

Cybersecurity is not very important

Andrew Odlyzko

University of Minnesota

odlyzko@umn.edu

<http://www.dtc.umn.edu/~odlyzko>

Revised version, March 18, 2019.

Abstract. There is a rising tide of security breaches. There is an even faster rising tide of hysteria over the ostensible reason for these breaches, namely the deficient state of our information infrastructure. Yet the world is doing remarkably well overall, and has not suffered any of the oft-threatened giant digital catastrophes. This continuing general progress of society suggests that cyber security is not very important. Adaptations to cyberspace of techniques that worked to protect the traditional physical world have been the main means of mitigating the problems that occurred. This "chewing gum and baling wire" approach is likely to continue to be the basic method of handling problems that arise, and to provide adequate levels of security.

1 Introduction

It is time to acknowledge the wisdom of the "bean counters." For ages, multitudes of observers, including this author, have been complaining about those disdained accountants and business managers. They have been blamed for placing excessive emphasis on short-term budget constraints, treating cybersecurity as unimportant, and downplaying the risks of disaster.

With the benefit of what is now several decades of experience, we have to admit those bean counters have been right. The problems have simply not been all that serious. Further, if we step back and take a sober look, it becomes clear those problems are still not all that serious.

All along, the constant refrain has been that we need to take security seriously, and engineer our systems from the ground up to be truly secure. The recent program of recommended moves [4] opens with a quote from the famous 1970 Ware Report that called for such steps. This demand has been growing in stridency, and has been increasingly echoed by higher levels of management and of political leadership. Yet in practice over the last few decades we have seen just a gradual increase in resources devoted to cybersecurity. Action has been dominated by minor patches. No fundamental reengineering has taken place.

This essay argues that this "muddle through" approach was not as foolish as is usually claimed, and will continue to be the way we operate. Cyberinfrastructure is becoming more important. Hence intensifying efforts to keep it sufficiently secure to let the world function is justified. But this process can continue to be gradual. There is no need to panic or make

drastic changes, as the threats are manageable, and not much different from those that we cope with in the physical realm.

This essay reviews from a very high level the main factors that have allowed the world to thrive in spite of the clear lack of solid cyber security. The main conclusion is that, through incremental steps, we have in effect learned to adopt techniques from the physical world to compensate for the deficiencies of cyberspace. This conclusion is diametrically opposed to the heated rhetoric we observe in the popular media and to the unanimous opinions of the technical and professional literature. No claim is made that this process was optimal, just that it was “good enough.” Further, if we consider the threats we face, we are likely to be able to continue operating in this way. But if we look at the situation realistically, and plan accordingly, we might

- enjoy greater peace of mind
- produce better resource allocations

The analysis of this essay does lead to numerous contrarian ideas. In particular, many features of modern technologies such as “spaghetti code” or “security through obscurity,” are almost universally denigrated, as they are substantial contributors to cyber *insecurity*. But while this is true, they are also important contributors to the imperfect but adequate levels of cyber *security* that we depend on. Although a widely cited mantra is that “complexity is the enemy of security,” just the opposite is true in the world we live in, where perfect security is impossible. Complexity is an essential element of the (imperfect) security we enjoy, as will be explained in more detail later. Hence one way to improve our security is to emphasize “spaghetti code” and “security through obscurity” explicitly, and implement them in systematic and purposeful ways. In general, we should adopt the Dr. Strangelove approach, which is to

stop worrying and learn to love the bomb.

In other words, not just accept that our systems will be insecure. Recognize that insecurity often arises in systematic ways, and that some of those ways can be turned into defensive mechanisms. We do have many incremental ways to compensate, and we have to learn how to systematically deploy them, so as to live and prosper anyway. The key point is that, in cyberspace as well as in physical space,

security is not the paramount goal by itself.

Some degree of security is needed, but it is just a tool for achieving other social and economic goals.

This essay is a substantial revision and expansion of the author’s earlier piece [6], which was an extended abstract of the WiSec’10 keynote, and also builds on the author’s other papers, such as [5]. However, no originality is claimed. While this piece is likely to strike many readers as very contrarian, many of the arguments made here can also be found elsewhere, for example in [1], and are not inconsistent with many of the recommendations of mainstream reports such as [4]. Historically, for many observers a serious reassessment of the traditional search for absolute security was provoked by Dan Geer’s 1998 post [2]. However, awareness of general risk issues, and growing perception that they were key, can

be traced much further back, to various research efforts in the 1980s, and the founding of Peter Neumann’s RISKS Digest in 1985. No attempt is made here to trace this evolution of attitudes towards security. That is a nice large subject that is left for future historians to deal with. This essay considers only the current situation and likely evolution in the near future.

2 The skewed view of the world among most technologists

The critics of the standard “business as usual” approach have been presenting to the public both a promise and a threat. The promise was that with enough resources and control over system development, truly secure information technologies systems would be built. The threat was that a gigantic disaster, a “digital Pearl Harbor,” would occur otherwise.

The promise of real security was hollow. If there is anything that we can now regard as solidly established, it is that we don’t know how to build secure systems of any real complexity. (There is another factor that is not discussed here, namely that even if we could build truly secure systems, we probably could not live with them, as they would not accommodate the human desires for flexibility and ability to bend the rules. But that is a different issue not in the scope of this essay.) Serious bugs that pose major security risks are being found even in open-source software that has been around and in extensive use for years, as with the Heartbleed defect. And some insecurities, such as those revealed in the recent Meltdown and Spectre attacks, not only go back decades, but are deeply embedded in the basic architecture of modern digital processors. They cannot be eliminated easily, and we will have to live with them for many years. The most we can hope for is to mitigate their deleterious effects.

