Adoption Challenges for Cryptographic Protocols

Konstantin Fischer | Ruhr University Bochum
Ivana Trummová | Czech Technical University in Prague
Phillip Gajland | Max Planck Institute for Security & Privacy, Ruhr University Bochum, and IBM Research Europe
Yasemin Acar | Paderborn University and The George Washington University
Sascha Fahl | CISPA–Helmholtz-Center for Information Security
M. Angela Sasse | Ruhr University Bochum

We interviewed cryptography experts from academia and industry to learn from their experiences with the design and deployment of cryptographic protocols. We present adoption challenges, including misaligned incentives in academia and standardization, mismatched assumptions, low-quality reference implementations, and usability issues.

Cryptographic primitives and protocols are the essential building blocks for modern security protocols. Designing and deploying these is a complex process, with pitfalls at nearly every step.

To understand what can be done to improve the processes by which cryptographic algorithms and protocols are designed, analyzed, implemented, deployed, and adopted, we interviewed 21 seasoned cryptography experts from academia, industry, nonprofit, and governmental organizations.

We report a set of identified challenges and suggestions for improvements in bringing cryptography innovations from research papers to real-world products. Our study answers the following research questions:

RQ1. What steps are involved in the adoption of cryptography, and who are the relevant stakeholders?
RQ2. What are the key obstacles hindering the widespread adoption and correct use of cryptography?
RQ3. What are potential ways to overcome these obstacles?

Through thematic analysis[1] of our 21 expert interviews, we first develop and propose a map to help make sense of, and argue about, the complex and ever-changing dynamics of cryptography adoption. The map illustrates the cryptography adoption ecosystem's relevant actors, processes, and artifacts. Then, we report on key blockers involved in cryptography adoption, investigate challenges to its effective adoption, and identify potential paths to improve future adoption.

This article is based on our USENIX Security 2024 paper,[2] summarizes its findings, and extends it by reporting on additional themes and challenges for security protocols sourced from our original interview dataset. We recommend reading the original paper for a more detailed description of our approach, sample, and analysis, as well as a more thorough discussion of our findings and quotes from the interviews.

Method and Sample

Our 21 study participants had at least ten years of experience researching, designing, standardizing, or implementing cryptography and had high standing and visibility in the cryptography community.

Digital Object Identifier 10.1109/MSEC.2024.3449809

Instrument Development

To scope the problem space and inform our study design, we started with three researchers who had a strong publication record and played a leading role in a cryptographic protocol or application that is widely used today. We conducted in-depth (80–110 min) interviewee-led interviews on their experiences with that successful deployment. The key question we posed was as follows: What obstacles or blockers did they encounter on the adoption path, and what did they have to do to overcome them? We prepared transcripts and shared them with the interviewees, together with several clarifications and follow-up questions, which they answered in writing. We developed a semistructured interview guide for the remaining expert interviews based on our analysis of these three interviews.

We then conducted 18 semistructured interviews, starting with experts mentioned in the initial interviews. Although the interview guide was provided up front, we allowed interviewees to focus on the questions they thought were most relevant. Since we were looking for insights beyond the published literature, we asked for their opinions on the causes of problems.

For a complete description of the method and sample, please refer to our USENIX Security 2024 publication.[2]

Ethics

We modeled the interview study after the ethical principles for human subject research involving information and communication technologies outlined in the Menlo report.[3] The research plan; interview procedure; data collection, storage, and analysis; and all involved researchers adhered to the strict German data and privacy protection laws and the General Data Protection Regulation (GDPR). We informed all interviewees about the study procedure and data handling before they signed up for the interviews. We encouraged them to get informed before deciding and offered to answer any potentially upcoming questions. We explained to interviewees that they could skip any question for any reason. We did not compensate our interviewees, as they were all highly successful individuals motivated to work on cryptography to make the digital world safer.

We removed parts of the interviews that participants flagged as too sensitive to transcribe, deidentified participants, and removed any information that would easily identify our participants from the transcripts. After checking the transcripts for correctness, we deleted all audio recordings.

The Cryptography Ecosystem Adoption Path

Our participants referred to stakeholders and processes that are part of what we have called the "cryptography adoption path" embedded in the cryptography ecosystem. Figure 1 shows a map of that ecosystem containing entities (actors), activities, and artifacts (products).

This map helped us structure the actors and processes in bringing cryptography from papers to products. We grouped stakeholders into entities that performed different roles or jobs in the cryptography ecosystem, which a single person or groups of persons may perform. Interviewees explicitly mentioned both the entities and the processes in the context of turning cryptography research output into products. The path sequence was pieced together from partial descriptions across the interviews, i.e., no single interview described the implementation path as such, and thus should be read as a hypothesis that may change after further evaluation.

We report our results, supported by our interview data, following the path of bringing cryptography theory from research papers to end-user products.

From left to right, grounded in our results, we identify the following areas on the map:

1. algorithm and protocol development
2. standardization
3. secure implementation (cryptography libraries)
4. product development
5. adoption and use of cryptographic products.

