Scientific research in cybersecurity

Claude Kirchner
&
Ludovic Mé,
INRIA

Abstract:
Making information systems secure means ensuring the confidentiality, integrity and availability of the system's resources and services. Both resources and services must be protected; eventual attacks, detected; and efficient actions, undertaken. While global progress has been made in the security of information systems and the protection of personal data during the past twenty years, much is yet to be done, both in operations and upstream in research and development. This article focuses on scientific research related to the protection of information systems, the detection of attacks against them and the reaction to these attacks. It does not fail to mention the challenges related to: security in certain very sensitive domains; the human, societal and economic aspects of cybersecurity; and our knowledge of these menaces.

Making an information system secure means ensuring the confidentiality, integrity and availability of its resources and services. Both resources and services must be protected; eventual attacks, detected; and efficient actions, undertaken. While global progress has been made in the security of information systems and the protection of personal data during the past twenty years, much is yet to be done, both in everyday operations and in R&D. This article focuses on scientific research related to the protection of information systems, the detection of attacks and reaction to them.[1]

Widespread digitization is leading our societies toward a global cybercivilization where cybersecurity will be a key issue of viability. This fundamental transformation has come out of advances in science and technology as well as the innovations and uses resulting from them, in particular (but not only) in the digital realm. Unparalleled by scale, this trend is paradoxically raising new scientific questions for traditional fields of digital technology: computer science, mathematics, electronics, robotics, etc. It is also, given its deep impact on human beings, society and the environment, raising questions for law and economics, for the human, social and environmental sciences. It challenges them to help us handle, control and understand the trend under way.

To better analyze research in cybersecurity, let us first examine the principal reasons why digital technology has such an impact.

[1] This article, including quotations from French sources, has been translated from French by Noal Mellott (Omaha Beach, France). The translation into English has, with the editor's approval, completed a few bibliographical references. All websites were consulted in April 2020.

A first, crucial, point is that information has emerged as a fundamental scientific concept, like matter, energy or life: "*Information is information, not matter or energy. No materialism which does not admit this can survive at the present day*" (WIENER 1948). Identifying information as a fundamental concept helps us understand the digital revolution's impact. We are systems that biologically process information. Of course, we are not just that), but as such, we interact with the systems that digitally process the information we have created. These interactions, still elementary a few decades ago, have reached the point that the systems for processing biological or digital information have become complementary. As a consequence, these systems are interacting and combining with profound effects on humanity and human organizations. Cybersecurity should be seen in relation not just to the security of information systems but to all the elements that come into play. We have to better understand the impact of social networks, the mechanics of fake news. We have to better detect attacks and defenses on all information systems. These attacks might be aimed at the system for processing electronic information — what we can call "cyberattacks". In addition however, a cyberattack might also be launched with the aim of affecting the processing of human, biological information (as in the case of Cambridge Analytica).

A second point has to do with the wide-ranging deployment of digitized systems. There is no organization or sector that is not, to some degree or other, digitized: industrial processes, finance, business, transportation, health care, systems for producing, carrying and consuming energy, cities and their buildings, automatic vehicles, R&D in science and technology, etc. This digitization of the whole environment around us has led to unprecedented complexity on all scales: personal, local, national, continental, global. It is not surprising that it is hard to exercise control over the defense of all these information-processing systems.

The third point to which we would like to draw attention is the rapidity both of the evolution of computerized systems and of their deployment. The pace of advances in knowledge and of the eventual innovations stemming from them has never been so fast in all of human history. This speed calls for agile, well-informed strategies for adapting defenses, adaptations that are never definitive.

In this complex, intimately human context, so pervasive and so quickly and deeply changing, cybersecurity is a fundamental factor for building confidence — confidence in our institutions and organization our lives (personal, family, group, professional). What are the major issues to be addressed during the next five to ten years so that, by 2030 (tomorrow!), we will still be able (and if possible better able) to establish our sovereignty, personal as well as collective?