The mantra, called Linus’s Law, that “given enough eyeballs, all bugs are shallow,” has been convincingly shown to be fallacious. There are only relative degrees of security. Still, we have to remember that this has always been true with physical systems. Furthermore, in both the cyber and the physical realms, the main vulnerabilities reside in people. Those creatures are not amenable to reengineering, and are only very slightly amenable to reasoning and education.

The threat of digital catastrophe has also turned out to be hollow. Sherlock Holmes noted that the “curious incident” in the *Silver Blaze* story was that the dog did not bark. In information technology insecurity, there are two curious “incidents” that have not attracted much notice:

- Why have there been no giant cybersecurity disasters?
- Why is the world in general doing as well as it is?

Skeptics might object and point out to any number of ransomware, identity theft, and other cybercrime cases. But those have to be kept in perspective, as is argued in more detail later. There have been many far larger disasters of the non-cyber kind, such as 9/11, Hurricane Sandy, the Fukushima nuclear reactor meltdown, and the 2008 financial crash and ensuing Great Recession. Has any cyber disaster inflicted anywhere near as much damage to any large population as Hurricane Maria did to Puerto Rico in 2017?

In the cyber realm itself, we have experienced many prominent disasters. But most of them, such as airlines being grounded for hours or days, or cash machine networks not functioning, have arisen not from hostile action, but from ordinary run-of-the-mill programming bugs or human operational mistakes. And of course we have the myriad issues such as cost overruns and performance disappointments which plague information as well as other rapidly evolving technologies. They have little to do with the lack of cyber security. Yet we suffer from them every day.

There is a third curious incident in information technology (in)security that also appears to be universally ignored. For several decades we have had simple tools for strengthening security that did not require any fundamental reengineering of information systems. A very conspicuous example of such tools is two-factor authentication. The widely cited and widely accepted explanation for this technology not having been deployed more widely before is that users disliked the extra bother it involved. So apparently decision makers felt that the extra security provided by two-factor authentication did not warrant the cost of inconveniencing users. The big “dog did not bark” question then is, given that this technology was not deployed, why did nothing terrible happen?

The general conclusion of this essay is that from the start, the “bean counters” understood the basic issues better than the technologists, even though they usually did not articulate this well. The main problem all along was risk mitigation for the human world in which cyberspace played a relatively small role. It was not absolute security for the visionary cyberspace that technologists dreamed of.

3 The state of the world

One could object that the world is not doing well, and point to climate change, rising inequality, civil wars, unemployment, and other phenomena that are cited as major ills of our society. But that has to be kept in perspective. Let’s put aside, until the next section, questions about issues such as long-term sustainability of our civilization. If we just look at where the human race is today from a long-term historical perspective, we find stunning advances by many measures, such as the number of people on Earth, how long they live, and how educated they are. There are more people today who are obese than hungry, which is unprecedented. Obesity is certainly not ideal, but can easily be argued to be an advance on the historically dominant feature of human lives.

Of course, there are a variety of threats for the future. But we need to remember that the progress that has occurred has relied often and in crucial ways on information systems that were, and are, insecure. Further, almost all of the most serious threats, to be considered next, are little affected by cyber security or lack of it.

4 Threats

We certainly do face many threats. In particular, we do face many cyberthreats. It seems inevitable that we will suffer a “digital Pearl Harbor.” What we have to keep in mind is that we have suffered a physical Pearl Harbor and other non-cyber disasters that large or

larger. Many occurred quite recently, as noted before. It seems absolutely certain we will suffer many more, and an increasing number of them will surely be coming from the cyber realm. On the other hand, it is questionable whether the cyber threats are yet the most urgent ones.

The human race faces many potentially devastating non-cyber dangers, such as asteroid strikes, runaway global warming, and large pandemics. These threats could have giant impacts, but are hard to predict and quantify, and are seemingly remote, so tend to be ignored by almost all people most of the time. However, we also face a variety of other still large dangers, such as those from earthquakes and hurricanes. Those occur more frequently, so the damage they cause is moderately predictable, at least in a long-run statistical sense. Yet we are not doing anywhere near as much to protect against them as we could, if we wanted to do so. We accept that they will occur, and rely on general resilience and insurance, whether of the standard variety, or the implicit insurance of governments stepping in with rescue and recovery assistance.

We also tolerate the ongoing slaughter of over a million people each year in automobile accidents worldwide (with about 40,000 in the U.S. alone). The horrendous losses of human life as well as property that involve cars arise mostly from unintentional mistakes. They result from our accepting the limitations of *Homo sapiens* when dealing with a dangerous technology. It's just that this technology has proven extremely attractive to our species. Hence we accept the collateral damage that results from its use, even though it far exceeds that from all wars and civil conflicts of recent times.

On top of accidents we also have the constant ongoing malicious damage, coming from crime in its many dimensions. Society suffers large losses all the time, and mitigates the threat, but has never been able to eliminate it. We have large security forces, criminal courts, jails, and so on. The U.S. alone has close to a million uniformed police officers, and more than a million private security guards.

Military establishments tend to be substantially larger than law enforcement ones. The main justification for them is to guard against the far rarer but potentially more damaging actions of hostile nations. One way or another, most societies have decided to prioritize protection against those external dangers over that of internal crime. Further, in recent decades, military spending (and therefore total security-related spending) has been declining as a fraction of the world's economic output. So when societies feel threatened enough, they do manage to put far more effort into security than is the case today.

