INVITED PAPER

Privacy and Cybersecurity: The Next 100 Years

This paper discusses the interrelationship of privacy and cybersecurity and provides brief retrospectives and speculative views of how the future may look for theory and algorithms, technology, policy, and economics from the point of view of five different researchers.

By Carl Landwehr, *Senior Member IEEE*, Dan Boneh, John C. Mitchell, Steven M. Bellovin, Susan Landau, *Member IEEE*, and Michael E. Lesk, *Member IEEE*

ABSTRACT | The past and the future of privacy and cybersecurity are addressed from four perspectives, by different authors: theory and algorithms, technology, policy, and economics. Each author considers the role of the threat from the corresponding perspective, and each adopts an individual tone, ranging from a relatively serious look at the prospects for improvement in underlying theory and algorithms to more lighthearted considerations of the unpredictable futures of policy and economics.

KEYWORDS | Cryptography; cybersecurity; privacy; security

I. INTRODUCTION (CARL LANDWEHR)

Why are privacy and cybersecurity joined? Perhaps because without some degree of privacy (control of release of personal data), people do not generally feel secure, and unless the security (confidentiality, integrity, and availability) of the data can be assured, control over the data is an illusion. Concerns about security and privacy in the context of information stored, processed, and transmitted by computers are nearly as old as the computing profession itself: the desire to break cryptographic codes (i.e., to violate the confidentiality of messages) motivated some of the earliest and most intensive developments of digital computers. As computers have insinuated themselves into sensors and control systems and society's dependence on them has grown, both the possibility and the reality of physical damage resulting from exploitation of security flaws have been demonstrated repeatedly.

The breadth of this topic is extraordinary. Through cryptography, it touches the algorithmic foundations of computer science. In aiming to build systems with as few security flaws as possible, it places strong demands on software architecture, software engineering, system engineering, and computing technology generally. All fields of engineering must contend with costs, but because security is often seen as an option, and frequently an unrecoverable engineering cost that may even impede system functions, the economics of privacy and cybersecurity is a key factor in determining deployment of even those technologies whose effectiveness has been demonstrated. Finally, without a specified security or privacy policy, systems never have security or privacy violations, only surprises. Hence policy, including public policy about allowed and prohibited data flows, is also an essential part of both the past and the future of privacy and cybersecurity.

An individual attacking a system to steal money or secrets, to plant illicit software, or to gain system control has to find a way to break in. Today, most successful attacks simply exploit flaws that were inadvertently placed in the system during development and were not removed by industrial quality control mechanisms prior to distribution. Usually, though not always, these flaws reflect errors in software implementations rather than fundamental design

Manuscript received January 23, 2012; revised January 27, 2012; accepted January 27, 2012. Date of publication April 12, 2012; date of current version May 10, 2012.
C. Landwehr was with the Institute for Systems Research, University of Maryland, College Park, MD 20742 USA. He is now an independent consultant.
D. Boneh and **J. C. Mitchell** are with the Department of Computer Science, Stanford University, Stanford, CA 94305-9045 USA.
S. M. Bellovin is with the Department of Computer Science, Columbia University, New York, NY 10027-7003 USA.
S. Landau is with the Department of Computer Science, Harvard University, Cambridge, MA 02138 USA.
M. Lesk is with the School of Communication, Information, and Library Science, Rutgers University, New Brunswick, NJ 08901-8554 USA.
Digital Object Identifier: 10.1109/JPROC.2012.2189794

flaws in, for example, cryptographic algorithms and protocols.

In the following sections, five distinguished researchers provide brief retrospectives and speculative views of how the future may look for privacy and cybersecurity theory and algorithms, technology, policy, and economics. Mitchell and Boneh address the progress in the theory and algorithms of security, including issues of system and property modeling and verification as well as cryptographic algorithms and protocols. They do not, however, address technology for producing software free of exploitable flaws. Bellovin discusses the history and prospects for the software and hardware technology for building systems that can provide, with some degree of assurance, the security and privacy properties desired of them, but does not hold out hope that, even a century hence, our software will be free of bugs or that attackers will go away. Landau reviews some of the past and present oscillations in policy and what sorts of policies the future may hold. Finally, Lesk addresses the economics of privacy and cybersecurity and how events in the coming decades may play out to influence the deployment of security and privacy enforcing technologies.

Now suppose that the average number of exploitable bugs delivered with each new product is substantially reduced over the next century, through a combination of advances in algorithms, theory, technology, and incentives from policy and economics. Such a trend may simply result in attackers striving to introduce backdoors or hidden, exploitable functions into systems as they are developed, distributed, or maintained (so-called "supply chain" attacks) or, alternatively, by tricking users into invoking legitimate operations that are against their own interests via "social engineering." We conclude that when the Proceedings of the IEEE celebrates its second century of publication, a section like this one will still be of interest.

II. THEORY AND ALGORITHMS (DAN BONEH AND JOHN C. MITCHELL)

A. Background

In the past half-century, Shannon's work on information theory and the emergence of computational complexity theory have made it possible to rigorously state and prove security properties of certain cryptographic constructions. We predict that in the next 50 years, the same underlying methodology will be used to bring the same degree of scientific rigor to other areas of computer security, replacing trial and error as a way of designing secure systems, with myriad benefits. Security modeling, generalizing the way that cryptography uses precise threat models and security conditions, will be used to capture an increasing range of security mechanisms, producing a science of security encompassing central aspects of cryptography, network security, access control, software system security, hardware security, and other branches of the field. While individual topics will retain their distinctive concepts and techniques, a unifying view will support broader use of rigorous security definitions, automated and handwritten security proofs, and comparative metrics for evaluating alternative security mechanisms.

To put this all in context, let us review some basic concepts. Computer systems are usually installed and operated to serve some set of intended users. However, many systems can also be misused, for a variety of reasons. Computer security involves designing ways to prevent misuse, such as stealing information from other users or interfering with their access to the system. We can think of a computer system and its security mechanisms, together with its intended users, as the overall system we would like to protect. We wish to protect it from malicious actions by an adversary, representing one or more forms of misuse. A set of actions an adversary may perform, together with the form of access the adversary has to the system, is called a threat model.

The theory of computer security therefore involves system models and threat models. For example, cryptography and network security involve a system model that transmits messages and an adversary who may intercept or modify these messages. A security definition combines system and threat models with a security condition. For example, we might wish to guarantee that a network adversary does not learn partial information about the content of any message.

B. Making Security Precise

It is important to realize that from a scientific standpoint, there is no absolute notion of security. Instead, meaningful statements about the security of any system must involve some characterization of the system, a precise threat model, and the conditions we wish to guarantee. This is well understood in the field of modern cryptography, but not always sufficiently well appreciated in other areas. Therefore, we consider cryptography a guiding example for the future development of a broader security science. To summarize, we must be clear about three things.

1) *System model*. There must be a clear definition of the system of interest, in sufficient detail to understand how the system behaves when subjected to both its intended operating conditions and unintended operating conditions. System models may be formulated in various ways, such as a mathematical formula for an encryption function, source code for a software system, or a formal specification of the possible states and state transitions of system.
2) *Threat model*. We must clearly identify the attacker's computational resources and the attacker's access to the system. For example, a cryptographic adversary may have access to

encrypted messages and a certain amount of computation time to devote to cryptanalysis. An operating system adversary may be able to place malicious code in a user process but unable to modify the operating system kernel. A web adversary may be allowed to build malicious websites, but assumed not to eavesdrop on the network.

3) *Security properties*. Security goals may be formulated as properties that we hope to prevent the attacker from violating. Such properties must be unambiguously associated with the system model so that it is clear whether a security property holds or fails when the system is operated by users and adversaries.

Given a system model, threat model, and security properties, we can determine whether the system is secure according to these conditions. Specifically, we achieve security if the system design guarantees the desired security properties, for all possible actions by any adversary operating according to the threat model.

Many commonly formulated security properties can be categorized as 1) confidentiality—no sensitive information is revealed; 2) integrity—the attacker cannot destroy the meaningfully operable condition of the system; or 3) availability—the attacker cannot render the system inoperable to intended users. However, there is currently no theory about why these properties are considered security properties. In addition, there is no standard way to decompose a given property into confidentiality, integrity, and availability components.

C. The Power of Security Policy Modeling

The reason we predict broader use of security modeling as a unifying approach is that formally stating our security goals lets us verify that a proposed implementation achieves the desired purpose. In addition, clear goals let us compare alternatives and look for the best possible solution. While it may seem obvious, this paradigm is not always followed; historically this failure has led to nonoptimal and sometimes insecure constructions. A good illustrative example is the story of *authenticated encryption*, a form of encryption that ensures both confidentiality and integrity. The security problem is to provide security against an active network adversary that, in addition to eavesdropping, can also tamper with messages en-route.