Menaces as challenges for research

The first step toward cybersecurity is to identify threats and then design means of detection and protection for countering them. An essential means is cryptography. Although cryptographic primitives and protocols are the basic elements of security, additional security services are necessary, such as authentication and access control. However the operating system or peripheral devices on the network usually provided these additional services; but they can be attacked and subverted. As a consequence, activities on information systems have to be supervised in order to detect security policy violations.

Since attacks can spread very fast, systems of protection must react automatically, or at least reconfigure themselves, so as to stem the tide. All security procedures must be meticulously integrated in critical applications for all systems (whether human or digital) that process and communicate information — systems comprising human beings, traditional information systems, industrial systems and, too, the new, "distributed" infrastructures of cloud computing and the Internet of things (IoT), which will soon connect more than a hundred billion devices.

Each step for improving security has its problems. Herein, we want to focus (but without claiming to be exhaustive) on the scientific problems that research must address in each step. Some of the problems to be discussed have been borrowed from the white book on digital security (KREMER *et al*. 2019), which presents the work done by INRIA's project teams. This white book offers us an overview of digital security that identifies key scientific issues.

Identify menaces: Know thy enemy

Search for points of vulnerability and systematically analyze them

A cybersecurity breach, once confirmed, is unfortunately much deeper than what meets the eye — than what has been detected. It can have serious consequences, even massive lethal effects (not yet observed but, as we know, possible). Cyberattacks in the coming five or twenty years will be able to wreak irreversible destruction on human constructions and the environment. Their design, the means used to execute them and their destructive potential are all increasingly sophisticated.

It is essential to know thy enemy. Academia (the "hard" sciences as well as the social and human sciences) should become massively involved in understanding attacks, those that have taken place and those that will take place. Geostrategic studies of current or coming (in the next twenty years) cybersecurity conditions could shed light on the decisions to be made. Thanks to this stronger involvement, experimental sciences could be developed in cybersecurity by applying scientific methods (ethics, reproducibility, sharing). A few high-security high-tech laboratories already exist in academia, similar to the P3 and P4 laboratories in biology. Bolstering them will be decisive enabling us to conduct observations, measurements, audits and certification and to eventually contribute to standardization.

Software attacks on hardware

A relatively new category of sophisticated attacks is growing. Thy use or produce points of vulnerability in the material components of information-processing systems, starting with processors. They typically take advantage of the physical properties of these components — of the means that modern processor cores use to optimize the management of caches and branch predictors, or to execute code in advance in order to gain time (speculative execution). Rowhammer and Spectre are recent examples of this trend. Rowhammer flips memory bits between neighboring cells while reading or writing. Spectre uses branch predictors and speculative execution to exfiltrate information through a backdoor via cache access. These very dangerous attacks can affect physical devices at a distance.

Attacks in this category share a common cause: abstraction. Typically, when proposing a security procedure at a given layer of abstraction, we tend to consider that the lower, underlying layers are solid and safe. Evidently, this is not always so. Attacks are increasingly focusing on these lower layers, closer to the system's hardware. They have gradually switched targets from software toward operating systems (OS), the kernel, firmware and, finally, hardware.

Preventing such attacks costs a pretty penny. It entails, for example, limiting the reduction of the component's surface or periodically refreshing cells through reading/writing operations. Such attacks are hard to detect since they leave no tracks on the OS or software layers.

These problems are very hard to handle and have not yet been solved. To find solutions, a clear typology of attacks must be drawn, and we have to better understand their *modus operandi* and design countermeasures at the software or hardware layers. This work will have to be done in a difficult context that might require reviewing the crucial optimizations that have long been used, such as speculative execution.

Protect

Confidence in encryption

Confidence in encryption is of primary importance. It is built on the control of cryptographic primitives and cryptanalysis. Cryptography is the science of designing cryptographic primitives, while cryptanalysis is the science of attacking these primitives. These two approaches lead to ameliorations and boost confidence in encryption.

The challenge is to organize the search for new attacks (whether using electronic or quantum computing power) on encryption algorithms by focusing on physical components correlated with secret keys, which algorithms manipulate, and formally establishing the properties of robustness for developing and installing algorithms.