Yet even military security at its very best is not water-tight, which has to be kept in mind when considering cyber security. Serious gaps have been uncovered on numerous occasions, such as a deep penetration of an American nuclear weapons facility by a pacifist group that included an 82-year old nun.

The bottom line is that society has always been devoting substantial and sometimes huge resources to security without ever achieving complete security. But those resources are still not as great as they could be. That's because, as noted above, *security is not the paramount goal by itself*. We make tradeoffs, and are only willing to give up a fraction of the goods and services we produce for greater safety. There is even extensive evidence

for human desire for a certain level of risk in their lives. When some safety measures are introduced, people compensate for that by behaving with less care.

Still, we do employ many people and extensive resources protecting ourselves from traditional physical world threats, far more than we devote to cybersecurity. Hence it is clear, and has been clear for a long time, that more effort could have been dedicated to cybersecurity, even without consuming productive resources. All we had to do was just shift some of the effort devoted to traditional physical security to the cyber realm. And indeed that is what is happening now, at least in relative sense. More attention and resources is being devoted to cybersecurity. One measure of the greater stress being placed on this area is the growing (but still very small) number of CEOs who have lost their jobs as result of security breaches. So the question arises, essentially the same question as before, just in a different form: Why was this not done before, and why has not much harm come from this?

5 Humanspace versus cyberspace

It is very hard for technologists to give up the idea of absolute cybersecurity. Their mind set is naturally attracted to the binary secure/insecure classification. They are also used to the idea of security being fragile. They are not used to thinking that even a sieve can hold water to an extent adequate for many purposes. The dominant mantra is that “a chain is only as strong as its weakest link.” Yet that is probably not the appropriate metaphor. It is better to think of a net. Although it has many holes, it can often still perform adequately for either catching fish or limiting inflow of birds or insects.

Technologists also tend to think of information systems as isolated. This attitude is represented beautifully by the famous 1996 creation of John Perry Barlow, “A Declaration of the Independence of Cyberspace.” This proclamation, which today seems outlandishly ludicrous, proclaimed the existence of a new realm, ‘Cyberspace,’ that was divorced from the physical world, and did not need or want traditional governments or other institutions. The key assumption was nicely formulated in the oft-quoted passage:

Cyberspace consists of transactions, relationships, and thought itself, arrayed like a standing wave in the web of our communications. Ours is a world that is both everywhere and nowhere, but it is not where bodies live.

Indeed, if cyberspace were totally divorced from human space, and if all the “transactions, relationships, and thought itself” depended just on some mathematical relationships, then cybersecurity would be of paramount importance. An opponent utilizing a clever mathematical idea to break a public-key system, or stealing a password, might wreak unlimited havoc.

And indeed, as the increasing number of incidents with Bitcoin and other cryptocurrencies proves, such dangers do lurk in pure cyber realms. Further, they cannot be avoided. As was discussed before, people are incapable of building completely secure systems, they do choose weak passwords or leak strong ones, they do fall prey to phishing attacks, and every once in a while a mathematical breakthrough does demolish a cryptosystem.

What makes our lives tolerable is that the Barlow vision is divorced from reality. Cyberspace is intimately tied to what we might call Humanspace, the convoluted world of physical objects and multiple relations, including institutions such as governments, and laws, and lawyers. In fact, we can say:

The dream of people like Barlow was to build a Cyberspace that would overcome the perceived defects of Humanspace. In practice we have used the defensive mechanisms of Humanspace to compensate for the defects of Cyberspace.

Those defensive mechanisms is what we consider next, starting with the limitations of attackers in both physical and cyber realms.

6 Pluses and minuses of natural stupidity

There are extensive ongoing discussions about the promises and threats of artificial intelligence (AI). Much less is said about natural stupidity and its positive aspects. Yet it is central to human life, and key to enabling society to function. (At an even more basic level, the astounding level of human credulity, which enables so many attacks, is an essential element of human psychology and sociology, and enables the cooperation that has led to modern civilization.) In particular, we are alive and living pretty well largely because

most criminals are stupid.

This includes terrorists. Most of them are stupid, too. They are in almost all cases more like the Shoe Bomber than the highly trained and highly proficient professionals that the multitudes of publicly prominent cyber Cassandras hold out as big threats to our lives. Most crimes are extremely mundane, and many more could easily be solved if more effort was devoted to them. Criminals constantly make foolish mistakes, such as leaving their fingerprints, or their DNA, on the scene, or driving their own cars. As a result, general crime has been kept within tolerable bounds for most of human history.

It is not just the most stupid people who make mistakes. Everyone does so. In fact, the mistakes of the smartest individuals are often the most disastrous, as they get entrusted with the most important jobs. Even the highly trained and highly proficient professionals in the military and intelligence agencies are fallible, including when at the peak of training and preparation. It is this fallibility that helps make cyberspace more similar to physical space than is commonly thought. Detecting where a network attack originates is harder than detecting where a ballistic missile is launched from. But digital forensics is a thriving field, largely because of human mistakes. Even the Stuxnet creators were not able to completely erase their “digital fingerprints,” leading to high confidence as to their identities.

Cybercrimes not only leave digital fingerprints. They are usually tied in one way or another to the physical world, most frequently through flows of money. Hence there are far more ways to trace them than would be the case if they happened purely in cyberspace. Once tracing is possible, measures to deter, prevent, and punish can be brought to bear. Those digital fingerprints also mean that natural stupidity of attackers has more opportunities to display itself. And that offers opportunities for defense and countermeasures, just as in the traditional environment.