Algorithm and Protocol Development

The adoption path starts left on the map with the design of cryptographic algorithms and protocols. Cryptography researchers create cryptographic algorithms and protocols, which they publish as academic research papers or specification drafts, thus making them available to the community for 1. cryptanalysis, i.e., looking for potential flaws and weaknesses, and 2. security proofs and formal verification, which can show that under a set of chosen assumptions, a given algorithm or protocol is infeasible to break.

Successful cryptanalysis, security proofs, and formal verification are commonly also published as research papers and inform the design of new or improved cryptographic algorithms and protocols. The result is a feedback loop that improves cryptographic designs through academic research.

Standardization

Cryptographic research output that has passed muster with the academic research community might become a standard. Standards development organizations can provide a platform and process for cryptography researchers and industry stakeholders to establish consensus and create standard specifications and documentation. This can, e.g., happen in the form of a drafting process

of open working groups, as in the Internet Engineering Taskforce (IETF), or the form of competitions, as often organized by the U.S. National Institute of Standards and Technology (NIST) for cryptographic primitives. Product vendors can choose to put resources into standards development for one of three reasons: they want to interoperate with their own or a competitor's systems, implement a standard for compliance reasons, or use it as an argument in marketing their product.

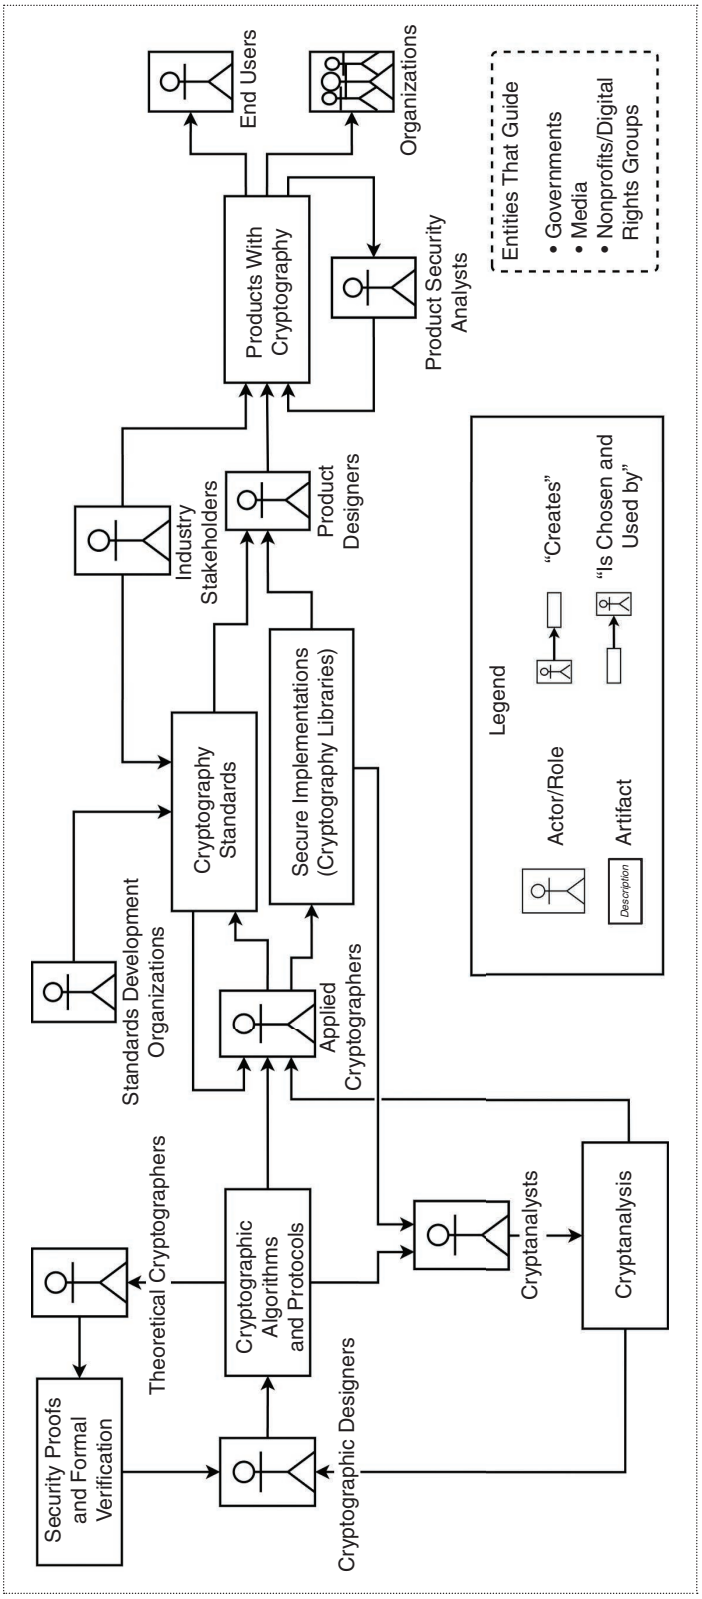

Figure 1. An illustration of the cryptography ecosystem and the path of cryptography adoption.[2]

Secure Implementation

Cryptography research papers can provide proof-of-concept implementations; these are generally not robust enough to be used in a product. Applied cryptographers take the cryptography research output and turn it into implementations that can be widely used by noncryptography experts, i.e., everyday software developers. Implementations may be provided in the form of cryptography libraries. Cryptanalysts and security researchers often scrutinize applied cryptographers' output for implementation flaws and feedback the results to improve robustness.