1) Early Days: The story begins with the development of the two relevant primitives. The first, called a message authentication code (MAC), is a data integrity mechanism that provides integrity, but no confidentiality. The second, called chosen plaintext secure encryption, or CPA-secure encryption for short, provides confidentiality against eavesdropping, but is not secure against an active attacker who tampers with traffic. Intuitively, combining the two primitives should provide both confidentiality and integrity against an active adversary, but how exactly should the two be combined?

Without a scientific understanding of how to combine the two primitives, every project invented its own method, hoping that a particular combination did the "right" thing. We give three examples from widely deployed systems: Transport Layer Security (TLS), the Internet Protocol Security (IPsec) protocol, and Secure Shell (SSH). TLS is a security protocol used to protect web traffic. IPsec is a protocol used to create private tunnels over the public Internet (for example, for connecting a branch office to a main office). SSH is a protocol used for securely logging in to a remote machine. In modeling all three widely used systems, we let k_e and k_m denote the encryption and MAC keys, respectively. Abstractly, these three protocols encrypt a message m as follows (here $x\|y$ denotes the concatenation of strings x and y).

- TLS: First, compute a checksum over the message, append it to the message, and encrypt the result. In symbols, compute

$$t := MAC(k_m, m) \text{ and output } c := E(k_e, m\|t).$$

- IPsec: First, encrypt the message and then output the resulting ciphertext followed by a checksum computed over the ciphertext. In symbols, compute

$$c_0 := E(k_e, m) \text{ and output } c := c_0\|MAC(k_m, c_0).$$

- SSH: Send the concatenation of the separately computed encryption and checksum. In symbols, compute and output

$$c := E(k_e, m)\|MAC(k_m, m).$$

During decryption, if the relevant integrity tag fails to verify, the decryption algorithm outputs a distinguished symbol (e.g., $\perp$) to indicate error.

Clearly the three approaches used in TLS, IPsec, and SSH combine integrity and encryption in three very different ways—the first encrypts the MAC, the second applies MAC to the encryption, and the third uses independent MAC and encryption. But which is right? Are they all secure? Without a precise characterization of the desired form of security, we cannot compare these constructions or even evaluate whether they achieve our goals.

2) Definitions: Long after these protocols were deployed, a number of papers [1], [2] defined the concept of authenticated encryption. In the threat model associated with authenticated encryption, the attacker is able to obtain the

encryption of arbitrary messages of its choice and the attacker's goal is one of the following two: learn information about the decryption of a well-formed challenge ciphertext (thereby defeating confidentiality), or generate a new well-formed ciphertext different from all ciphertexts previously given to the attacker (thereby defeating integrity). If the attacker cannot do either then we say that the system provides authenticated encryption.

Armed with a precise definition we can now compare the three constructions above.

- The IPsec construction can be shown to provide authenticated encryption for any MAC and CPA-secure encryption. The basic reason is that the MAC "locks" the ciphertext so that any modification of the ciphertext en-route will be detected by the decryptor.
- The TLS construction is not generically secure: there are specific examples of encryption and MAC such that the TLS combination does not provide authenticated encryption [3]. However, for specific encryption systems, such as randomized counter mode encryption, the TLS method provides authenticated encryption even if the MAC is only weakly secure (so called, one-time secure). The reason is that the MAC is protected by the encryption and therefore need not be a fully secure MAC; weak MAC security is sufficient.
- The SSH construction is known to be secure when a very specific MAC is used, but may not be secure for a general purpose MAC. To see why, recall that a MAC need not preserve confidentiality and therefore $MAC(k_m, m)$ may leak information about the encrypted plaintext.

Based on these comparisons, a designer can choose the appropriate method for the application at hand. When countermode encryption is used, the TLS construction is adequate even if a simple MAC is used. Otherwise, one should use the IPsec construction. This clear understanding is only made possible thanks to the precise formulation of authenticated encryption.

Using the definition of authenticated encryption, the National Institute of Standards and Technology (NIST) was able to publish precise encryption modes, called CCM and GCM, designed to meet the definition.

3) The Best Constructions: The last chapter in the story is that once the goals of authenticated encryption were clearly spelled out, it turned out that authenticated encryption can be built far more efficiently than by combining encryption and MAC algorithms. The reason is that encryption and MAC systems are often built from an underlying primitive called a block cipher [for example, the Advanced Encryption Standard (AES) is a widely used block cipher]. Building authenticated encryption by combining an encryption and MAC means that every block of the message is processed twice by AES: once for encryption and once for computing the MAC. By using the block cipher directly, it is possible to construct authenticated encryption by processing every block of the message only once [4]. Initially this seemed counterintuitive: without a precise definition it seemed that defeating an active adversary must rely on a combination of encryption and integrity. The precise security definition enabled cryptographers to prove theorems showing that the goals can be achieved more efficiently using a block cipher directly. Had this been known earlier, protocols like TLS and IPsec could have been twice as fast without affecting security.

D. Looking Forward

1) Composition: One of the most vexing basic problems in computer security is the problem of secure composition. While almost all interesting contemporary systems are built up from smaller components, it is accepted folklore that security properties do not compose. Even if each component is secure in isolation, a system composed of secure components may not meet its security requirements. Attacks using properties of one component to subvert another have shown up in practice in many different settings, including network protocols and infrastructure, web browsers and infrastructure, and application and systems software and hardware. We can divide the composition problem into two separate forms: nondestructive composition and additive composition.

Nondestructive composition is the problem of ensuring that if two system components are combined, then neither degrades the security properties of the other. This is particularly complicated when system components share state. For example, if an alternative mode of operation is added to a network protocol, then some party may initiate a session in one mode and simultaneously respond to another session in another mode, using the same public key (shared state) in both. Unless the modes are designed not to interfere, there may be an attack on the multimode protocol that would not arise if only one mode were possible. As another example, new attacks became possible when trusted computing systems were augmented with a new hardware instruction that could operate on protected registers (shared state) previously accessible only through a prescribed protocol.

Additive composition supports a combination of system components in a way that accumulates security properties. A basic example is the problem of authenticated encryption described above. Given a CPA-secure encryption function providing confidentiality and a MAC providing integrity, we would like to compose these two to produce authenticated encryption, which has both integrity and confidentiality properties. As we saw with the examples of TLS, IPsec, and SSH, there are several ways to combine these two parts that look reasonable, but only one of them is secure on general grounds.

We predict that, in the next decade, there will be significant progress on both additive and nondestructive composition. If we want a system with the positive security features of two components A and B, we need nondestructive composition conditions to be sure that we do not lose security features we want, and additive composition conditions to make sure we get the advantages of A and B combined. Since this is such an important core problem in computer security, we predict that secure composition will receive the increasing attention that it deserves.

2) Overcoming Limitations of Security Modeling: By formulating a security model with a precise threat model, it is possible to prove that attackers cannot mount successful attacks. However, in reality, attackers may be successful by using more powerful attacks that are outside the threat model. Two good examples are side-channel and fault attacks. In a side-channel attack, the adversary obtains more information than the model allows by measuring the running time or power used by the running system. This may reveal information about the secret key used by the decryptor enabling the attacker to break the system [5], [6]. In a fault attack, the attacker causes the decryptor to malfunction and observes how the decryptor fails [7]. Again, this may reveal information not covered by the model.

Attacks outside the threat model, however, are not a failure of the security modeling paradigm. Once side channels and faults are recognized as legitimate threats, they can be addressed within the paradigm. The challenge is to identify side channels appropriately, extend the threat model by modeling side channels abstractly, and then design systems that remain secure under this threat model [8], [9]. In extending security modeling beyond cryptography, there may be good reasons to consider several different threat models. In cybersecurity, for example, some mechanisms are designed only to protect against an adversary that operates malicious websites, but has no control over the network. We would like security against such attackers and network attackers, of course, but in practice the overhead of additional network protection may make simpler mechanisms more commonly used.

3) Future Computing Paradigms—Quantum Computing: Another paradigm change may come from future success in building quantum computers. While at this time it is not clear that sufficiently powerful quantum computers will ever be built, success in this area would have a direct impact on deployed cryptographic systems. First, traditional public-key systems such as RSA, ElGamal, and Elliptic Curves used for key exchange in protocols like TLS would become insecure, no matter what key size is used. Fortunately, in recent years, researchers have come up with new public-key systems that remain secure against adversaries that possess a quantum computer. Currently, the best candidates are public-key systems based on hard problems on lattices [10]. We emphasize that these are classical public-key systems that operate on standard classical computers, and as far as we know, cannot be broken by a quantum computer.

The second impact of quantum computers is a generic exhaustive search attack on all symmetric key ciphers that lets one find a k-bit secret key in time $2^{k/2}$. On a classical computer exhaustive search takes time 2^k, but a quantum computer achieves a generic square-root speedup. For example, a 128-bit AES key can be found in time 2^{64}, which today is considered insecure. To defend against these attacks, designers will need to double the key length of all symmetric key systems. Concretely, this means moving to symmetric ciphers where the best classical exhaustive search attack takes time 2^{256}, so that the quantum attack will take time 2^{128}. Of course, there may be better quantum attacks on the cipher and there will likely be a need to design new classical symmetric ciphers that can resist clever quantum attacks. This is a fascinating area for future research.