7 Smart and stupid criminals

The reasons most criminals are stupid are worth considering. An important one is that we mostly hear of the criminals who get caught, and that is not a perfectly representative sample. The smart ones avoid detection and capture. But the really smart ones mostly figure out it is far safer and more comfortable to stay close to the line of legality. Serious damage to the system as a whole, or even to many individual players, tends to provoke strong countermeasures. Some criminals even learn to be symbiotes, and contribute positively to society.

An insightful analogy can be drawn with biology. A virus that kills the host instantly tends to perish, as it has little chance to spread. The more successful viruses (more successful in terms of being widespread) are like those for the common cold, which cause relatively small annoyances that serve primarily to help them propagate. Many parasites evolve to become symbiotes, and the study of commensal relationships is a thriving field with a variety of examples.

8 The cybercrime ecosystem

Most criminals, even among those on the extreme edge of the stupidity spectrum, have no interest in destroying the system they are abusing. They just want to exploit it, to extract value for themselves out of it.

An amusing and instructive example of illicit cyber behavior that maintains the functioning of the system is provided by the ransomware criminals. Studies have documented the high level of “customer care” they typically provide. They tend to give expert assistance to victims who do pay up, and have difficulty restoring their computers to the original state. After all, those criminals do want to establish “reputations” that will induce future victims to believe that payment of the demanded ransom will give them back control of their system and enable them to go on with their lives and jobs.

An extreme example of exploitation of cyber insecurity without causing noticeable damage is that of national intelligence agencies. They carry out extensive penetrations of a variety of government and commercial systems, but are usually just after limited pieces of information, and try (and usually succeed) in staying inconspicuous. In most cases they exploit only a tiny fraction of what they acquire, precisely in order not to raise suspicions about their activities. Of course, their activities do create other dangers, when they acquire control of systems for future large-scale hostile activities. But such penetrations by state actors have to be handled at state levels, similarly to what occurs in the physical realm.

There are certainly some malicious actors who simply want to inflict damage, whether it is against a person against whom they have a grudge, or, especially in case of terrorists, against society at large. But even such people are generally not as dangerous in cyberspace as they could be. First of all, there are not that many of them. Second, they generally have limited skills and resources, and are mostly very foolish, and engage in foolish activities. The more rational among them choose their targets and methods for maximal effectiveness in achieving whatever nefarious purposes they have in mind. For terrorists, say, cyberspace is generally not very attractive as a target. Blocking people from withdrawing money from

cash machines, or even causing a blackout in a city does not carry as strong a message as blowing up airplanes, bringing down buildings, or causing blood to flow among spectators in a sports arena.

There is much concern about ongoing technology developments making the lack of cyber security far more dangerous, especially as more devices go online, and IoT (the Internet of Things) becomes more pervasive. Those are valid concerns, but let us keep in mind that those ongoing technology developments are also creating or magnifying many physical dangers even without taking advantage of cyber insecurity. Just think of drones (or possibly imaginary drone sightings) shutting down airports recently, or drones or self-driving cars delivering bombs in the future.

In general, and reinforcing earlier discussions, society has always faced manifold dangers from its members misusing various technologies. Deterrence, detection, punishment, in addition to general social norms, is what has enable civilized human life to exist. Contrary to the cyberlibertarian visions of people like Barlow (or many modern advocates of Bitcoin and blockchain) they are likely to be just as crucial in the future, if not more so.

Of course, as the old saying goes, bank robbers went after banks because that is where the money was. But now the money is in cyberspace. So that is where criminals are moving. And that is also where security resources are being redirected. Completely natural and expected, and happening at a measured pace.

9 Black swans vs. long tails

Cybersecurity efforts are dominated by very mundane work, monitoring the automated probes of the network, or attacks of the “script kiddies.” And, perhaps most prominent and most boring, but absolutely critical: assisting legitimate users who have forgotten their passwords. Which is exactly analogous to the state of traditional physical security. Much of the time of firefighters and police officers is devoted to rescuing kittens stuck high up trees, or handling temporarily inebriated but otherwise perfectly respectable citizens.

The evolution of the cybersecurity field over the last few decades has led to wide recognition among its practitioners that threats cannot be entirely eliminated. There are frequent references to minimizing “the attack surface,” for example. This reflects the reality that one can limit attacks and the damage they can do, but not get rid of them. More resources can be used to lessen threats. But those resources are costly, either in terms of the pay and equipment of the security professionals, or, what is typically much more important, in terms of constraints on the legitimate users. So one is led to look at optimizing the allocation of resources and studying and modifying the incentives. One outgrowth of such thinking on the academic side has been the rise of the field of economics of information security. It has produced a flourishing literature and a series of annual workshops. Together with all other academic and industry efforts, it fits into the basic philosophy that animates modern economics, namely of studying systems in equilibrium. There is ongoing hostile activity that is counteracted by security measures, and the task is to select the optimal combination of those measures that fits within some budget constraints.

One could view such approaches as concentration on the “long tail” of security threats. There are many of them, they require large resources in the aggregate to deal with, but indi-

vidually they pose limited and reasonably well understood dangers. Overall, their potential impact can be estimated and constrained by standard approaches.

At the other end of the spectrum, though, there are the “black swans,” the giant security breaches that cause major damage. Those don’t fit into the equilibrium framework, just as catastrophic financial collapses don’t fit into the standard economic equilibrium framework (and have been almost entirely ignored by mainstream economists). But neither do the giant physical disasters, such as Pearl Harbor or Hurricane Katrina. Their damaging effects basically can only be mitigated by designing in general resilience.