Product Development

Hardware or software vendors and developers may want to protect the data their products handle. To do that, they can select a relevant cryptographic library, which becomes part of the product they create. This product can then again be analyzed by "product security analysts" for security vulnerabilities stemming from, e.g., unintended or unexpected usage of the API of the chosen cryptography library.

Adoption and Use of Products With Cryptography

The last actors on the cryptography adoption path are everyday end users and organizations who choose and use specific cryptographic products. They can choose inherently insecure products or use products in unintended ways, thus risking their privacy or security.

Entities That Guide

We find the following additional entities to be worth noting. They can influence cryptography adoption in multiple ways.

Governments can create legislation and regulations that impact the funding of research and critical Internet infrastructure, decide on standards that need to be implemented by certain industries, and specify security and privacy requirements for cryptographic products.

Media outlets can influence public opinion and focus public attention on certain topics. This might impact what research projects get funded, which standardization organizations are deemed trustworthy, how high-security aspects are prioritized in software companies, and what application an end user installs on their device.

Consumer advocates and digital rights groups can aid end-user decisions and support the development of secure implementations and infrastructure. For example, the Electronic Frontier Foundation's (EFF) Certbot (https://certbot.eff.org) and the "Let's Encrypt" initiative (https://letsencrypt.org) were driving forces for the widespread adoption of HTTPS on the web.

Adoption Challenges

We report and discuss the common themes that emerged from our thematic analysis of the expert interview data. These include misaligned incentives of cryptography ecosystem actors, misguided standardization efforts, low-quality reference implementations, and missing resources. We acknowledge that our findings may be limited by interviewing only cryptography experts. The challenges faced by nonexpert developers, who often lack access to or awareness of relevant standards and may resort to insecure custom implementations, may be underrepresented. We refer to complementary research that explores these challenges in greater detail, like Acar et al.[12] and Huaman et al.[13]

Incentives in Academia and Beyond

Incentives hold significant sway over the trajectory of human endeavors, both in their presence and in their absence. When examining the path of cryptography adoption, a central aspect to investigate is the actors' incentives. They often do not align perfectly with our overarching goal of secure cryptography adoption. We also identified missing and even conflicting incentives.

Academic research is mainly funded through grants. Our interviewees generally agreed that the most important metric for a significant majority of grants is the number of published top-tier research papers. Thus, for many researchers, there is not much to gain from putting rigorous work into participation in cumbersome standardization processes or into usable and secure implementations of their cryptography. At the same time, our interviewees also noted that cryptography conferences are starting to reward the "engineering side" of research more. Some interviewees highlighted the Real World Crypto (RWC) Symposium for its efforts to bring academia and industry together. One interviewee argued that RWC is not as inclusive as some would want it to be: "RWC, even by its name, it conveys what the message is: 'Don't bring your theoretical nonsense here. We don't want to hear about it!'" (P13).

Some interviewees expressed favor for the concept of *boring cryptography*, meaning crypto that "simply works, solidly resists attacks, never needs any upgrades" as opposed to *interesting cryptography*, which is more complex, or even flawed, thus offering opportunities for ample amounts of cryptography research. One participant mentioned a controversial incentive that comes with the concept of boring cryptography: a conflicting community incentive for cryptographers themselves not to design secure cryptography but to be "exactly at the edge of things being not broken, but, hopefully, some of them do get broken so that you can keep writing papers" (P2). However, when we asked interviewees if cryptography adoption would thrive if all researchers worked toward "boring cryptography," nearly all interviewees refused:

> "I don't want to denigrate blue-sky, crazy, research, because that is what academia is for. It should be somewhat insulated from the real world. As an industry-person it feels churlish to complain that these people, who we don't pay for, are not doing our work for us." (P9)

When, on the other hand, incentives align well, we can see cryptography adoption thrive. This happened, for example, in the drafting of TLS1.3, which is arguably *the* security protocol of the Internet. The protocol is renowned enough that on the academic side, resources for standards development can be justified in grant applications, and papers on the analysis of protocol drafts are deemed impactful enough to be accepted at top-tier conferences. On the industry side, there are economic incentives for browser vendors and cloud providers to help develop a faster, more efficient, and more secure version of TLS. Another example of aligning incentives is when the Signal team helped WhatsApp realize that it would actually be beneficial for WhatsApp to start using the end-to-end encrypted Signal protocol to both have a better security stance for their users and also be less easily held accountable for what users are sending using their platform by potentially overreaching governments.