III. TECHNOLOGY (STEVEN M. BELLOVIN)

A. The Past

There is an oft-misunderstood quote from Santayana on the fate of those who ignore history. The full paragraph is more interesting [11]:

> Progress, far from consisting in change, depends on retentiveness. When change is absolute there remains no being to improve and no direction is set for possible improvement: and when experience is not retained, as among savages, infancy is perpetual. Those who cannot remember the past are condemned to repeat it. In the first stage of life the mind is frivolous and easily distracted; it misses progress by failing in consecutiveness and persistence. This is the condition of children and barbarians, in whom instinct has learned nothing from experience. In a second stage men are docile to events, plastic to new habits and suggestions, yet able to graft them on original instincts, which they thus bring to fuller satisfaction. This is the plane of manhood and true progress. Last comes a stage when retentiveness is exhausted and all that happens is at once forgotten; a vain, because unpractical, repetition of the past takes the place of plasticity and fertile readaptation. In a moving world readaptation is the price of longevity. The hard shell, far from protecting the vital principle, condemns it to die down slowly and be gradually chilled; immortality in such a case must have been secured earlier, by giving birth to a generation plastic to the contemporary world and able to retain its lessons. Thus old age is as forgetful as youth, and more incorrigible; it displays the same

> inattentiveness to conditions; its memory becomes self-repeating and degenerates into an instinctive reaction, like a bird's chirp.

In other words, we need to know the deficiencies of the past in order to improve on them, but if we ever relax and get complacent, we will start making the same mistakes again.

The history of cryptology, a crucial security and privacy technology, is no exception. Reading Kahn's magisterial work [12], [13], we see several trends that have persisted over the centuries.

1) *Security*. Historically, the offense has generally had the upper hand. From the Black Chambers of Renaissance Europe to Bletchley Park, Arlington Hall, and beyond, cryptanalysts have been able to read messages in systems that were thought to be very strong. This is *a priori* quite reasonable: why strengthen a system that has not proved susceptible?
2) *Speed*. Cryptosystems (and cryptanalysts) have continually had to adapt to communications speeds. Machine encryption was developed to deal with the speed of the telegraph; similarly, the Black Chamber of 18th century Vienna had a schedule dictated by the cycles of the Post Office [12, p. 261].
3) *Usability*. Over the years, usability has been a major driver in improvement, originally to aid in encryption speed and accuracy (and hence reduce the need to send some material in plaintext), and later to help avoid the mistakes that would aid enemy cryptanalysts.
4) *Amateurs*. Many of the major advances in cryptology have come from people outside the cryptologic mainstream, including Jefferson's wheel cipher, public-key cryptography, and the one-time pad [14], [15]. Often, though, these discoveries were not appreciated until after later reinvention.

When trying to predict the future, it is reasonable to use these four pillars for our analysis. That said, there seems to be little on which to base predictions about cryptography. Encryption speed, at least in dedicated hardware, is generally adequate today; increasing parallelism should let speeds continue to increase as needed. Similarly, though no one has ever deployed a system they knew to be weak, today's systems have proven remarkably strong; we do have a much better mathematical understanding of the problem today. In the 35 years that the Data Encryption Standard has been with us, only one attack significantly stronger than brute force has been found in the open literature [16]; the other known weakness, a brute force attack, was realized from the beginning and was arguably deliberate [17], [18]. It is impossible to foresee when a brilliant amateur will come up with a fundamentally new cryptologic idea. While the advance of the field makes it unlikely that such a person will conceive of a dramatic new improvement in today's technologies, someone who asks a very different sort of question—public-key cryptography is the obvious parallel—may indeed change the field. But genius occurs when it occurs; it is not predictable.

There is less to say about the history of other security technologies, simply because they are so much younger. Fundamentally, all consist of a "trusted computing base" (TCB) (e.g., the operating system) that mediates access to protected resources, and a policy model that says who can have access to what.[1] By and large, this approach has failed for four reasons.

1) The policies were not compatible with the kind of work people needed to do in the real world.
2) The trusted software was buggy, allowing for policy violations.
3) Many security violations now take place outside the TCB. For example, document viewers are never considered trusted, but many a system has been compromised when a document viewer attempted to render a booby-trapped file.
4) Security systems are not usable; neither software developers, system administrators, nor end users can adequately carry out their intentions, even when the underlying system is sufficient.

Again, these four points should underlie our prognostication. It is worth noting that usability is the only point in common.

One more aspect of security and privacy should be touched upon: attacks are perpetrated by people. Any solution depends on understanding both what must be protected, and against whom. The latter may be random hackers, common criminals, national intelligence services, or corrupt insiders working for any or all of the above.

B. Plus Ten Years

In fields as mature as security and cryptology, radical changes are rare. It took NIST four years to select the AES [19]; selecting a new hash function will take five [20]. New operating systems take longer still; the design/code/test cycle, the hardware replacement interval, and the need for backwards compatibility all suggest that there will not be many radical changes in the next decade. Rather, the changes that we see will be driven more by technology (e.g., the rise of tablets) and the operational environment, rather than by new scientific or engineering insights. There is the chance of new cryptanalytic attacks on AES (given results such as [20], few would be shocked by a major success); if so, the attack is likely to be quite expensive and not a major

[1]The actual definition of the TCB is more complex than we can go into here, and always includes the type of security model it enforces. As such, the following discussion uses mildly incorrect terminology; the essential message, however, is correct.

threat to most users, but instead to be grounds for picking a newer standard.

There is more room for usability improvements in the near term. Both security and cryptology technologies have a long history of user interface problems [13], [22]–[26]; these issues are finally getting significant attention in the mainstream research and development communities. The most likely significant change would be the advent of "encrypted by default"—encryption is always used, with little or no user input needed. A similar trend in privacy technologies seems unlikely, given the (commercially driven?) drive for people to share more of their lives.

By this time, we should start to see a better division between the TCB and untrusted code. Rather than a simple binary choice or even a hierarchy, we will start to see more complex graphs of the trust relationships between different components. Trying to manage these graphs will not be easy; the increased complexity will likely more than balance the theoretically improved security.

There is no reason whatsoever to think that either dishonesty or buggy software will be gone by then.

C. Plus Twenty Years

The last ten years have made two technological trends increasingly clear: more and more devices are being controlled by microprocessors, including objects as mundane as coffee makers; in addition, more and more of these controllers are being networked. In 20 years, **everything** will be connected. Even today, there are network-connected light bulbs for sale. The security implications of this are frightening.

However, 20 years is long enough for fundamental new advances in system design to enter into commercial use. A number of new technologies that address today's limitations have been proposed, such as content-based access control [27]; perhaps these will help. Whether this will cope with the complexity problem is quite unclear, since there will be exponentially more user/device pairings to protect. In addition, many of these new devices will leak private information; while this can be beneficial, too often we see things deployed without adequate consideration of the question. It may be that we will see a societal shift away from privacy [28]. We are even told today that "all issues of privacy are old people issues" [29].

It is all but certain that there will still be dishonest people, and our software will still have bugs.

D. Plus Fifty Years

Fifty years hence is an interesting time scale in which to work, since one can easily look back at the history of the last 50 years of computing. Not only does that interval cover most of the commercial history of computers, most of the basics we rely on today were either present or about to appear, including remote access, large (by the standards of the day) mass storage devices, and video games on screens [30]. Packet switching [31] and protected mode (and hence the TCB) for operating systems had been invented [32]. Handheld mobile phones had been described, at least in science fiction [33]–[35]. What is new are the combinations and the unimaginable complexity of today's systems. It is plausible, then, that many of the basic concepts of the systems of a half-century from now exist today—but of course we do not know which they are or how they will be combined.

One intriguing technology area to watch is the use of quantum phenomena. There are already encryptors based on the physical impossibility of undetectably measuring certain properties (such as polarization) of a photon while simultaneously not affecting it [36]. There have been a number of critiques of the concept on technical grounds, but one of the most vexing problems is that such secure communications are inherently point to point, since to a quantum state a legitimate router is indistinguishable from an eavesdropper. Will scientists and engineers find a way around this problem, perhaps by large-scale optical switching to provide direct, end-to-end paths for actual photons? Will there be some sort of quantum routing device, embedded in a tamper-proof chip, that will preserve the right properties while excluding eavesdroppers?

A related notion is quantum computation. As noted, it is unknown at this time whether large, reliable quantum computers can be built, especially the massively parallel ones needed to attack symmetric ciphers via exhaustive search. Within 50 years, we should know, and—if they are indeed feasible—be able to build them and use them to attack real problems, with all the implications that we will have for cryptographic security [37].