Measures that provide resilience against cyber attacks are often the same as those against traditional physical attacks or against natural disasters. As just one example, there is much concern about the damage to the electric power grid that might be caused by malicious actors. But the worst scenarios along those lines are similar to what we are sure to suffer when something like the Carrington Event recurs. This was the giant geomagnetic solar storm that hit the Earth in 1859. It caused widespread failures of the telegraphs, the only electrical grids in existence at that time. Estimates are that if it were to occur today, it would cause damages in the trillions of dollars. And it is bound to recur some day!

The conclusion that emerges is again that cyberspace is not all that different from the more traditional physical space we are more used to. And security measure for the two are again similar.

10 Neglect of obvious security measures

The main thesis of this note, that cybersecurity is not very important, is illustrated nicely by the phenomenon of two-factor authentication. This technique is spreading. It is not a panacea, but there is general agreement that it offers significant enhancement to security.

But why is it only now that two-factor authentication is coming into widespread use? The basic technique is ancient by the standards of the information technology industry. Two and a half decades ago it was used at my employer of that time. The hardware tokens came from one of several suppliers that were already in that line of business.

Yet even at my former employer, two-factor authentication was abandoned after a while, and in most places it was never put into service in that era. So what has changed to finally make this technology used more widely? As often happens, it was likely a combination of factors:

- threats have increased
- implementing two-factor authentication has become easier

The old hardware tokens of the 1990s were not very expensive, but they had to be carried around (as opposed to receiving a text on a mobile phone that people have with them almost all the time, say), and they required typing in strings of arbitrary symbols. Now we can use short texts, or hardware tokens that plug into a computer, or else mobile phones that communicate with a nearby computer wirelessly. So while the monetary costs of the basic system have not changed dramatically, the costs to users have declined significantly. And, of course, the threats have increased, as noted above, so the incentives to use two-factor authentication have grown.

Yet even now, two-factor authentication is nowhere near universal. Further, most deployments of it at this time appear to use the least secure version of it, with texts to mobile phones. Practical attacks on this version have been developed and applied. The more secure versions with hardware tokens are used much less frequently. Obviously what is happening is that choices are being made, the additional inconvenience to users being weighed against the likely losses from hostile penetrations. Even without any new technology breakthroughs, more secure versions of two-factor authentication can be deployed when they are seen as necessary. But they are clearly not being seen as necessary at present.

There are many more examples of relatively easy steps that have been available for a long time, and can strengthen security without any fundamental reengineering of information systems, or rearranging how society functions. Consider the adoption of chip credit cards. They have been universal in much of the world for years, but are only now taking over in the U.S. The costs have been understood by the banking industry, and it was decided, through a messy process by various stakeholders, that they were too high until the perceived threats increased.

Electronic voting is another prominent example where simple and well-known steps would have provided greater security a long time ago. Experts have been arguing from the start that purely electronic voting basically cannot be made secure, at least not with feasible technology and the financial resources that are available or are likely to be made available. All the evidence that has been gathered over the years supports this view. Further, all the advantages of electronic voting (convenience, accessibility for those with handicaps, quick collection of results, ...) can be obtained very easily, together with a much higher degree of security, through the use of printed records that are preserved in physical form. The additional costs that are involved are very modest, and seem well worth it to most people who have examined the situation, including this author. Yet in many jurisdictions this simple solution is being ignored. And it has to be admitted that so far no serious abuses have been documented. What is likely to happen is that if some big scandal surfaces that is based on a cyber breach, political leaders will swing into action, and find the resources to provide the obvious solution. (We should remember that big voting scandals do occur all the time, based on other aspects of the voting system, and they lead to responses that vary with circumstances.) But, as seems typical in human affairs, it will likely take a big scandal to cause this to happen.

Electronic voting provides an interesting illustration of a cyber insecurity that is not difficult to fix, but is not being fixed. It also provides an example of a common phenomenon, namely that the fix involves stepping back to the traditional physical world, in this case of messy paper ballots. (The same could be said of chip cards.) In other words, the insecurity of the cyber realm is compensated by a measure from the brick-and-mortar world.

An even better example of reliance on physical world to compensate for defects in cyber security is that of passwords. They have been pronounced obsolete and dead many times, but are still ubiquitous. A key element in making them more tolerable in spite of their well-known weaknesses is the use of paper for users to write them down (or, preferably, to write down hints for those passwords or passphrases). The security field has finally been forced to admit that asking users to remember scores of complicated passwords (and change them

every few months) is not going to work, not with the bulk of human users. But paper slips work out quite well, as physical wallets and purses do not get stolen all that often.

Notice that there are many other direct physical methods for increasing security. Air-gapped systems, isolated from the Internet, have been standard in high-security environments. They are again not absolutely secure, as the Stuxnet case demonstrates. But they do provide very high levels of security, as breaching them requires special skills and extensive effort (as the Stuxnet case demonstrates, again). At a simpler level, allowing certain operations (such as resetting the options on a router or another device) only through the press of a physical button on the device also limits what attackers can do.

Frequent backups serve to mitigate ransomware and many other attacks. They can be automated, so that they do not impose any significant mental transaction costs on the users. They increase the reversibility of actions, which is a key component to security (but seems not to be understood by the advocates of Bitcoin and other cryptocurrencies). And they are not expensive in terms of hardware. Of course, backups increase security only if they are not subverted. But there are a variety of ways to make backups more trustworthy, such as using write-only media (such as some optical disks), or special controllers that limit what operations can be done.