Cryptographic protocols

Cryptographic protocols, via exchanges of encoded information, define security properties, such as the authenticity of an identification given by an entity active on a network. These protocols, such as those used to validate banking transactions via a mobile telephone, raise sensitive questions of security. Proofs of them are long and complicated, and involve many interactions between different cases. “Hand-written” proofs, even those made by experienced computer scientists or mathematicians, might contain errors. Formalizing and automating the protocols and the properties to be proven are the only way to obtain proofs without errors and, therefore, with a high degree of security.

The challenge here is multifaceted. Protocols have to be formally specified at the appropriate layer of abstraction. This means building a model of the environment in which a protocol will be executed and of its level of abstraction (*e.g.*, whether or not at the level of machine language). Models must also be built of the hacker’s capacity for acting on the environment where a protocol is executed, and of his eventual connivance with other malevolent entities. Flaws in the design and eventual implementation of protocols could thus be detected and corrected. The advent of 5G has made this very important.

Calculations using encrypted data

The pervasiveness of cloud computing has led to the development of homomorphic encryption, which allows, for example, the operation of addition to be compatible with an encryption function. Thanks to homomorphic encryption, work can be done directly on encrypted data without having to transmit data or make them available on a cloud in an unencrypted state. Confidence depends on the quality of the encryption algorithm. The main difficulty is to design homomorphic encryption primitives that are as universal as possible in the sense that they are homomorphic for all operations, useful or imaginable — what is called full or universal homomorphic encryption. We know how to make such primitives, but their performance in time and space is so weak that it is not realistic to use them on today’s machines and networks.

The difficult but crucial challenge for boosting confidence in cloud computing is to discover homomorphic encryption primitives that are reliable and efficient in time and space for large, if not all, classes of operations.

Encryption faced with quantum computing

The shift from von Neumann's classical model to quantum computing will change the complexity of a function executed by a program in each of these models. Algorithms such as RSA with an exponential complexity when executed in a classical model are of a polynomial complexity in a system with quantum architecture. Once quantum machines are available with enough qubits, RSA will, for sure, become unusable; and the secrets now memorized with RSA will be easy to read. New cryptographic "post-quantum" primitives with enough complexity for the two computing models have been discovered. They involve different mathematical challenges (*e.g.*, find a vector with few dimensions in a Euclidian network or decode an arbitrary linear code).

The task, clear and important, is to find and analyze proofs of complexity for these new cryptographic primitives. In addition, it will be necessary to manage ahead of time the fact that, within the next twenty years (by 2040 at the latest), most current cryptographic primitives should no longer be used, since all the stored information encrypted with them will become vulnerable.

Apart from the foregoing issues, quantum computing can be used to communicate secrets safely owing to the quantum properties of matter, which are thought to be inviolable, whence a different challenge, namely: the capacity of implementing "perfectly safe" communication protocols. The latter are now operational for distances of less than one hundred kilometers (between certain Swiss banks for instance). The challenge, in particular for physicists, is to make this possible over longer distances (*e.g.*, for communications between the earth and satellites or between ground locations several thousand kilometers apart).

Formalization and proofs for security

At present, information system security is mainly based on classical engineering solutions that are not formal. This approach has proven its mettle and, too, its limits, as reports released about several cyberattacks have shown. Formalized methods are the key to "security by design". This means building security into the information system and using formal, automated proofs that a system and its security procedures have given properties (*e.g.* that only a specified user can read the information contained in a certain file). An information system has procedures for both processing and transporting information, which, it should be pointed out, tend to be confused owing to the increasing virtualization of software defined networks (SDNs).

The challenges are important and difficult in this context, perhaps more than in others. Formal methods have, as indicated, demonstrated their effectiveness in correcting cryptographic protocols. However much is yet to be done to prove that software (a complete OS, supervisors, hypervisors, etc.) is operating correctly. One problem comes from the change of scale, since volumes of highly complex codes have to be validated. In addition, the proof has to be made for the whole stack of computing layers, from software to hardware (since, as stated, hackers are more frequently targeting the lower layers), while, of course, taking account of interactions between different layers. Both preventive and reactive security procedures will have to be proven, as discussed in the next paragraph. To cite a single example, let us imagine trying to prove that a system for detecting breaches works for a given category of attacks. This complicated task has a cost that has to be precisely calculated and compared with the cost of insecurity. This is the only way that formal methods will prevail (unless regulations make them mandatory) in security and the protection of data (in particular personal data), as is the case for the access to services in critical environments.