One can hope that by 50 years from now, the complexity problem will have been tamed; alternatively, it is quite possible that complexity will continue to increase faster than our ability to cope with it.

We do feel that there will still be dishonest people, and our software will still have some bugs.

E. Plus Hundred Years

In order to predict what security and privacy technologies will be like 100 years hence, it helps to envision the dialog one might have via a "chronophone" with Herman Hollerith, the inventor of the punch card, around 1910 or so. The most important basic concepts existed: there were telephones, radios, the ancestor of fax machines, voice recording and playback, cryptography, and of course the high-speed data processing devices that he invented. It would still be extremely challenging to explain things like Facebook, operating system access controls, Google, etc., let alone their challenges and failures. One can just imagine the response: "you're telling me I should 'drag and drop' a 'file'—do you mean a metal tool or a bundle of papers from my desk, and by the way do you know what happens if someone drops a deck of my punch cards?—to an encrypted 'flash drive' (I can drive a car, but what does

that have to do with a camera flash?); alternatively, I can wander down to the nearest Orthodox church and hope that the priest doesn't get angry with me for clicking at—'on'?—his icons, all in order to protect something from a worm? A worm? One of those things I stick on my hooks when I go fishing? And why do you spell 'fishing' with a 'ph' and imply that it's a bad thing?"

Privacy and authentication are not new issues. The use of a mother's maiden name as an authenticator goes back to at least 1882 [15] and though that particular scheme may be challenged by social changes [38], some form of secret-based authentication will likely persist.

Usability concerns may be eased by the advent of very competent user agents. They may not implement the proverbial "Do What I Mean" (DWIM) instruction; they will, however, know individual's preferences well enough to make sophisticated judgments. Similarly, system file protections will be set automatically based on very high-level statements of how a system should act. [On the other hand, this may be seen as a form of artificial intelligence (AI). We note that AI has been just around the corner for about the last 50 years. . .]

Cryptography is more problematic. It may not be needed (is it even possible to intercept modulated dark energy beams?). A more interesting question is the mathematical underpinning. Will it be possible to prove that some practical systems are unconditionally secure? Alternatively, will it be possible to prove that, say, public-key cryptography is not possible? If the latter, will the proof be of a form that leads to a real attack, or simply demonstrates that one must exist?

Finally, there are two predictions we can be quite certain about: there will still be dishonest people, and our software will still have some bugs.

IV. POLICY: NEVER QUITE SURE WHAT IT IS GOING TO BE FROM ONE MINUTE TO THE NEXT (SUSAN LANDAU)

Fifty years ago phones stayed put, computers were the size of a room, and a call to Europe from the United States required booking in advance. Then, in 1970, one could dial the United Kingdom directly from the United States. Shortly afterwards, the same could be done for the European continent, and then the rest of the world. Now, calling someone no longer requires knowing where they are. But if you do not know where the person you are calling is, others do. And they see ways to use it that you never intended.

In 2011, the Chinese government announced that it would track people's movements through their cell phones for better traffic control [39], while a recent study of the Haitian population after the 2010 earthquake showed that similar tracking is extremely useful in informing where people are—and where relief aid should go [40]. In some cases, shutting off a phone does not prevent one from being tracked: in France, in a location of interest—say where charges are being made on stolen credit cards—police note the appearance of shutoff cell phones in the vicinity [41].

The combination of the Internet, social networking sites, and data aggregators have enabled a perfect storm of capability for identifying individuals as they go about their daily lives: who you are [42], where you live and work [43], even whether you are catching the flu [44]. More than a half century ago, Isaac Asimov [45] premised that predicting the flow of human civilization is possible, but the author disavowed predicting actions of individuals. With the data we now have, the latter seems almost within reach. People's ubiquitous use of communications devices reveals their daily habits, their use of social networking sites makes public present activities and future plans, and sensors everywhere—in buildings, cars, even the great outdoors—shows what they are doing, where they are doing it, and with whom [46]. Meanwhile information about their genome reveals information not only about themselves, but also their relations: parents, siblings, even unborn offspring.

Technology giveth, but technology also taketh away. Encrypted and peer-to-peer communications tools thwart wiretapping. Tor, an anonymizing overlay network that hides transactional information complicates investigations [46]. Anonymized SIM cards mean you cannot track when or to whom a target is calling. Criminals—and hackers who work in support of nation states—use the lack of IP packet-level attribution to hide who and where they actually are.

For law enforcement, the world has become very complex. There are multiple legal requirements on providers—retain communications transactional data, do not retain search data—and multiple ways for the determined user to secure her information. There are multiple computer operating systems, multiple smartphone operating systems, multiple rich sources of information, complicating the investigatory landscape.

But while digital forensics may have become more complex to untangle, the change in the way people live means that there are ready electronic trails everywhere for these investigators to mine. The only way to really secure anonymity is to disconnect. In fact, that alone may cause one to stand out. A prime example of this is Osama bin Laden, whose expensive compound in Abbottabad, Pakistan, was notable for its lack of telephone and Internet connectivity.

We are at a time when individuals' privacy resides in an Alice-in-Wonderland state [48]. One moment Alice imbibes from the bottle labeled "DRINK ME" [48, p. 31] and she is in the world of cell phones and location tracking with a huge electronic profile. Her privacy almost disappears. Then Alice bites from a small cake [48, p. 63] of end-to-end encryption, Tor, and anonymizing cell phone cards, and her electronic profile shrinks. Her privacy becomes ENORMOUS. A fan appears. Hot and lonely, Alice

waves her Facebook fan. Before she knows it, Alice—and 750 million others around the world—are sharing information on what they are doing for the weekend, what movie they are watching on Netflix (and whom they are watching it with)—and hundreds of thousands of other intimacies. Personal privacy has become as minute as the five-inch girl that Alice has become [48, pp. 37–39].

Police pull one way, other government agencies another. Europe, Canada, Australia, and New Zealand have long had activist government efforts that control both government and private enterprise use of individuals' data. The United States has become more activist as well: the U.S. Federal Trade Commission (FTC) has assumed a role as activist privacy regulator [49], while the government's effort for secure online identity management, the National Strategy for Trusted Identities in Cyberspace (NSTIC) [50], has explicitly pressed for privacy protections. And even while the FBI was pressing for stronger surveillance capabilities, the U.S. State Department was simultaneously lauding communication tools enabling secure communication by human rights activists. The United States is not the only place in which this occurs. At almost the same time that the European Union passed a data retention requirement for telecommunications providers, a European Union advisory group challenged the length of time that Google retained user search data [51]. Or as Alice puts it, "I'm never quite sure what I'm going to be, from one minute to the next!" [48, p. 77].

In the battle over privacy, one thing is sure: governments will follow self-interest. Consider the following example. In the late 1990s, European governments were concerned that the U.S. government was eavesdropping not only on diplomatic and military communications, but also on civilian ones, including those of industry. Having been on a trajectory to limit the use of *strong cryptography* (cryptography that cannot be broken using current technology), the European Union lifted export controls on systems containing strong cryptographic systems. Meanwhile, having spent from 1975 seeking to control civilian use of cryptography, the U.S. position made its own U-turn. The rationale was: partially competition (the government preferred that systems sold containing strong cryptography be domestic in origin), partially pressure from Congress, and partially a realization that it was time to focus elsewhere [52].

What really happened was a divergence of interests between the national security community, which supported the loosening of export controls, and law enforcement, which did not, enabling this shift. National security's dual role of securing information as well as conducting surveillance meant that by the late 1990s it saw the world differently, and that difference meant it supported the wider use of strong encryption. That split has implications for the future.

Now predicting anything as capricious as the wild swings of privacy and government policy is hard, but certain drivers of government policy on privacy are clear.

1) Everyone—businesses, governments, individuals—will be collecting and storing data. People will reveal and devices will store information that once was essentially inaccessible or entirely ephemeral: the number of steps they walk and foods they eat each day, the strand of music they listen to while waiting for the subway, the lanes they traverse as they drive to work.
2) Data collected on individuals will continue to be an extraordinarily rich source of information. Its value to private enterprise means that such collection will continue.
3) Capabilities once only available to governments—tools for tracking and tools for hiding aspects of data—will become increasingly available to private industry and even private citizens.
4) In many cases, use of transactional data combined with other rich data sources will obviate the need for content.
5) Remote access will confuse the ability to distinguish inside and outside of an organization (or a nation). Although the distinction is real, this blurring will have large implications for cybersecurity—and for privacy.

In the absence of regulation by government, data collection will proliferate madly. Sophisticated users will have tools to hide their tracks—note that the use of complex encryption or anonymizing tools will mark a user as a "party of interest"—while the majority of the populace will have less privacy unless government intervenes. Because of the ready availability of open source data (information from publicly available sources), protection of individuals' privacy will become a national security issue.

