while my primary interest for the last thirty-five years is concerned with reducing the risk that nuclear deterrence will fail and destroy civilization. This Turing Lecture draws connections between those seemingly disparate areas as well as Alan Turing's elegant proof that the computable real numbers, while denumerable, are not effectively denumerable. For reasons of space, this abstract treats only the first part of the lecture.

There is a striking similarity between my work applying risk analysis to a potential failure of nuclear deterrence [1] and whether public key cryptography will remain secure as now practiced. In both cases, high confidence is derived from what appears to be too little data.

Proponents of nuclear deterrence often point to the fact that over seventy years have elapsed without a Third World War. They see that as evidence that the strategy is safe, with noted IR expert Prof. John Mearsheimer stating that "nuclear weapons are weapons of peace." [2]

Risk analysis takes a more nuanced, statistical approach to the question of safety. It pays attention to the nuclear near misses that have occurred in those years as potential early warning signs. The approach goes further and looks at events with just the potential to initiate a nuclear crisis.

For example, in my 2008 preliminary risk analysis [3], I cited three events over the preceding fifty years that had the potential to precipitate a nuclear crisis over Cuba. I then estimated the arrival rate of potential initiating events at 6 percent per year. That kind of track record raises questions about the assumption that nuclear deterrence will work perfectly forever, which is what is what it must to for civilization to survive. That is especially true since one of those potential initiating events (the deployment of American nuclear-armed Jupiter missiles in Turkey in the Spring of 1962) led to a full-blown nuclear crisis (the October 1962 Cuban missile crisis) that President Kennedy estimated had at least a "one-in-three" chance of leading to war.

The thinking of Prof. Mearsheimer and other proponents of nuclear deterrence bears a similarity to that of cryptographers who point to the lack of any significant progress on factoring over the last twenty-five years as evidence that factoring algorithms appear to have hit a brick wall. The lessons I learned from applying risk analysis to nuclear deterrence paint a very different picture.

Roughly speaking, advances in factoring algorithms that doubled the key size required for the RSA algorithm (and, when extended to discrete logs, for Diffie-Hellman Key Exchange as well) occurred in 1970, 1980, and 1990. I am referring, of course, to continued fractions, sieving, and the number field sieve. However, since 1990, there have not been any similar advances.

Thinking of each decade as a coin toss that shows Heads when an advance is made and Tails when that does not occur, we have observed HHHTT and are in the process of watching the coin on its sixth toss. Since this decade is slightly more than 50 percent over, that toss looks more likely to produce T than H, but the outcome is far from certain.

Of course, advances in factoring each decade are not independent Bernoulli trials. But that is not an unreasonable model, in which case the evidence for trusting both RSA and Diffie-Hellman becomes suspect. Even if no advance in factoring is discovered during the rest of this decade, the coin tosses will be HHHTTT. A prudent person seeing a sequence like that would not dare predict that only Tails will occur into the indefinite future. For that reason, I have argued – and continue to argue – that backup systems should be in place so that yet another advance in factoring does not devastate the digital economy supported by public key cryptography.

I should emphasize that my concern about nuclear deterrence and public key cryptography does not mean that either is likely to fail in the near future. Rather, my concern is that the high confidence that a prudent person would demand over the appropriate time span seems inconsistent with the available data. Given the consequences of either failing, would 95 percent confidence over the next decade be adequate? What about 99 percent confidence over the next century?

Fortunately, the current interest in public key algorithms that are resistant to possible advances in quantum computing is pushing research in the right direction – such algorithms also will be resistant to any advances in factoring. Now we need comparable research on national security strategies. While history indicates that is unlikely to happen, I hope that including these thoughts in the ACM Turing Lecture might help.

CCS'16, October 24-28, 2016, Vienna, Austria
ACM 978-1-4503-4139-4/16/10.
http://dx.doi.org/10.1145/2976749.2976757

REFERENCES

[1] M. E. Hellman, “How risky is nuclear optimism?,” Bulletin of the Atomic Scientists Vol. 67, No. 2, pp. 47-56, March 2011. Accessible online at http://www-ee.stanford.edu/%7Ehellman/publications/75.pdf

[2] PBS NEWSHOUR, July 9, 2012, Judy Woodruff Interviews John Mearsheimer, “Nuclear-Armed Iran Would Bring ‘Stability’ But Risks.” Transcript accessible online at http://www.pbs.org/newshour/bb/world-july-dec12-iran2_07-09/

[3] M. E. Hellman, “Risk Analysis of Nuclear Deterrence,” The Bent of Tau Beta Pi, Vol. 99, No. 2, pp. 14-22, Spring 2008. Accessible online at http://www-ee.stanford.edu/%7Ehellman/publications/74.pdf