Detect, diagnose and contain attacks

Detect intrusions and anomalies

As explained, activities on an information system have to be supervised; and security violations, detected. Two approaches are now in use: detect breaches of security (*i.e.*, the known signals of known attacks), or detect anomalies (*i.e.*, deviations from the normal uses of the information system's services, deviations that might be evidence of an attack, whether known or unknown). Detection procedures are widely used by antivirus software, intrusion detection systems (IDS) or endpoint detection and response (EDR); but their effectiveness is often wanting. There is no guarantee that all attacks will be detected (the risk of false negatives); and above all, feedback from the field has shown that false alarms (false positives) are sometimes so numerous that they submerge true alarms, thus making it harder for the system administrator to notice them.

At present, the detection of intrusions mainly relies on analyzing network traffic packet by packet. This approach is inadequate, since a packet as such is too poor in information. This limits the effectiveness of detection systems, even though various procedures for aggregating information have been proposed. In any case, the proportion of encrypted traffic is increasing, now estimated at 50-80% of all traffic; and this trend will eventually make obsolete an approach based on searching for traffic patterns.

Two challenges crop up. By drawing on the possibilities for computing encrypted data (as previously mentioned), detection procedures could, we imagine, work directly on the network's encrypted traffic. Or, assuming that the information is no longer available on the network, we can imagine using other (semantically enhanced) sources of information, at the application or operating system layer, for example.

Methods and procedures for detecting anomalies are less used. Nonetheless, various methods have been devised for building benchmark models of an information system's behavior such that activities on the network can be compared with the benchmark, and an alarm sounded in cases of anomalies. Successful solutions based on machine learning are likely to set off a revolution in the detection of anomalies, just as they have done in image-processing. However it is no simple matter to apply these techniques in cases of data security. One task is to clearly define what may and may not be done. Two difficulties spring up. First of all, the data used for machine learning are seldom public. Secondly, many AI approaches (*e.g.*, deep learning) have a major flaw when examined with security in mind: we are still unable to explain AI's results. We have to skirt around these two difficulties. A direct way to do so is to avoid machine learning. In this case, the benchmark model is fixed in relation, for example, to the technical specifications of the services offered or to the system's security policy. The activities observed on the network can then be analyzed in relation to these specifications.

Another important challenge has to do with the testing or eventual certification of detection systems. On the practical level, there is, at present, no platform where academics may test their ideas and compare them with those of other researchers. Were such a platform built, experiments could be reproduced — hardly the case nowadays, since access to the code of detection systems or to the test data they have used is very seldom available. We mentioned earlier the possibility of using formal methods for proofs. The properties to be certified might be, for example, that a given category of attacks is detectable or, more broadly, that a given procedure is capable of detecting any violation of the system's security policy. Since analyzing all user activities might violate privacy, research will have to come up with detection procedures that uphold privacy.

Diagnose security violations

Security operation centers (SOCs) receive alerts (many of them false, as pointed out). Human operatives try to understand and qualify these alerts by using correlations, a major component of security information and event management (SIEM), which groups under a single meta-alert the information available about the same attack but as detected by several detection tools. This form of correlation (a merging of information), though useful, does not offer a granular analysis that reconstructs the phases of the attack and identifies the hacker's actual objectives.

For a granular analysis, the very nature of the system (machines, their connections, the services offered, the security procedures in used and their configuration, known points of vulnerability that cannot be or have not yet been patched, etc.) must be taken into account along with more general information (intensity of activities in such and such a place or the recrudescence of certain forms of attack). The correlator will also need a description of possible attack scenarios of the sort provided, for example, by a risk tree analysis.