Democracies and authoritarian nations naturally have different self-interests when it comes to privacy. In democracies, the need to protect large swaths of industry as well as many who work in sensitive but unclassified parts of society—members of the judiciary, family members of the law enforcement and the military, etc.—may create an argument for the government to step in, as it has done in identity management and activity tracking, to protect privacy. Widespread revelations of personal information may not really be much to society's value. Authoritarian regimes function by controlling their population and are less likely to protect user privacy.[2]

Many security breaches—and security breaches are often indistinguishable from privacy breaches—cross borders, making privacy an international issue. But international investigations, whether of criminal activity, spying, or hacking, are complicated, and in a world where some nations assiduously protect privacy and others do not,

[2]While Israel is a democratic state, its use of Facebook to identify pro-Palestinian activists and thus prevent them from flying to Tel Aviv is one such example [53].

these investigations are often stymied. Many nations have signed the Council of Europe cybercrime treaty, which assures mutual law-enforcement aid in such investigations, but notably Russia and China have not.[3] Given the wide disparity of views on whether privacy is worth protecting, global, or near-global, treaties enabling better privacy protections are unlikely. It is much more likely that instead there will be international standards, e.g., such as those emerging on identity management, and cross-national agreements, such as those promulgated by the European Union.

Government action will also depend on perceptions of security. A greater ability to measure effectiveness of solutions, something proposed both in National Research Council studies [54], and in conferences such as the Workshops in Economics of Information Security [55], may lead to better governmental privacy protections.

The real watchword, though, is volatility.

> "Are their heads off?," shouted the Queen.
>
> "Their heads are gone, if it please your Majesty," the soldiers shouted in reply [48, p. 104].

In Alice in Wonderland, the heads were still right where they belonged. In our world though, we will have to keep running as fast as we can to stay in place if we are to have any privacy at all.

V. ECONOMICS (MICHAEL E. LESK)

A. Background

There is a traditional Hollywood summary of an actor's career.

1) Who is Joe Blow?
2) Get me Joe Blow.
3) Get me a Joe Blow type.
4) Get me a young Joe Blow.
5) Who is Joe Blow?

Right now, most people are still wondering "what is cybersecurity?" As time goes on, we predict people will start to demand it; then want it cheaper, and then eventually forget it. This section will focus not on how the technological problems of security will be addressed as much as who will wind up paying for it. Our lack of any general agreement on how we measure either the cost or value of security makes economic discussions particularly fuzzy and unsatisfying. Nevertheless, with only slightly better cost models, motorists pay for safer cars, with government regulation forcing enhancements such as antilock brakes; but taxes pay for safer roads. What model makes sense for cybersecurity over a long time?

1) Should You Believe Anything I Write About the Next Hundred Years?: Either Yogi Berra, or Neils Bohr, or somebody, said something like "Prediction is difficult, especially about the future." My favorite forecast is an 1888 book by Edward Bellamy entitled *Looking Backward* [56], which is a description of Boston as it will be in the year 2000. He imagined little technological progress, with everything still powered by steam or compressed air; but he foresaw great social progress, with crime and poverty eliminated. Thus, he had things pretty much backwards. This is often what happens with predictions. The IEEE has decided it wants a view of the future economics of privacy and cybersecurity, and I am replying perhaps in sympathy with Horatio Hornblower, who said "I'd rather be in trouble for having done something than for not having done anything." And the IEEE has a history of publishing technology forecasts, such as its enthusiasm in 1999 for "wearable computers" and a 340-MB floppy drive [57], not to mention its explanation that year that energy price spikes were caused by bad weather [58]; Enron was mentioned only as an investor in photovoltaics. So, with your understanding that this may be even less reliable than the usual articles in this magazine, here is a guess at cybersecurity and economics.

In discussing cybersecurity, there are two distinct problems: accidental losses and malicious behavior. Perhaps the most familiar noncriminal cybersecurity problem is "a disk crash ate my homework." When there is nothing malicious about the problem, it is much easier to design a solution. General robustness should diminish as a problem. Technology is getting more reliable, and more importantly, cloud storage will provide the redundancy and general service. I am writing this on Google Docs, precisely to allow general access from many locations, and as a side effect freeing me from worry about backups. The primary issue with cloud storage is not robustness but privacy, which is some other author's problem.

The economics of the "cloud" seem so attractive that even within ten years, I would expect that "lost files" will be a relatively solved problem at a relatively low cost. Some cloud services are so cheap that they are given away, or provided free in exchange for advertising. Just as nuclear power was going to make electricity "too cheap to meter," gigabit bandwidths have made e-mail almost free (although this has not stopped text messaging companies from charging enormous amounts per byte for a short message).

Even without the cloud, the ease of duplication has meant robust data storage. Surprisingly few important files have been lost. The 1960 census tapes, in the end, were mostly OK; any losses were small and unimportant [59]. Perhaps the most important example is that the National Aeronautics and Space Administration (NASA) lost the higher resolution video recording of the moon landing

[3]The United States has signed the treaty but not the separate section dealing with hate speech; such speech is protected under the First Amendment.

[60]. As more and more people get their computing from cloud resources with professional management, lost files should cease to be a problem. Yes, we will see the cloud companies outsourcing their operations to countries with inadequate infrastructure and poorly paid staff, but the steadily lower costs of computing will allow them to maintain robustness through redundancy.

Criminal behavior will be more difficult to deal with. Cybercrime is of course a new activity. In the same way, during the 19th century we saw the new crime of train robbery, and it was eliminated through better policing. Car theft arose in the 20th century but has now declined as a result of better car locks. Similarly, I would expect that cybersecurity will improve over time. What I do not know is how. Unlike Bellamy, I do not think we will eliminate crime in general. Even if we were able to reduce some kinds of crime within one society, as with the general reduction in murder in the United States between the 1920s and the 1950s, and then, after a resurgence, another reduction during the 1990s and the 2000s, cybercrime is international. Much of the cyberfraud observed in developed countries is associated with foreign origins. Ask anyone what they think of when you say both "computer" and "Nigeria."

At the moment, cybercrime being new, we have not quite faced up to the need to do something about it, and we spend relatively little money on it. For example, a U.K. report suggested that the annual loss to cybercrime was £27 000 million, but that expenditures on coping with cybercrime were £650 million [61]. My expectation is that as we take cybercrime more seriously, we will need to substantially increase what we spend preventing it. But when will that happen?

Based on past history, what we should expect is that the problem will get worse until we decide to spend the resources to deal with it. What I do not know is whether we are going to need an incident such as Pearl Harbor or 9/11, or whether we can begin to take steps in a reasonable way.

As comparisons, consider things like sports scandals. We live through regular episodes of discovering that matches are being fixed. There was the "Black Sox" baseball scandal of 1919, the college basketball scandal of 1950, and as I write in 2011 there are allegations that world soccer matches are being fixed. Each produces some kind of reaction, involving new governance and investigative procedures. In each case (at least until now) something is done to make the matches more honest. There are also doping scandals, but in several of those cases it did not appear that the authorities showed much initial interest in ending doping (see baseball or cycling).

So what is likely to happen? There will be some sufficiently distressing event to cause us to decide that we have to invest the effort needed to fix the problems. Again, look at history. We started air traffic control across the United States after a 1956 midair collision above the Grand Canyon [62]. Britain decided that railway signals should have a default position of red (stop) after an ice storm froze semaphore arms in the green position and caused an accident at Abbots Ripton in 1876 [63]. That particular kind of accident never happened in North America [64], and our signals remained default green (until it became common to turn them off to save on bulb lifetime).

So our optimistic scenario is that within the next ten years, some relatively unimportant hacking scenario causes us to insist on more robust software, perhaps a more dramatic version of the episode in which somebody hacked the PBS website but only chose to announce that Tupac Shakur is alive and living in New Zealand [65]. My pessimistic scenario is a collapse of something like our electricity infrastructure or our banking infrastructure. Unfortunately, I suspect that it will take something dramatic to raise public awareness: the wiretapping of several Greek cabinet ministers, for example, has passed almost without notice in the rest of the world, and even wiretapping by Murdoch newspapers has not produced many calls for improvements to the security of voicemail systems. I fear we will need a disaster that affects far more people.

Cybersecurity is an intricate mixture of actions taken by criminals, manufacturers, governments, and users. The economics of cybersecurity is about who should pay how much for what steps. Will cybersecurity economics be primarily money spent by the government on police efforts to put malefactors in jail? Or by manufacturers to build more robust software? Or by users to install virus checkers and password managers? Or by users in the form of losses to crime?

Initially, the response is going to be from users. After 9/11, lots of people stocked up on canned food and bottled water. Similarly, after the trigger incident, whatever it is, people will start buying anything sold as virus protection or bank account insurance. Next, government action will come, with legal programs to investigate and prosecute the criminals. I am assuming that whatever the trigger event is, it will be worldwide, and so it will be possible to get international cooperation in chasing down the perpetrators. It may be expensive, but I do not see anyone quibbling about the cost of killing Osama bin Laden. Finally, it will become private industry—vendors will harden their products, and users will pay a bit more for better and more robust service.