We should also remember there is one piece of advice that applies in both cyberspace and physical space: If it's dangerous, don't use it! Some very cautious organizations disable USB ports on their computers, but such organizations are rare. Email attachments are a notorious carrier for all sorts of malicious software. They could be blocked, but seldom are. All these examples show how society has in effect accepted obvious risks in order to get benefits of insecure information technology solutions.

11 Surveillance capitalism and loss of privacy

The analogy between cyber and physical security is strong, but there are certainly substantial differences. The one that appears to be cited most frequently is privacy. There was no absolute privacy in the past. In particular, there was always the most intractable problem of all, namely that of insider disclosure. (According to an old saying, “two people can keep a secret, as long as one of them is dead.”) But modern threats to privacy are orders of magnitude larger than those faced in the past. Further, as we move forward, our central and giant problem is that potential leakers are proliferating at a rapid pace. Individuals can convey far more information now than in the past, as the Manning, Martin, and Snowden information torrents from NSA demonstrate. For the majority of people, though, the main threat comes in the shape of the many devices we use, which are increasing in numbers and in their capability to transmit information about us to others. The cell phone is the premier example, but increasingly so is our fitness tracker, our TV set, and our electric meter. Practically nothing that we will be doing can be assumed to be secret in the future. This will even apply to our physiological reactions, even ones we do not express, or may not consciously be aware of, since they might be discerned by various sensors.

Already today, the old mantra that “on the Internet, nobody knows you are a dog,” has in practice been turned on its head. Many organizations know not only that you are a dog, but also what breed of dog you are, and what kind of fleas you have.

For the purposes of this essay, the key counterpoint to this line of argument is that this erosion of privacy we experience has little to do with cyber insecurity. Some of that erosion does come from illicit hacking of our systems, which is indeed facilitated by the insecurity of our information systems. But most of it comes by design, as providers of services and devices purposely build them to collect data about users for exploitation by those providers and their (almost universally concealed) networks of partners. (Even the illicit hacking of those devices, databases, and so on, can occur only because of this huge and legal, even though usually obfuscated, data gathering.) Hence there are no improvements in cybersecurity that would by themselves make a measurable difference to the erosion of privacy that we experience. To the extent that society wants to preserve some semblance of privacy, other methods will have to be used, which likely will have to be based on laws and regulations, and to some extent on technologies for users to protect themselves.

On the other hand, the erosion of privacy is a key element to maintaining tolerable levels of security in general. Tens or sometimes hundreds of millions of credit cards are routinely captured by criminals by compromises of databases. Yet the overall damages are limited, and often dominated by the cost of arranging for replacement cards. The prices of stolen credit card credentials on the black market are low, on the order of a dollar or so each. The reason is that banks have developed techniques for detecting credit card fraud. Those are based on knowledge of users' patterns of behavior. A typical card holder is not an anonymous "standing wave" of Barlow's imagination, or some account even more anonymous than those involved in the not-all-that anonymous Bitcoin operations. Instead, such a person is usually an individual who follows a staid routine in life and in commercial transactions, say stopping by a particular coffee shop on the way to work, or dropping in at a grocery store on the way back from work.

There are many measures that erode privacy, such as cross-device tracking (in which users are identified even though they use different gadgets) or identifying users by the patterns of their typing, that are often regarded as objectionable or even creepy. Yet they do serve to identify users, and thereby to prevent mischief, even if this is incidental to the main purposes for which they are deployed. Organizations that operate these systems can get a high degree of assurance as to the person they are dealing with, and in such circumstances stealing a credit card or cracking a password is often of limited use.

It should also be remembered that since enterprises do want to track customers or potential customers for their own business reasons, they have incentives to develop and deploy those privacy-invasive methods in preference to providing more direct security. This is a case where general economic incentives skew what security methods are used. But those methods are very effective in compensating for cyber insecurity.

12 The deceptively transparent but opaque world

The development of information technology does mean that nothing can be assured of staying secret. (The Manning, Martin, and Snowden security breaches at NSA cited above are only some of the most prominent examples.) There are just too many vulnerabilities in our systems, and too many tools to capture and extract information, such as cameras in our cell phones, and miniature cameras that are getting ever smaller and harder to detect.

But neither can it be assumed that all relevant information will be available in forms that lead to proper action. The technique of “hiding in plain sight” was popularized by Edgar Allan Poe two centuries ago. Modern technology creates so much more information that this often works with minimal efforts at concealment, or even without any such effort. Even when information is known, it is often not known widely, and is not known by people who might or should act on it. Just consider Dieselgate, where various groups had obtained measurements of emissions exceeding legal limits years before the scandal erupted. Or think of the Danish bank that laundered over \$200 billion through a small Estonian branch over a few years. Not to mention all the various sexual harassment cases that took ages to be noticed publicly.

In general, information that can be captured by information systems is becoming more detailed and far more extensive. But it is still limited in many ways. One of the most important ones is that human society is a messy affair, and much that goes on is hard to codify precisely. In particular, tacit knowledge is crucial for individuals and organizations. Hence even complete penetrations of computer systems of an organization are seldom sufficient to be able to replicate that organization’s functioning. Studies have been carried out on the effects of East German espionage in West Germany. It was extremely effective at penetrating almost all targeted commercial organizations. But it allowed only a small narrowing in the performance gap between East and West German companies in the same industry. Especially when technology is advancing rapidly, the time to fully exploit information about current state of the art means that the intruders, who acquire the formal knowledge that is recorded, end up behind when they master those technologies.