While cryptography adoption can certainly be accelerated by addressing many manageable problems, widespread adoption of a protocol remains an excruciatingly big task that is best addressed by funding large, long-term research projects.

> "You need to accept that some things are going to take 20 years. We accept that in the physical world. We accept that the highways in a city are really horrible and it's going to take 5–20 years to fix that. We accept that when you need to repair a bridge, sometimes the way that you have to repair a bridge is that you have to build a brand-new bridge next to the one that you want to repair. Those are necessary." (P21)

One crucial insight is that research papers do not make it into practice, not because cryptographers do

not care but because they are assessed by the number of top-tier scientific articles they produce. At the same time, there is a significant portion of theoretical work in cryptography that does not have practical application as its end goal. Phillip Rogaway[4] argued that cryptography researchers should focus more on real-world impact. Still, the rewards of doing so are much less certain than, assuming you have learned to do so successfully, the rewards for producing yet more papers.

Cryptographers' Assumptions and End-User Expectations: Key Management Is Hard

When drafting a secure protocol, the designers make certain assumptions. While these assumptions are essential for the design phase, they should not be considered set in stone but should be reconsidered later in the process. Such assumptions must be flagged when the protocol is implemented in a given product. A theme our interviewees kept circling back to was end-user key management, which we want to use to illustrate some assumptions and end-user expectations. End-user key management was described by many interviewees as an unsolved problem that does not get the attention it deserves.

> "Key management in general is a really big failing in both the usable security and cryptography area. If you look at most of the cryptography research that is out there, [...] they say: 'If key management can be done, look at this cool thing we can do!' [...] I would argue that nobody has actually solved key management." (P19)

We use the broad term *key management* to describe the tasks and duties that come from security protocols relying on cryptographic keys as the identities of individual end users or to protect potentially sensitive data. For cryptographic communication, these tasks can typically include key generation, secure key storage, keeping track of a private key, key extraction, key import, key synchronization, protecting one's private key, distributing public keys or identities, discovering contacts' keys, and validating contacts' keys.

Example: Access Recovery

For a cryptographic communication protocol between Bob and Alice, Bob and Alice are assumed, at any time, to be in control of their respective private keys, to decrypt any messages that were encrypted for them or to prove authenticity by signing the messages they send. While assumptions like these are integral to protocol design, several of our interviewees agreed that these assumptions will not always apply in real-world circumstances, which can lead to problems or hinder adoption. In the example above, Bob could accidentally break his phone and thus lose access to his private keys stored on the phone. Our participants argued that such real end-user concerns should be considered already at the protocol design stage and that not doing so can make or break the success of a security protocol.

If we look at real-world solutions, we see backup strategies for user keys, e.g., password managers like 1Password who, upon setup, provide their users with an "emergency kit": A PDF document with the account details, the secret key to decrypt their password vault, and a blank field to write down their account password. They recommend printing it out and locking it away in a safe or giving it to a trusted peer. There is currently no empirical evidence of whether this approach for recoverability is viable for end users.

The FIDO Alliance chose a different solution when adoption of the FIDO2 passwordless authentication standard stalled. For security reasons, a private key would never be allowed to leave the hardware device on which it was created. This strong security requirement led to a similar situation, where losing a device could mean losing access to all online accounts. When it became clear that adoption was not taking off, the FIDO Alliance moved forward and proposed *multi-device FIDO credentials*,[5] more commonly known as *passkeys*. Passkeys effectively allow for cloud synchronization of private keys among devices via the user's OS provider: Google, Apple, or Microsoft. The FIDO Alliance arrived at this security compromise since they deemed adoption unlikely if users were given the non-trivial burden of 1. *owning* and 2. *managing* one or more backup devices. End-user key management is complex, and we know that users will not adopt security measures that have little perceived benefit but threaten to introduce a new burden onto them.

> "Users are accustomed to that whenever on a website you forget your password, you click on 'I forgot my password' and it's being sent to you again. This is the expectation that people have of recovering anything they're using." (P5)

In contrast, the messaging app Signal does not offer backups by default. This assumes a security model that, again, is helpful in guiding protocol development but can lead to issues down the road when everyday users adopt Signal but are not aware of possible duties put onto them.

> "[Signal] was designed to be used by people who are issued a phone at birth and they keep that one phone until they die. They never lose the phone. This is an artificial world where no human being sits. And they

> deliver the best security possible to those who do not exist. You have to make compromises, because you have human beings using your stuff." (P21)

This quote is, of course, hyperbole: Signal does offer migration features and manual backup options; however, these are nondefault and need manual interaction and knowledge. One could argue that *Johnny* will not benefit from anything that is not default or automated.