Stopping cybercrime should be easier than other law enforcement activities. Cybercrime is not as frightening as crimes of personal violence like robbery or kidnapping; nobody is afraid to leave their house at night because they might meet a spammer. It usually works through the financial system, meaning that its victims are sophisticated enough to have bank accounts and e-mail services. There is a problem with permanence; problems where you might in the past have said, "that was a mistake, but at least it is behind me" now live forever on the Internet. This offers new blackmail possibilities.

Cybercrime is also not something that attracts public sympathy. IMDB reports that 17 movies have been made about John Dillinger and 80 about Jesse James; 115 have been made about Robin Hood. Even more recently, there have been four movies about the "Great Train Robbery" which took place near Aylesbury, U.K., in 1963. By contrast, I do not know the name of any cyberfraud hero. And I cannot see any cybercriminal featuring in a video game, a museum, or a festival, all of which exist for the James brothers. There are innumerable movies about geeks who use their hacking skills to defeat criminals and/or save the world (while finding true love at the same time), but they are the sheriff, not the bad guys.

Assuming that somehow we do decide that we are going to pay for decent cybersecurity, who is likely to bear the cost? There are three possibilities: users, vendors, or the government. As always, the loudest voice is from the vendors, explaining why this is not their problem. For example, a *Network World* item from 2009 is headlined "User education key to IT security: Microsoft" [66]. Similarly, vendors of electronic medical records systems regularly attribute all problems to user errors, rather than bad product design. The right response is a line by Masys about aviation safety: "pilot error is not an explanation but rather is something to be explained" [67]. However, in practice, it has been relatively easy for software systems to place the blame on users, rather than themselves.

This is not a new discussion. In 1909, the railway engineer Wilson wrote "the Machine is now perfected, and these pages will, without doubt, force the reader to the conclusion that present day accidents are due to failure of the Man" [68]. Then, in 1925, the same author said "the record of 40 to 50 years ago, which now appears to us so terrible a tale of destruction, was due less to the sins of omission and commission of the men than to the neglect of the companies to provide safeguards that were available" [69]. He even had the chutzpah to write in the preface to the 1925 book "in innumerable instances blame was cast on signalmen that should have been laid at the door of the officers and directors of the company, who failed to appreciate the benefits that concentration, the block, and uniform signaling would have provided," not mentioning that he had been one of those casting the blame.

Given that we decide we have to do something about cybersecurity, who will wind up paying for it? We can imagine several scenarios.

1) Everybody is persuaded that they have to buy more robust systems or better virus checkers, and spend their own money on it. This would be in the same way, for example, that large numbers of people now sign up for intensively advertised credit reporting services. Vendors would have to be advertising that their systems were safe and there would need to be a way for purchasers to believe that the system was safe.
2) The government might impose either liability or standards on the vendors. As an example, it is a result of government standards that every new car has seat belts and air bags. As an alternative, products such as asbestos insulation, IUD contraceptives, and climbing gyms have been driven out not by direct regulation but by liability lawsuits.
3) Police agencies might get better funding to seek out cybercriminals, in the way that a rash of publicized bank robberies encouraged the creation of the modern FBI.

The payment mechanism is likely to change over time, as the threat level rises and falls. Today, for most people, we are at the "what is cybersecurity?" level. Pretty soon people will want it, then they will want it but cheaper, then it will become routine, and then it will be forgotten.

So right now, cybersecurity is a specialty for a few computer experts. We are in the "who is Joe Blow?" phase. What happens next?

B. Plus Ten Years

We are now in the stage of "Get me Joe Blow!" Some event causes society to decide that cybersecurity is essential. Unfortunately, it is probably going to have to take some form that affects a lot of people: a continent-wide blackout in the United States or Europe, a complete gridlock of road traffic when every traffic light stops functioning, or a freeze-up of all financial markets because no transaction can be believed in.

After this, we see some steps taken. The first reaction will be user panic, widespread demand for security, much of which will be wasted money, but that is typical. Then, some regulatory and police responses will come. These are dictated by government action, since vendors will continue to explain why whatever happened was somebody else's problem. Vendors are about as likely to take responsibility as one's cat is to explain why the curtain is now on the floor. What actions might be taken? These might be dedicated networks for critical infrastructure (unlikely; we do not have a separate set of bridges for critical road traffic), or perhaps a combination of serious testing of software, lack of tolerance for continued buggy shipments, and more tracking and detection of malefactors (with some loss in personal privacy, as with the airlines). Perhaps the average browser no longer allows arbitrary downloads, there is no more anonymous e-mail, and some kind of tax on Internet Service Providers (ISPs) is used to pay for the police efforts. The cost is some combination of product costs, taxes, and hassle and delay of the users.

C. Plus Twenty Years

This is the stage of "Get me a Joe Blow type!" where people want a less annoying solution to cybercrime. The situation has improved to the point where we are complaining about the problems it makes for daily life, in the same way that we complain about air travel security.

Fortunately, the increasing decline in system costs means that we do not complain about the price of cybersecurity. Major disasters no longer happen, and people think that avoiding them should be cheaper and easier. A few vendors and politicians continue to emphasize the threats and encourage people to spend more on cybersecurity products. A few nations try to capitalize on a reputation for honesty and reliability to encourage cloud computing vendors to set up on their territory, but they lose out to places that offer the cheapest cost and the least enforcement of tax laws (or any other laws).

Costs shift from the public to the industry, and in fact start to arrive at the vendors, as they are producing decent software.

D. Plus Fifty Years

This is the stage of "Get me a young Joe Blow," when the solutions have become routine. The system is now running fairly smoothly. To the extent that people can remember when you could send anonymous e-mail and there were no fingerprint or iris scanners, it has gone the way of other nostalgia.

E. Plus Hundred Years

Who is Joe Blow? Cybersecurity incidents are now rare and since the typical year goes by without anything on the news about it (when was the last time you heard about a U.S. train robbery?), nobody except a few specialists knows anything about cybersecurity. Perhaps the Smithsonian includes a "Museum of Spam" which proudly displays the "Green Card Lottery" information e-mail from 1994.

This has been a basically optimistic view. You may be more pessimistic. I can only say if I could really predict the future, I would not be writing articles for the IEEE. I would be at the racetrack. ■

Acknowledgment

The authors would like to thank G. Cybenko for his expeditious review and helpful comments which led to improvements in the manuscript. Susan Landau would like to thank Dame G. Beer for her talk "Alice in Time" at the Radcliffe Institute for inspiring these ideas. Carl Landwehr would like to thank the Proceedings of the IEEE staff for their patience.

References

[1] M. Bellare and P. Rogaway, "Encode-then-encipher encryption: How to exploit nonces or redundancy in plaintexts for efficient encryption," *Advances in Cryptology—ASIACRYPT'00*, vol. 1976, T. Okamoto, Ed. Berlin, Germany: Springer-Verlag, 2000, pp. 317–330.

[2] J. Katz and M. Yung, "Unforgeable encryption and adaptively secure modes of operation," *Fast Software Encryption, 7th International Workshop, FSE 2000*, vol. 1978, B. Schneier, Ed. Berlin, Germany: Springer-Verlag, 2000, pp. 284–299.

[3] H. Krawczyk, "The order of encryption and authentication for protecting communications (or: How secure is SSL?)," *Advances in Cryptology—CRYPTO 2001*, vol. 2139, Lecture Notes in Computer Science, J. Kilian, Ed. Berlin, Germany: Springer-Verlag, 2001, pp. 310–331.

[4] P. Rogaway, M. Bellare, J. Black, and T. Krovetz, "OCB: A block-cipher mode of operation for efficient authenticated encryption," in *Proc. ACM Conf. Comput. Commun. Security*, 2001, pp. 196–205.

[5] P. C. Kocher, "Timing Attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems," in *Proc. CRYPTO*, 1996, pp. 104–113.

[6] P. C. Kocher, J. Jaffe, and B. Jun, "Differential power analysis," in *Proc. CRYPTO*, 1999, pp. 388–397.

[7] D. Boneh, R. A. DeMillo, and R. J. Lipton, "On the importance of checking cryptographic protocols for faults," in *Proc. EUROCRYPT*, 1997, pp. 37–51 (Extended Abstract).

[8] A. Akavia, S. Goldwasser, and V. Vaikuntanathan, "Simultaneous hardcore bits and cryptography against memory attacks," in *Proc. 6th Theory Cryptogr. Conf.*, 2009, pp. 474–495.

[9] M. Naor and G. Segev, "Public-key cryptosystems resilient to key leakage," in *Proc. CRYPTO*, 2009, pp. 18–35.