As technology advances, the level of information that can be acquired increases, and so one might argue that the importance of tacit knowledge decreases. But that is very questionable. Systems are increasingly complicated, so it is harder to formally describe their functioning and their various failure modes and special features.

Further, modern technology allows for significant enhancements to the basic technique of “hiding in plain sight.” Obfuscation techniques can be improved, and deployed much more widely and systematically, since we have increasing ability to create fake information. Looking forward, we are likely to see an arms race, with AI systems used to create “alternate realities” on one hand, and to try to penetrate and deconstruct them on the other. The “post-truth” world is regarded as a danger, but it seems inevitable, and does have positive angles.

Note that there are many examples of primitive instances of such developments. The impenetrable legalese in the Terms of Service that users have to accept to use online services is a frequently encountered instance of what one recent paper referred to as “transparency [as] the new opacity.” In general, “speed bumps,” steps which offer some protection, rather than absolute security, proliferate. Non-Disclosure Agreements, or NDAs, are one such example. Silicon Valley, home of both privacy-abusers and transparency advocates, uses them widely. Though far from impenetrable, NDAs do substantially limit the spread and use of information.

13 The virtues of messiness

Lack of cyber security is universally regarded as just one aspect of the generally poor quality of our software, much of which is blamed on the “spaghetti code” nature of that software. But one should note that this poor quality also has positive aspects. Software piracy is not all that serious a problem, for example. Unpatched systems that are exposed on the Internet get easily penetrated. So frequent patching is required, and that means the software producer has to be in contact with systems running that code, and has a handle on illicit copies. Further, systems that are barely stable, and require constant upgrades to deal with bugs and improve functionality cannot be easily adopted by competitors, in another aspect of the tacit knowledge argument.

At a more mundane level, messiness of code, along with logging, is the primary reason digital forensics is as effective as it is. Attackers have difficulty covering up their traces. Much more can be done in this direction through intentional design.

Note that there are already successful examples of such approaches in the physical world. For example, color copiers generally have Machine Identification Codes (MICs) which leave a digital watermark on every page, identifying the printer and the date. (This case provides also another instance of successful “security through obscurity,” since this technology was in wide commercial use for almost two decades, and was not particularly secret, before it was widely publicized.)

A related approach is that of protecting consumer transactions by using diverse communication channels. Banks increasingly require confirmation of large and suspicious transactions through voice calls or texts. Not as simple, quick, and cheap as letting web entries go through, but capable of deployment in a flexible fashion, depending on the level of risk.

14 Speed, reach, and cost for offense and defense

At a very high level, information technologies have been revolutionary primarily because they offered quantum leaps in the three main measures of infrastructure effectiveness. They enabled actions or communications to be carried out much faster than was feasible before. They also allowed actions or communications to take place on a much wider scale. Finally, they did all of this at much lower cost.

These same advantages of information technologies, which led to so much progress in society, have also been attractive to criminals. Expert burglars could get into practically any dwelling, but it would usually take them some time to do so for every place. Automated probes can find and penetrate unpatched computers in seconds. Even an accomplished burglar needs some time, minutes or more typically hours, to rob a house. Hackers can commandeer thousands or even millions of computers in that time. Finally, all those attacks can be carried out at very low cost by hackers, who often don’t even need much in the way of computers, as they can rely on ones they manage to seize control of.

But those same advantages of information technologies have also aided defense (just as happened with numerous earlier technologies). Defense can act much faster, as communication channels can be blocked, or software patched, far faster than physical locks could be changed. Centralized defense teams can provide security for global organizations, without

the need to station an armed guard at each location. And the costs are far lower than for physical protective measures.

Finally, there is that basic approach that was mentioned before: If it's too dangerous, don't use it. If high speed is a problem (as it is, as cryptocurrency enthusiasts keep discovering over and over, and fail to learn from), slow things down. Don't allow large money transfers to occur until a day or two have passed, and there is a chance for monitoring systems (possibly ones involving loss of privacy) to collect and analyze data about the behavior of the entities involved. And so on.

These basic techniques underlie the usual approach taken by operators when faced with serious problems: Bring down the network, repair (by reinstalling basic operating systems if necessary) all the machines that might be affected, and start bringing up functionality in sections of the network. That is how the now-ancient Morris Worm infestation was dealt with. It is also how the collapse of campus network at a prestigious college was cured recently [3]. The ability of modern technology to operate in a decentralized fashion, with multiple ways of providing at least some basic functionality, is very helpful. As the report on that college's information systems debacle notes, when the basic network stopped functioning, the people involved "got creative." Not something that one would undertake voluntarily. But it demonstrates the resilience of the system, and, among other things, makes it that much less attractive a target for attackers.

15 The increasingly ambiguous notion of security

Obfuscation, cited earlier, whether deliberate or accidental, will surely be an unavoidable and prominent feature of the "post-truth" world we are moving into. This world, full of information and misinformation, will create new challenges for security. To repeat the point made before, security is not the paramount goal by itself. But even beyond that dictum, we have to deal with the most fundamental questions of what security is, and how it is to be provided. Increasingly it is not just about keeping out some well-defined "bad guys" out of the physical or cyber systems of an organization. The erosion of individual privacy tends to overshadow in the public mind the general explosion of information about organizations. Customers, suppliers, and partners legitimately possess an immense amount of information about any given enterprise. This information is being assembled in easily accessible format (for example, in the various customer relationship packages), which makes it easier to acquire and exploit. Therefore any enterprise is becoming less of a cohesive and isolated entity (physical or cyber), and more like a diaphanous web that overlaps other similar diaphanous webs. The problem of security in such a setting is then of managing the information flows to and from numerous other organizations, a much harder task than keeping out burglars or terrorists from a building.