Example: Endpoint Security

It is often assumed that the users' private keys will never be stolen and plaintext messages will never be accessed. To secure conversations from more than passive machine-in-the-middle attackers trying to breach message confidentiality, a holistic secure design approach would take a step back and look at additional attack vectors that might need innovation. Examples named for this were self-destructing messages, properly using security hardware like TPMs, or properly encrypted device cloud backups. The reliance on endpoint security could decrease with access to new usable and easily deployable protocols for matters of this kind.

Example: Contact Key Authentication

Whitten and Tygar[7] first explored the usability of end-to-end encrypted e-mails in "Why Johnny Can't Encrypt." After 20 years of "Johnny" papers, problems related to end-user key management turned out to be a main issue every time. The concept of public keys is not intuitive and was impossible to automate without sacrificing security by, e.g., introducing the need for a trusted third party. The pattern continued with research on the usable security of end-to-end encrypted chat apps, which can handle key distribution and key discovery entirely transparently for the user. Still, it can only offer opportunistic end-to-end encryption this way. Opportunistic end-to-end encryption still allows active network attackers to mount machine-in-the-middle attacks by distributing attacker-owned public keys to users. Chat apps started offering key authentication ceremonies to achieve authenticated end-to-end encryption, where users can be sure that no active or passive attacker is reading along or interfering with messages. These ceremonies are too cumbersome and are rarely conducted by end users: key management is complex. Recent innovations by Whatsapp and iMessage, in the form of key transparency,[8] promise to offer automated lookups in auditable key directories. In the current implementation, WhatsApp's auditable key directory does not yet protect its users from a rogue service provider, but it does add security from certain previously possible threat models, like nation-state attackers trying to compromise key distribution.

The essence of the examples above is that even in the early design stages of a security protocol, it is imperative to consider your end users' possible goals, capabilities, restrictions, and worries. If a security protocol introduces new risks, like losing access (availability is one of the security goals), users might perceive the risks as higher than the benefits and thus not adopt your protocol or product. This includes private end users, other project maintainers, and developers downstream.

Standardization Processes

Both our interviewees with industry backgrounds and those with academic backgrounds noted that over the past twenty years, the IETF standards drafting process has become slower and more cumbersome. Our interviewees who directly participated in standards development at IETF or ISO described the processes as slow, exhausting, or tedious. We additionally observe skepticism toward the motivations of certain parties involved in standardization processes. While the objective might seem to focus on creating efficient and secure standards, the process can be influenced by commercially driven companies and government agencies protecting their interests, which, in some cases, can lead to compromised standards. The prime example of a compromised standard is Dual-EC, a backdoor-able random number generator that was standardized by ANSI and NIST. The NSA secretly paid RSA, a leading security firm, US$10 million to make Dual-EC the default method in RSA's BSafe software. The NSA promoted Dual-EC by leveraging RSA's early adoption to secure approval from NIST.[11] If an attacker can predict the output of a pseudorandom number generator, they can, e.g., deduce the keys that one or both sides of a secure connection will use and decrypt communications.

In a less subversive instance, in the TLS1.3 drafting phase, the financial industry resisted enforcing forward secrecy, advocating for the reintroduction of RSA static keys, which allows them to cheaply break open otherwise secure TLS connections within their networks to inspect traffic for malicious activity, which is mandated by regulations. This example underscores how industry motivations can be centered on minimizing short-term costs and engineering efforts, which often conflicts with achieving the most secure outcomes.

In all other cases, standardization work was considered unattractive to most academic researchers due to its tedious nature and difficulty securing funding. Nonetheless, all our interviewees generally prefer open standardization processes over closed ones. On the other hand, industry stakeholders usually lack incentives to invest resources beyond the minimum required, as additional investment, making the open standard more straightforward for everyone, would also benefit their competitors.

For some cryptography standards, particularly from closed standardization processes, our interviewees attribute the missing adoption to a failure to identify the standard's users and use cases meaningfully. Several interviewees agreed that standards go wrong if real-world use cases are missed or misunderstood.

> "Often, the standardization process is done in a vacuum, where they don't talk to the potential customers in any meaningful way [...] then it gets deployed and the customers think that this is not very good, and either the deployers then end up with the bad system, or they skip the standard completely." (P1).

Standardization processes represent an essential part of the adoption path, but this is where cryptography research output meets stakeholders from government and industry, who have goals beyond just selecting the best cryptography; plus, there are many misconceptions and misunderstandings. The time and effort investment for cryptography researchers who participate in those processes is significant, and the outcomes are uncertain, so it is not surprising that only a small number take it upon themselves to participate.

Reference Implementations

Interviewees highlighted the significance of reference implementations of algorithms and protocols, noting their role at various levels, including research ideas, standards, and cryptographic library API usage. Reference implementations serve as practical examples for developers, aiding protocol implementation and verifying interoperability. However, the quality of these reference implementations varies greatly. This leads to potential disaster when developers incorporate them into products without comprehending their limitations or vulnerabilities.