[10] O. Regev, "On lattices, learning with errors, random linear codes, and cryptography," *J. ACM*, vol. 56, no. 6, 2009, DOI: 10.1145/1568318.1568324.

[11] G. Santayana, *The Life of Reason*. New York: C. Scribner's Sons, 1905. [Online]. Available: http://www.gutenberg.org/files/15000/15000-h/vol1.html

[12] D. Kahn, *The Codebreakers*. New York: Macmillan, 1967.

[13] D. Kahn, *Seizing the Enigma: The Race to Break the German U-Boat Codes, 1939–1943*. Boston, MA: Houghton-Mifflin, 1991.

[14] W. Diffie and M. E. Hellman, "New directions in cryptography," *IEEE Trans. Inf. Theory*, vol. IT-22, no. 6, pp. 644–654, Nov. 1976.

[15] S. M. Bellovin. (2011, Jul.). Frank Miller: Inventor of the one-time pad. *Cryptologia* [Online]. *35(3)*, pp. 203–222. Available: http://dx.doi.org/10.1080/01611194.2011.583711

[16] M. Matsui, "The first experimental cryptanalysis of the data encryption standard *Advances in Cryptology—CRYPTO'94*, vol. 839, Berlin, Germany: Springer-Verlag, 1994, pp. 1–11.

[17] J. Gilmore, Ed., *Cracking DES: Secrets of Encryption Research, Wiretap Politics & Chip Design*. New York: O'Reilly, Jul. 1998.

[18] W. Diffie and M. E. Hellman, "Exhaustive cryptanalysis of the NBS data encryption standard," *Computer*, vol. 10, no. 6, pp. 74–84, Jun. 1977.

[19] NIST, *Advanced Encryption Standard*, Federal Information Processing Standards Publication 197, 2001. [Online]. Available: http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf

[20] NIST, *Cryptographic Hash Algorithm Competition*. [Online]. Available: http://csrc.nist.gov/groups/ST/hash/sha-3/index.html

[21] A. Bogdanov, D. Khovratovich, and C. Rechberger, "Biclique cryptanalysis of the full AES," in *Proc. ASIACRYPT*, 2011. [Online]. Available: https://lirias.kuleuven.be/handle/123456789/314284

[22] A. Whitten and J. Tygar, "Why Johnny can't encrypt: A usability evaluation of PGP 5.0," in *Proc. Usenix Security Symp.*, 1999. [Online]. Available: http://db.usenix.org/publications/library/proceedings/sec99/whitten.html

[23] S. Clark, T. Goodspeed, P. Metzger, Z. Wasserman, K. Xu, and M. Blaze, "Why (Special Agent) Johnny (still) can't encrypt: A security analysis of the APCO Project 25 two-way radio system," in *Proc. Usenix Security Symp.*, 2011. [Online]. Available: http://www.usenix.org/events/sec11/tech/full_papers/Clark.pdf

[24] R. W. Reeder and R. A. Maxion, "User interface dependability through goal-error prevention," in *Proc. Int. Conf. Dependable Syst. Netw.*, 2005, pp. 60–69.

[25] R. W. Reeder, P. G. Kelley, A. M. McDonald, and L. F. Cranor, "A user study of the expandable grid applied to P3P privacy policy visualization," in *Proc. 7th ACM Workshop Privacy Electron. Soc.*, 2008, pp. 45–54. [Online]. Available: http://doi.acm.org/10.1145/1456403.1456413

[26] H. R. Lipford, A. Besmer, and J. Watson, "Understanding privacy settings in Facebook with an audience view," in *Proc. 1st Conf. Usability Psychol. Security*, Berkeley, CA, 2008, pp. 1–8.

[27] M. Hart, C. Castille, R. Johnson, and A. Stent, "Usable privacy controls for blogs," in *Proc. Int. Conf. Comput. Sci. Eng.*, Aug. 2009, vol. 4, pp. 401–408. [Online]. Available: http://www.cs.sunysb.edu/~rob/papers/socialcom-final.pdf

[28] D. Brin, *The Transparent Society: Will Technology Force Us to Choose Between Privacy and Freedom?* Reading, MA: Addison-Wesley, 1998.

[29] R. Hoffman, "The growing influence of social networks," Session of the World Economic Forum, Davos, Switzerland, Jan.27, 2010 [Online]. Available: http://www.forumblog.org/blog/2010/01/davos-social-media-session.html

[30] K. J. Nyitray. (2011, Feb.). William Alfred Higinbotham: Scientist, activist, and computer game pioneer. *IEEE Ann. History Comput.* [Online]. *33(2)*, pp. 96–101. Available: http://www.bnl.gov/bnlweb/history/higinbotham.asp

[31] P. Baran, "Introduction to distributed communications networks," Rand, Tech. Rep. RM-3420-PR. [Online]. Available: http://www.rand.org/about/history/baran-list.html

[32] T. Kilburn, R. Payne, and D. Howarth, "The Atlas supervisor," in *Proc. AFIPS Comput. Conf.*, 1961, pp. 279–294. [Online]. Available: http://dl.acm.org/citation.cfm?id=1460786

[33] R. A. Heinlein, "Waldo," *Astounding Mag.*, Aug. 1942.

[34] R. A. Heinlein, "Jerry was a man," *Thrilling Wonder Stories*," Oct. 1947.

[35] R. A. Heinlein, *The Star Beast*. New York: Scribner, 1954.

[36] C. H. Bennett, F. Bessette, G. Brassard, L. Salvail, and J. Smolin. (1992). Experimental quantum cryptography. *J. Cryptology*. [Online]. *5*, pp. 3–28. Available: http://dx.doi.org/10.1007/BF00191318

[37] P. W. Schor, "Algorithms for quantum computation: Discrete logarithms and factoring," in *Proc. 35th Annu. Symp. Found. Comput. Sci.*, Los Alamitos, CA, 1994, pp. 124–134.

[38] L. Newman, *Heather Has Two Mommies*. Boston, MA: Alyson Wonderland, 1989.

[39] M. Kan, "Beijing to track people's movements via their mobile phones," *Computer World*, Mar. 4, 2011.

[40] L. Bengsston, X. Lu, A. Thorson, R. Garfield, and J. von Schreed. (2011). Improved response to disasters and outbreaks by tracking population movements with mobile phone network data: A post-earthquake geospatial study in Haiti. *PLoS Med.* [Online]. *8(8)*. Available: http://dx.doi.org/10.1371\%2Fjournal.pmed.1001083

[41] V. Gratzer and D. Naccache, "Cryptography, law enforcement, and mobile communications," *IEEE Security Privacy*, vol. 4, no. 6, pp. 67–70, Nov.-Dec. 2006.

[42] M. Barbaro and T. Zeller, Jr., "A face is exposed for AOL searches," *New York Times*, Aug. 9, 2006.

[43] P. Golle and K. Partridge, "On the anonymity of home/work location pairs," in *Proc. 7th Int. Conf. Pervasive Comput.*, 2009, DOI: 10.1007/978-3-642-01516-8_26.

[44] J. Ginsberg, M. Mohebbi, R. Patel, L. Brammer, M. Smolinski, and L. Brilliant, "Detecting influenza epidemics using search engine query data," *Nature*, pp. 1012–1014, Feb. 19, 2009.

[45] I. Asimov, *Foundation*. New York: Gnome Press, 1951.

[46] N. Eagle and A. Pentland, "Reality mining: Sensing complex social systems," *J. Pers. Ubiquitous Comput.*, vol. 10, no. 4, Jun. 2005, DOI: 10.1007/s00779-005-0046-3.

[47] *Tor Project: Anonymity*. [Online]. Available: http://www.torproject.org

[48] L. Carroll and M. Gardner, *The Annotated Alice: Alice's Adventures in Wonderland and Through the Looking Glass*. New York: Bramhall House, 1960.

[49] K. Bamberger and D. Mulligan, "Privacy on the books and on the ground," *Stanford Law Rev.*, vol. 63, no. 2, pp. 247–316, Jan. 2011.

[50] NIST, *National Strategy for Trusted Identities in Cyberspace*. [Online]. Available:http://www.nist.gov/nstic http://www.whitehouse.gov/sites/default/files/rss_viewer/NSTICstrategy_041511.pdf

[51] K. O'Brien and T. Crampton, "E.U. probes Google over data retention," *New York Times*, May 26, 2007.

[52] S. Landau, *Surveillance or Security? The Risks Posed by New Wiretapping Technologies*. Cambridge, MA: MIT Press, 2011.

[53] J. Last. (2011, Jul. 8). Israel blocks airborne protest, questions dozens. *Forbes Mag.* [Online]. Available: http://www.forbes.com/feeds/ap/2011/07/08/general-ml-israel-palestinians8555618.html

[54] National Research Council, *Protecting Individual Privacy in the Struggle Against Terrorists: A Framework for Program Assessment*, National Academies Press, 2008.