With this information, research can take up a major challenge, namely: design automated procedures for analyzing the stream of alerts by, for example, adopting methods related to symbolic AI. In addition, research should be conducted on the visualization of information related to security (including, of course, the alerts). Some of this quite varied and highly structured information is much more important than the rest. Human operatives must have the most relevant image possible of what is happening on the information system. They must be able to browse extremely large volumes of data efficiently. Beyond this work of visualization, research should also be carried out on "good" forms of interaction.

Automate the deployment of countermeasures

Since attacks can spread fast, systems of protection must react automatically or at least be automatically reconfigured to keep attacks from spreading. Existing procedures allow, for instance, for automatically closing a port on a firewall (in order to block an attack) or ending a system process (to stop an attack under way via the process). They do not, however, assess the impact of such countermeasures and, above all, do not offer us thoughts about drawing up or modifying the security policy. If an attack is successful, then prevention procedures have not been correctly configured; and their parameters must be reviewed. Typically, the security policy itself is incorrect or incomplete, and must be amended; and new parameters for security procedures, set. There are two types of reactions, the one having to do with the configuration of security procedures, the other with the policy of security itself.

The challenge for research is to be capable of making a very rapid diagnosis of underway incidents so as to react as fast as possible. Another challenge is to prove, on the one hand, that the properties that the security policy and its implementation are supposed to guarantee are actually valid and, on the other hand, that the proposed modifications will not interfere with the information system's services. Furthermore, it would be worthwhile to be able to automatically implement a security policy (to set its parameters). For all these tasks, formal methods are indispensable for building systems capable of defending and automatically adapting themselves to moving threats (a form of "autonomic computing" that we have called "autonomous security").

Let us conclude this section by mentioning another sort of reaction but without dwelling on it, namely: counterattacks. The latter raise extremely sensitive ethical, technical and geopolitical issues. Given our current knowledge and technical capacity, there is no reason to automatically counterattack; and it would be indispensable for human beings to be present in the loop.

Protect personal data and privacy

Enforce the General Data Protection Regulation

The EU's General Data Protection Regulation (GDPR) lends support to fundamental concepts and objectives, in particular for upholding privacy and protecting personal data (health data in particular). Its enforcement must be broadened to move from words to deeds. Too many services and procedures are still black boxes lacking the transparency required by the GDPR. Furthermore, users do not have the appropriate information and interfaces for expressing their consent or opposition to uses of their data.

The challenge for research is, first of all, to design tools for analyzing the risks of privacy violations and to set up formal structures for oversight and auditing. Another challenge is to design means that will enable individuals to control their personal data and manage the fine balance between privacy and the data to be used or shared. This means designing new, above all automated, procedures for expressing consent or refusal that are robust and ergonomic.

Anonymize personal data

The privacy of data depends on the security of the procedures for managing access to them. Direct access must be maintained. We thus fall back on techniques based on encryption or appropriate security procedures. Nonetheless, data might be indirectly disclosed when they are transmitted for processing purposes or used to train machine learning algorithms. The first problem calls for procedures to anonymize the data (typically k-anonymization or differential privacy), whereas the data used for AI can be divulged, at least partially, via access to the classification algorithm.

Research must thus take up the challenge of designing robust anonymization procedures. This is hard to do given the diversity of the data available and the possibilities of crossing data. In principle, the distribution of data deprives hackers of the benefits of centralized access, but it forces us to come up with strategies that minimize the cost of access to the data for machine learning or data-processing purposes.

See to the security of…

It is important to take account of certain sensitive contexts related to applications. We shall take three examples from: the Internet of thing (IoT), industrial systems and artificial intelligence (AI). Of course, other sensitive domains should also be taken under consideration, such as health or robotics.

…the Internet of things

It is relatively easy to attack devices on the IoT, mainly because security has usually not been designed into such devices and their features. In addition, the number of devices multiplies the possibilities for attacks. Attacks can have quite serious effects on personal data and in the physical world, since connected devices are already (and will even more be) present in all aspects of our lives and in all settings: homes, offices, cars, cities, factories, hospitals….

This situation presents several challenges for research. First of all, security absolutely must be taken into account from the very design of these devices, of their hard- and software, of their OS and of their capacity for short-distance communication and for low energy consumption. Since these

devices have a limited computing capacity, security procedures will have to be frugal and lightweight. This also holds, of course, for any cryptographic procedures. Studies must be made of the versions of encryption suited for this new environment.