The distinction between reference and optimized implementations was emphasized, pointing out that the latter, although performance-enhancing, is unsuitable for general use. Misuse of reference implementations can lead to security vulnerabilities, most commonly by copying insecure configurations or hard-coded credentials from such deployments.

Our interviewees also observed code labeled as a reference implementation when, in reality, these were nothing more than proof-of-concept implementations with little regard to real-world security concerns. These and low-quality reference implementations can highly impact the adoption of secure cryptography by rendering them unusable for practical applications. Some suggest communicating the purpose and qualities of good reference implementations clearly and labeling proofs of concept distinctly from actual reference implementations. Successful and secure adoption of newly designed protocols requires "flooding the space" with high-quality implementations to ensure everyday developers wanting to adopt the protocol have access to secure and hard-to-misuse APIs.

The skills required to create high-quality reference implementations differ strongly from those needed for successful academic research and writing. This can explain the observed variability in implementation quality. Our interviewees note that only a few people on this planet are experts in cryptography *and* secure coding. Yet, we find that precisely those "double-experts" tended to be the ones enabling and driving the secure adoption of cryptography in the past. Faced with this limit, the alternative proves to be a concerted effort among multiple experts from these different fields, working together to achieve a greater number of high-quality crypto libraries and reference implementations. Having more experts work on the same challenge then again implies the need for better funding in standards development.

Need for Automation: Computational Cryptography

Since the pool of available "double experts" needed for successful security protocol development is very limited, and resources are sparse, our interviewees did not refrain from suggesting that we automate everything we can: implementation attacks, code audits of existing implementations, and security proofs for standard specifications.

Next to recommendations for automating protocol attacks (e.g., with fuzzing), our interviewees called for more research into automating code audits. They drew a picture of an everyday developer specifying their security needs and giving those as input to an auditing tool that analyzes how the developer puts together cryptography and protocol building blocks.

> "What you should have is a specification of the desired cryptographic properties and automated ways of verifying that the way that you used the cryptography meets those things." (P4)

Multiple interviewees told us they see a need to automate the testing and proofs of standard specifications. Current standards are usually written in prose and pseudocode snippets. This makes testing and security proofs hard to achieve and not easily scalable. Standards usually come with reference implementations in one or multiple languages, which are crucial for adoption. Our interviewees, however, went one step further and called for machine-readable and even executable protocol standard specifications.

> "Standards are large, complicated bunches of, essentially, code within a language that we can not execute and can not write unit tests against. And that is problematic. Only humans can read them and try to run them in their heads. I have abstract problems with the way we build standards in general. I think they should be reference implementations and should be executable." (P9)

Automating protocol audits and protocol proofs would lead to crucial increases in efficiency in these areas, where we know that resources and expertise are limited. The research field of high-assurance cryptography and computer-aided cryptography[9] is up and coming in this regard. High-assurance cryptography combines program verification and cryptography engineering techniques to create efficient cryptographic software with machine-verified proofs. It ensures memory safety, functional correctness, provable security, and no timing leaks. High-assurance cryptography is already used in, e.g., some integral parts of the Mozilla Firefox Browser, the Linux kernel, and the Microsoft TLS Library *msQuic*.

> "The future is to have proof-carrying code. So that you really have a computer verified proof that the code is correct. Then you don't need to read the code anymore, you basically need to read the specification that it's compared to by this proof. [...] Ideally, we have computer readable formal specification within standards." (P5)

Discussion and Recommendations

We discuss our findings in the context of our research questions and make recommendations how to improve the process by which security protocols are designed, analyzed, implemented, and deployed.

The Cryptography Ecosystem

Our participants' descriptions of adoption paths and the ecosystem were diverse. They identified a multitude of actors and activities that may help or hinder the process of bringing cryptography from papers to products. We hope that clearly identifying the involved roles and the steps that cryptographic innovations have to go through enables parts of the cryptography and security communities to focus future research on better understanding, describing, and overcoming the challenges of bringing cryptography from papers to products.

Challenges of Cryptography Adoption

To achieve wide adoption of cryptography, collaboration and communication are key. It is essential that different clusters of experts communicate with each other and ensure that the needs, requirements, and results of their work are understood. Theoretical cryptographers make assumptions when designing cryptographic algorithms. However, once these algorithms are turned into an end product, the initial assumptions are often lost. This can lead to security issues in the adoption of cryptography. We need to find clear ways to communicate these assumptions to groups of people outside the theoretical landscape, such as developers and users.

Of course, not everyone needs to be an expert in multiple areas. However, our interviews have shown that the role of a translator, "a crypto plumber," or a person in the middle, is often poorly rewarded and insufficiently incentivized. Our results suggest that there is certainly a need for people to step into this role. We have also identified pain points: gaps in terminology, documentation that is not understandable by the people using it, and developers having a hard time using reference implementations. A problem adjacent to communication is unclear responsibilities. Our interviewees observe that in multiple areas, unusable products of research, reference implementation, or vulnerabilities result from unfinished work. Research ideas end up unmaintained, not "production ready."