[55] *Tenth Workshop on Economics of Information Security*. [Online]. Available: http://weis2011.econinfosec.org

[56] E. Bellamy, *Looking Backward*. Boston, MA: Ticknor, 1888.

[57] A. Dutta-Roya, "Computers [1999 technology forecast and analysis]," *IEEE Spectrum*, vol. 36, no. 1, pp. 46–51, Jan. 1999.

[58] W. Sweet, "Power and energy [1999 technology analysis and forecast]," *IEEE Spectrum*, vol. 36, no. 1, pp. 62–67, Jan. 1999.

[59] M. A. Brown, "Myths and Reality about the 1960 Census," *Prologue*, vol. 32, no. 4. Winter, 2000.

[60] D. Kushner, "One giant screwup for mankind," *Wired*, vol. 15, no. 1, Jan. 2007.

[61] BBC News, *Cyber Criminals "Should Get Tough Sentences" Say Police*, Nov. 10, 2011. [Online]. Available: http://www.bbc.co.uk/news/uk-15680466

[62] S. J. Levy, "The expanding responsibility of the government air traffic controller," *Fordham Law Rev.*, vol. 36, no. 3, p. 401, 1968.

[63] L. T. Rolt, *Red for Danger*. London, U.K.: John Lane, 1955.

[64] R. Shaw, *A History of Railroad Accidents, Safety Precautions and Operating Practices*. Potsdam, NY: Northern Press, 1978.

[65] X. Jardin, *PBS Hacked in Retribution for Frontline Wikileaks Episode*, May 29, 2011. [Online]. Available: http://www.boingboing.net/2011/05/29/pbs-hacked-in-retrib.html

[66] M. Cheung. (2009, Apr.). User education key to IT Security: Microsoft. *Network World* 13. Available: http://www.networkworld.com/news/2009/041309-user-education-key-to-it.html2009

[67] A. J. Masys, "Pilot error: Dispelling the hegemony of blamism—A case of de-centered causality and hardwired politics," *Disaster Prevention Manage.*, vol. 17, no. 2, pp. 221–231, 2008.

[68] H. R. Wilson, *The Safety of British Railways*. London, U.K.: King, 1909.

[69] H. R. Wilson, *Railway Accidents*. London, U.K.: Self-published, 1925.

ABOUT THE AUTHORS

Carl Landwehr (Senior Member, IEEE) received the B.S. degree in engineering and applied science from Yale University, New Haven, CT, and the M.S. and Ph.D. degrees in computer and communication sciences from the University of Michigan, Ann Arbor.

He has conducted research in computer science for more than three decades, primarily in what is now referred to as cybersecurity and was earlier called computer security, information security, and information assurance. For the past decade he has assisted in creating, funding and managing cybersecurity research programs for the U.S. National Science Foundation, the Intelligence Advanced Research Projects Activity (IARPA), and the Defense Advanced Research Projects Agency (DARPA) while serving as a Senior Research Scientist at the University of Maryland, College Park, and a Senior Fellow at Mitretek Systems, Falls Church, VA. For more than two decades prior, he led a research group in computer security at the U.S. Naval Research Laboratory, Washington, DC. He is now employed as an independent consultant.

Dr. Landwehr has received awards from the IEEE Computer Society, the Association for Computing Machinery (ACM) Special Interest Group on Security, Audit and Control (SIGSAC), the International Federation for Information Processing (IFIP), and the National Science Foundation (NSF) for distinguished service.

Dan Boneh received the Ph.D. degree from Princeton University, Princeton, NJ, in 1996.

He is a Professor of Computer Science at Stanford University, Stanford, CA. His research focuses on applications of cryptography to computer security. His work includes cryptosystems with novel properties, security for handheld devices, web security, digital copyright protection, and cryptanalysis. In 2002, he cofounded Voltage security, a security startup commercializing identity-based encryption.

Dr. Boneh is a recipient of the Packard Award, the Alfred P. Sloan Award, the RSA conference award for mathematics, the Ishii award for industry education innovation, *ComputerWorld*'s Horizon award, and the Terman Award. He was program chair of several technical conferences, including CRYPTO 2003 and Usenix Security 2002.

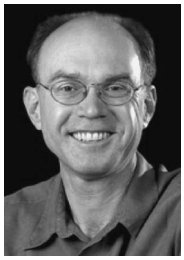

John C. Mitchell received the B.S. degree in mathematics from Stanford University, Stanford, CA and the M.S. and Ph.D. degrees in computer science from the Massachusetts Institute of Technology (MIT), Cambridge.

He is the Mary and Gordon Crary Family Professor in the Computer Science Department, Stanford University, Stanford, CA. His research focuses on web security, network security, privacy, programming language analysis and design, formal methods, and applications of mathematical logic to computer science. He has led research projects sponsored by the Air Force Office of Scientific Research (AFOSR), the Defense Advanced Research Projects Agency (DARPA), DHS, DHHS, the National Science Foundation (NSF), the U.S. Office of Naval Research (ONR); he is the Stanford Principal Investigator for the Team for Research in Ubiquitous Secure Technology (TRUST) NSF Science and Technology Center and Chief Information Technology (IT) Scientist of the Department of Health and Human Services (DHHS) Strategic Health Advanced IT Research Projects on Security (SHARPS) project on healthcare IT security and privacy.

Prof. Mitchell is the Editor-in-Chief of the *Journal of Computer Security*. His past awards include a Director's Award from the U.S. Secret Service for his efforts in connection with the Electronic Crimes Task Force.

Steven M. Bellovin received the B.A. degree from Columbia University, New York, NY, and the M.S. and Ph.D. degrees in computer science from the University of North Carolina at Chapel Hill, Chapel Hill.

He is a Professor of Computer Science at Columbia University, where he does research on networks, security, and especially why the two do not get along. He joined the faculty in 2005 after many years at Bell Labs and AT&T Labs Research, where he was an AT&T Fellow. He is the coauthor of *Firewalls and Internet Security: Repelling the Wily Hacker* (Reading, MA: Addison-Wesley, 2003), and holds a number patents on cryptographic and network protocols.

Dr. Bellovin is a member of the National Academy of Engineering and is serving on the Computer Science and Telecommunications Board of the National Academies, the Department of Homeland Security's Science and Technology Advisory Committee, and the Technical Guidelines Development Committee of the Election Assistance Commission. He has also received the 2007 National Institute of Standards and Technology (NIST)/National Security Agency (NSA) National Computer Systems Security Award. While a graduate student, he helped create Netnews; for this, he and the other perpetrators were given the 1995 Usenix Lifetime Achievement Award (The Flame). He has served on many National Research Council study committees, including those on information systems trustworthiness, the privacy implications of authentication technologies, and cybersecurity research needs. He was also a member of the information technology subcommittee of a National Research Council (NRC) study group on science versus terrorism. He was a member of the Internet Architecture Board from 1996 to 2002. He was codirector of the Security Area of the Internet Engineering Task Force (IETF) from 2002 through 2004.

Susan Landau (Member, IEEE) received the B.A. degree in mathematics from Princeton University, Princeton, NJ, the M.S. degree in mathematics from Cornell University, Ithaca, NY, and the Ph.D. degree in applied mathematics/theoretical computer science from the Massachusetts Institute of Technology (MIT), Cambridge.

Her research was originally in theoretical computer science and algebraic algorithms, but she has spent the last two decades focusing on issues related to cybersecurity, privacy, and public policy. She is currently a Visiting Scholar in the Computer Science Department, Harvard University, Cambridge, having previously been at the Radcliffe Institute for Advanced Study, Sun Microsystems, University of Massachusetts, and Wesleyan University. She is the author of *Surveillance or Security?: Risks Posed by New Wiretapping Technologies* (Cambridge, MA: MIT Press, 2011) and coauthor, with Whitfield Diffie of *Privacy on the Line: The Politics of Wiretapping and Encryption* (Cambridge, MA: MIT Press, rev. ed., 2007).

Prof. Landau is a Fellow of the Association for Computing Machinery (ACM) and the American Association for the Advancement of Science (AAAS), and a 2008 winner of the Women of Vision Social Impact Award (Anita Borg Institute).

Michael E. Lesk (Member, IEEE) received the B.A. degree in chemistry and physics and the Ph.D. degree in chemical physics from Harvard University, Cambridge, MA, in 1964 and 1969, respectively.

He is currently on the faculty at Rutgers University, New Brunswick, NJ, having previously worked at Bell Laboratories and Bellcore doing research in application areas such as digital libraries, UNIX software, and information retrieval technology. For four years at the National Science Foundation he helped manage the digital libraries initative, and he is the author of *Understanding Digital Libraries* (San Mateo, CA: Morgan Kaufmann, 2004).

Prof. Lesk is a Fellow of the Association for Computing Machinery (ACM), a member of the National Academy of Engineering, and a recipient of the Usenix "flame" award.