A very sensitive question is the updating of the software on these devices after, for example, a flaw has been found. How to make updates secure? This is, for sure, a question for cryptography.

Finally, as is usual, prevention, though indispensable, will probably not suffice. How can a light, autonomous but efficient supervision of the IoT (where many devices are used in contexts without any administrator) be set up in order to detect attacks or anomalies from the hundreds of millions of devices that can be used for attacks.

…industrial systems

Since industrial systems increasingly rely on software and open standards, they can be attacked like any other information system. This is a highly sensitive issue; and the impact of an attack might be catastrophic. Some currently used procedures will still be in use for many more years, even though they were not designed with security in mind. Given their limited computing resources, it is hard, or even impossible, to add cryptographic procedures onto them in order, for example, to protect exchanges. Since their technical specifications are not always public, standard security procedures (firewalls, detection intrusion systems) cannot process the data streams they generate on networks.

An obvious challenge for research is to adapt security procedures to this specific context. This requires real-time operations. The coexistence between modern, secure procedures and older procedures that cannot be modified is a question for research. Communication protocols must be meticulously examined, and interoperability, ensured. When it is difficult to integrate new procedures and devices, supervision will be essential. It is crucial to study detection procedures that are efficient, specific and capable of being rolled out without upending the current setup.

…machine learning

Often based on statistical learning processes, systems of artificial intelligence bring two major threats. The first has to do with protecting personal data. Which information from the training data used for AI can be extracted from the neural network — depending on whether or not the hacker has access to internal values in this network? The second threat concerns confidence in AI's output. As we know, adding carefully chosen, indiscernible "noise" to an image causes it to be classified incorrectly and thus leads to an erroneous decision — what is called adversarial machine learning.

The challenge for research is to find ways to protect AI's training data. This calls for studying how these data can and are to be modified before storage and utilization. Modifications must not affect (too much) the elements indispensable for machine learning or for the performance of the operations expected of the trained network. One line of research would be to study a distributed form of machine learning so that all training data are not stored in a single same place. This would limit the impact of an attack.

The prevention of adversarial machine learning necessitates, as a first step, understanding the weaknesses of learning strategies so as to determine which attacks are possible, how they operate, and how to counter them. It is also worthwhile studying how the supervision of the interactions between the layers in a network can help us observe and qualify eventual illegitimate activities. Such a study would also help us understand interactions between the more robust layers in the system.

Interactions with human beings and human organizations

In interactions between machines and people, either party might be the hacker, the vector of an attack, or the victim. As pointed out in the introduction, it is crucial to exercise control over how human systems for processing information and digital systems interact, cooperate and combine with each other. Several questions thus arise for research in a variety of disciplines, ranging from what Michel Serres has called the hard to the soft sciences. Let us succinctly present four questions we deem important.

In relation to information-processing by human beings, a first challenge is to understand social interactions in the context of a constant evolution of electronic media and the global trend toward the digitization of societies. Since digital information systems can (via manipulation) induce cognitive biases, the instillation of disinformation (often called fake news) is to be studied. The granular adaptation of disinformation to its human targets involves interactions on the social networks, as the Cambridge Analytica affair has confirmed. It is important its detect and analyze such phenomena, but it will be difficult to adopt countermeasures, although they might be based on the advances discussed hereafter.

A second challenge is to understand and anticipate cybersecurity's geopolitical, economic and societal effects. Though at the center of articles in English, this impact has received less, much less, attention in Latin countries, in particular France. A few studies exist; but national strategies need to be drafted and defended internationally in consonance with our allies. This is the condition for defending our values, know-how and firms. A game-changer would be to have a well-prepared, well-formed, coherent team to represent France in standard-making organizations.