One common challenge that arises across all of the topics discussed in the interviews is the frequent lack of alignment of incentives. This results in the need for people to work on tasks that are crucial for the adoption of cryptography but require significant effort while offering limited rewards. Examples include members of academia being involved in standard-creating processes, researchers focusing on practical tasks to solve real-world problems, and cross-disciplinary work being challenging or implementations stemming from quality research ending up unusable.

We would like to stress that the amount of problems we identified in this work is in no way a reason for despair: we see that the cryptography ecosystem is steadily evolving, and like most of our interviewees, we are quite optimistic about reaching a future of better collaboration among actors on the path of cryptography adoption. To give an example, in 2022, Kannwischer et al.[10] created the PQClean project to improve the quality of reference implementations submitted to the NIST Post Quantum Cryptography competition. They find that when properly implemented, a set of guidelines together with a testing framework can increase code quality; reduce necessary efforts; and thus benefit submitters, the community, and the standardization body itself.

Recommendations

Based on our findings, we provide recommendations for the academic community, industry stakeholders, and standardization organizations.

Academic Community

We described how existing incentives in cryptography research can either foster or inhibit cryptography adoption. Ideally, more academic funding should not single-handedly rely on publication count when evaluation is due. If grant givers, universities, and research institutes aim to support the adoption of secure cryptography, they might consider funding and rewarding academic participation in standardization bodies or the secure implementation of cryptography.

We found that there is a lot to gain from funding academics to attend nonfree standardization meetings, and there are examples of this funding being beneficial. Academics should be actively incentivized to engage in standardization processes by participating in drafting committees and providing their expertise, with an appropriate reward. This might help the standards being developed with a better understanding of requirements and security theory.

Our results imply communication challenges among different stakeholders in the cryptography ecosystem. We recommend the cryptography research community try to establish common terminology and language to help cryptographers communicate better. We also identified communication issues among cryptography experts and engineers and recommend fostering cross-disciplinary collaboration by developing communication skills in, e.g., joint university and industry workshops and rising awareness of the potential risks of misunderstandings and the differences in terminology used by each discipline.

When building reference implementations, cryptographers should distinguish among proof-of-concept implementations that support results in academic papers and ones that support reuse by software engineers. We recommend cryptographers clearly mark proof-of-concept implementations on one hand but, more importantly, keep code readability and comprehensible documentation of their code in mind if they want to see it picked up by others. Additionally, we recognize that most cryptographers are not also experts in secure coding and thus encourage them to consult experts in software engineering when implementing protocols and algorithms, to improve implementation quality.

Industry Stakeholders

We encourage companies and organizations to consider investing in core infrastructure maintenance projects like the Open Source Security Foundation (OSSF). We also encourage standard implementers and users to reach out upstream, communicate problems and needs, or actively contribute to open standards development instead of developing and standardizing cryptographic solutions behind closed doors.

Our interviews imply that wide cryptography adoption is hindered when different stakeholders' goals and requirements do not align. Until regulations catch up, researchers interested in adoption might want to try to identify privacy improvements that benefit not only end users but also service providers, to have them help. If service providers adopt secure solutions, end users are not burdened with the choice between something secure and something (possibly) more usable. The biggest security and privacy gains for end users come from existing products adopting transparent security. There are a few well-known examples: WhatsApp implementing E2EE made a huge difference in end-user security, and Google opted to encrypt all data in transit after the Snowden revelations.

Our interviewees also pointed out the need to better automate code audits and work toward frameworks that would allow security experts to assign properties to cryptographic functions that are easy to understand for everyday developers. After developers specify what security needs they have, an automated analysis of how the developer "plumbed" the crypto functions together could greatly help avoid security-critical mistakes.

Standardization Organizations

To ensure a security protocol standard will be adopted in the real world, we recommend putting more emphasis on thinking through to the end and considering end-user requirements from early on in the design process. Specifically, we advocate following Paterson and van der Merwe's[6] push toward establishing proactive rather than reactive protocol development cycles. Proactive protocol development involves all relevant stakeholders and fosters collaboration between academia and industry, allowing one to find and fix flaws in the specification before the protocol standard is released. Standardization organizations should acknowledge contributors' resource constraints and streamline the process while providing effective scaffolding. To enhance the usability of standards, organizers should consider employing usability research methods, such as UX testing standards drafts with participants from the actual target user group.

We also recommend exploring the possibilities of machine-readable and possibly executable standard specifications to support automated security proofs and testing. We encourage close cooperation with experts from the field of computational cryptography to benefit from recent advances in theoretical cryptography on security proofs. Standardization organizations should make sure they are deemed trustworthy by further engaging with the academic community, emphasizing open communication and open competitions. The proprietary nature of some standardization organizations

can hinder analysis and adoption. A "seal of approval" of a trusted standardization organization can be a major driver for cryptography adoption.