In addition, security has always involved a very large dose of what Bruce Schneier has called "security theater." Security is often more about perceptions of security than about any quantifiable and solidly established measures of security. Therefore security will increasingly overlap with public relations, and the generation of "spin."

16 Conclusions

This essay is a brief and very high level view of the cybersecurity area, in particular of how society has managed to thrive in spite of reliance on insecure information systems. The main conclusion is that, contrary to the public perception and many calls from prominent business and government leaders, we are not facing a crisis. This does not mean, though, that cybersecurity can be neglected, nor that all the effort that has been devoted to new security technologies has been wasted. Threats have been proliferating, and attackers have been getting more sophisticated. Hence new measures need to be developed and deployed. Firewalls are widely claimed to be becoming irrelevant. But they have been very useful in limiting threats over the last few decades. Now, though, we have to migrate to new approaches.

Furthermore, just as in the physical realm, dramatically different levels of security are called for in different organizations. The military and the intelligence agencies can naturally be expected and required to devote far more attention and resources to security than civilian enterprises. And they can also be called upon to deal with powerful state actors that threaten ordinary businesses. We don't expect hotels to protect against foreign agents bringing rare and ultra-lethal poison agents to the premises. That is what government agencies are for, as they can marshal the expertise and resources to deal with such threats.

Still, much can be done even at the level of small civilian enterprises. We do not know how to build secure systems of substantial complexity. But we can build very secure systems of limited functionality. Those can be deployed for specialized purposes, such as monitoring large systems for signs of penetrations or corruptions, or ensuring integrity of backups.

We can also improve laws, regulations, and security standards. Cybersecurity is particularly rife with problems arising from the "tragedy of the commons" and negative externalities, and those problems can be mitigated. Microsoft dramatically improved the security of its products early in this century as a result of pressure from customers. Much more can be done this way. For example, it has been known that it is important to perform array bound checking, and how to do it, for half a century. It would not be too difficult to close that notorious hole that is key to numerous exploits.

The buffer overrun issue cited above brings up one of the main points of this essay, namely that there are many ways to improve cybersecurity even without new inventions. As a recent piece notes, "[m]ost of our security vulnerabilities arises from poor practice, not from inadequate technology" [1]. What that means is that one has to be modest in expectations for anything truly novel. It may be a worthwhile goal to try for a "moonshot" or "silver bullet" technological solution, in order to inspire the designers. But even if some dramatic breakthrough is achieved, it will still have to compete with a slew of other, more modest "Band-Aid" style approaches. So other factor than pure effectiveness, such as ease of use, may easily dominate, and result in slow or no adoption.

This essay does suggest some contrarian ideas for increasing security. They are based on increasing complexity, to enable many of the "speed bumps" that limit what attackers can do and help trace them. "Spaghetti code" has already been helpful, and can be deployed in more systematic ways. In general, we should develop what Hilarie Orman has suggested calling a "theory of bandaids."

This essay does not claim that a “digital Pearl Harbor” will not take place. One, or more, almost surely will. But that has to be viewed in perspective. Given our inability to build secure system, such events are bound happen in any case. So all we can affect is their frequency and severity, just as with large physical dangers. Further, the likelihood of a “digital Pearl Harbor” has to be considered in comparison to all the other threats we face. The issue is risk management, deciding how much resources to devote to various areas.

Acknowledgments

The author thanks Ross Anderson, Steve Bellovin, Dorothy Denning, Ben Gaucherin, Balachander Krishnamurthy, Peter Neumann, Hilarie Orman, Walter Shaub, Robert Sloan, Bart Stuck, Phil Venables, Richard Warner, Bill Woodcock, and the editors of *Ubiquity* (Rob Akscyn, Kemal Delic, Peter Denning, Ted Lewis, and Walter Tichy) for their comments. Their providing comments should not be interpreted as any degree of endorsement of the thesis of this essay.

References

1. P. J. Denning, “The Profession of IT: An interview with William Hugh Murray,” *Communications of the ACM*, vol. 62, no. 3, March 2019, pp. 28–30. Available at <https://cacm.acm.org/magazines/2019/3/234920-an-interview-with-william-hugh-murray>.
2. D. Geer, “Risk management is where the money is,” *Risks Digest*, vol. 20, issue 06, Nov. 12, 1998, available at <https://catless.ncl.ac.uk/Risks/20/06>.
3. L. McKenzie, “Amherst students incredulous about going for days without services they consider absolute necessities,” *InsideHigherEd*, Feb. 21, 2019, <https://www.insidehighered.com/news/2019/02/21/almost-week-no-internet-amherst-college>.
4. New York Cyber Task Force, “Building a defensible cyberspace,” Sept. 2017 report, available at <https://sipa.columbia.edu/ideas-lab/techpolicy/building-defensible-cyberspace>.
5. A. Odlyzko, “Cryptographic abundance and pervasive computing,” *iMP: Information Impacts Magazine*, June 2000, available at https://web.archive.org/web/20030415005519/http://www.cisp.org/imp/june_2000/06_00odlyzko-insight.htm.
6. A. Odlyzko, “Providing security with insecure systems,” Extended abstract. *WiSec’10: Proceedings of the Third ACM Conference on Wireless Network Security*, ACM, 2010, pp. 87–88. Available at <http://www.dtc.umn.edu/~odlyzko/doc/wisec2010.pdf>.