A third challenge is education. Users who are not well aware of security issues are often the weak link in the global chain of cybersecurity. Technical solutions alone do not suffice. Along with them, a strong "cybersecurity culture" must be instilled; and for this, education is a key. Major efforts must be made to diffuse knowledge and raise the awareness of the public: citizens (including children and adolescents), technicians, engineers and security experts as well as economic and political decision-makers. The shortage of skills in cybersecurity is a major handicap on sovereignty (national, digital, entrepreneurial and individual). Schools must introduce pupils to the basics of computing and cybersecurity. Throughout life, every citizen must be (re)made aware of "good practices" and "cyberhygiene". Obviously, professional users should be capable of understanding the risks of cyberattacks at the workplace and knowing how to react to them. They should be trained with this in mind. System administrators should have their training regularly updated with information on new menaces and new means of defense. Finally, France and Europe need more experts in cybersecurity. While several public or private institutions do provide training, major efforts are still to be made.

The last challenge we would like to mention is the need for multidisciplinary studies of people-machine interactions. As pointed out, techniques, though indispensable, do not suffice. They have to be as simple to use as possible, since human errors are a major source of security problems — errors often due to the poor quality of people-machine interactions and interfaces. These interfaces should always be designed to avoid involuntary errors and see to it that users are well aware of the consequences of their actions. The design of such systems requires more multidisciplinary research by computer and cognitive scientists.

Conclusion

There is no "small" challenge in cybersecurity. The chain of protection is as solid as its weakest link. The challenges for research in each link on this chain represent feats with have varied consequences. For example, AI or quantum computing will generate visible disruptions.

France has an academic system that contributes, at the highest international level, to advances in knowledge on cybersecurity, especially in cryptology and formal methods (GROUPE… 2017). The advances, actual or potential, nourish a very fertile web of small, medium-sized and big firms that are well recognized internationally for their qualifications and know-how. A cultural and organizational challenge is for academics innovators to work together and stimulate each other. Firms, research centers, and universities should work and innovate together.

We have frequently referred to our society's underlying values. For all the aforementioned links on the security chain, a general point for consideration is to reflect on the ethics of all aspects of cybersecurity. Input from the individual, entrepreneurial, local, regional and national levels should be collected and coordinated by the national ethics advisory board (CCNE: Committee Consultatif National d'Éthique) on science, technology, and the uses and innovations of digital technology (GANASCIA *et al*. 2018). These reflections should nurture European and international discussions in a process that allows each entity to state its hierarchy of values. This should make it possible for referring new uses and innovations to a shareable and, if possible, consistent corpus of reflection on ethics.

The thoughts presented in this article are a short, necessarily schematic and incomplete, account based on INRIA's white book (KREMER *et al*. 2019) and on several roadmaps to which interested readers can contribute.[2] In conclusion, we would like to point out that, at the European level, projects like SPARTA[3] are trying to work out an overarching view that encompasses analyses of national roadmaps.

References

GANASCIA H.G., GERMAIN L. & KIRCHNER C. (2018) "La souveraineté à l'ère du numérique. Rester maîtres de nos choix et de nos valeurs" (CERNA ALLISTENE) 39p., available via http://cerna-ethics-allistene.org/digitalAssets/55/55708_AvisSouverainete-CERNA-2018.pdf

GROUPE DE TRAVAIL CYBERSÉCURITÉ D'ALLISTENE [Alliances des Sciences et Technologies du Numériques] (2017) "Cartographie de la recherche académique française en cybersécurité", June, 38p., available via https://www.allistene.fr/files/2018/03/VF_cartographie_2017-06-13.pdf.

KREMER S., MÉ L., RÉMY D. & ROCA S. (2019) Editors of *Cybersecurity: Current Challenges and INRIA's Research Directions* (Le Chesnay: INRIA) White book n°3, June, 172p., available via https://files.inria.fr/dircom/extranet/LB_cybersecurity_WEB.pdf.

WIENER N. (1948, revised edition 1961), *Cybernetics, or Control and Communication in the Animal and the Machine* (Cambridge, MA: MIT Press).

[2] Such as the ECS work group on cybersecurity (www.ecs-org.eu/working-groups/wg6-sria-and-cyber-security-technologies) & SecUnity (https://it-security-map.eu/en/home/).

[3] For "*re-imagining the way cybersecurity research, innovation, and training are performed in the European Union*" (https://www.sparta.eu).