Designing and deploying cryptographic protocols is a difficult task, with pitfalls and challenges at every step of the path. By analyzing 21 interviews with cryptography experts, we report on challenges regarding the design, implementation, deployment, adoption, and use of cryptographic protocols. We report on miscommunications and misunderstandings among cryptographers and engineers and missing responsibilities and resources, as well as misaligned goals and incentives. We are confident that our insights can inspire future research into the adoption processes of cryptographic algorithms and security protocols and also help the different stakeholders in the cryptography ecosystem learn about each other, inspiring more cross-disciplinary collaboration and engagement in protocol development. ■

Acknowledgment

We thank our interviewees for their valuable time, their helpful insights, and their openness toward our research methods. This work is funded by the Deutsche Forschungsgemeinschaft (German Research Foundation) under Germany's Excellence Strategy, EXC 2092 CASA–390781972. This work was also supported by the Grant Agency of the Czech Technical University in Prague, Grant SGS23/211/OHK3/3T/18, funded by the MEYS of the Czech Republic. This work is also supported in part by the U.S. National Science Foundation under Grant 2206865. Any findings and opinions expressed are those of the authors and do not necessarily reflect the views of the funding agencies. This work involved human subjects or animals in its research. The authors confirm that all human/animal subject research procedures and protocols are exempt from review board approval.

References

1. G. Terry, N. Hayfield, V. Clarke, and V. Braun, "Thematic analysis," *SAGE Handbook Qualitative Research in Psychology*. London, U.K.: SAGE Publications Ltd., 2017, vol. 2, nos. 17–37, p. 25.
2. K. Fischer, I. Trummová, P. Gajland, Y. Acar, S. Fahl, and M. A. Sasse, "On the challenges of bringing cryptography from papers to products: Results from an interview study with experts," in *Proc. 33rd USENIX Secur. Symp., USENIX Secur.*, 2024.
3. E. Kenneally and D. Dittrich, "The Menlo report: Ethical principles guiding information and communication technology research," *SSRN Electron. J.*, Aug. 2014. [Online]. Available: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2445102.
4. P. Rogaway. "The moral character of cryptographic work." Cryptology ePrint Archive. Accessed: Apr. 12, 2024. [Online]. Available: https://eprint.iacr.org/2015/1162
5. "White paper: Multi-device FIDO credentials." FIDO Alliance. Accessed: Apr. 12, 2024. [Online]. Available: https://fidoalliance.org/white-paper-multi-device-fido-credentials/
6. K. G. Paterson and T. van der Merwe, "Reactive and proactive standardisation of TLS," in *Proc. Int. Conf. Res. Secur. Standardisation*, Cham, Switzerland: Springer-Verlag, 2016, pp. 160–186.
7. A. Whitten and J. D. Tygar, "Why Johnny can't encrypt: A usability evaluation of PGP 5.0," in *Proc. 8th USENIX Secur. Symp.*, 1999, pp. 169–184.
8. H. Malvai et al. "Parakeet: Practical key transparency for end-to-end encrypted messaging." Cryptology ePrint Archive. Accessed: Apr. 12, 2024. [Online]. Available: https://eprint.iacr.org/2023/081.pdf
9. M. Barbosa, G. Barthe, K. Bhargavan, B. Blanchet, and C. Cremers, "SoK: Computer-aided cryptography," in *Proc. 42nd IEEE Symp. Secur. Privacy (SP)*, San Francisco, CA, USA, 2021, pp. 777–795, doi: 10.1109/SP40001.2021.00008.
10. M. J. Kannwischer, P. Schwabe, D. Stebila, and T. Wiggers, "Improving software quality in cryptography standardization projects," in *Proc. IEEE Eur. Symp. Secur. Privacy Workshops (EuroS&PW)*, Genoa, Italy, 2022, pp. 19–30, doi: 10.1109/EuroSPW55150.2022.00010.
11. J. Menn. "Exclusive: Secret contract tied NSA and security industry pioneer." Reuters. Accessed: Apr. 12, 2024. [Online]. Available: https://www.reuters.com/article/2013/12/20/us-usa-security-rsa-idUSBRE9BJ1C220131220/
12. Y. Acar et al., "Comparing the usability of cryptographic APIs," in *Proc. IEEE Symp. Secur. Privacy (SP)*, San Jose, CA, USA, 2017, pp. 154–171, doi: 10.1109/SP.2017.52.
13. N. Huaman et al., "You have to read 50 different RFCs that contradict each other': An Interview Study on the Experiences of Implementing Cryptographic Standards," in *Proc. 33rd USENIX Secur. Symp.*, 2024, pp. 7249–7266.

Konstantin Fischer is a Ph.D. student at the Chair for Human-Centred Security at Ruhr University Bochum, 44801 Bochum, Germany. His research interests include humans and cryptography, as well as end-user adoption of secure tools. Konstantin received a master's degree in IT Security from Ruhr University Bochum. Contact him at konstantin.fischer@rub.de.

Ivana Trummová is a Ph.D. student in cryptography at the Department of Information Security at the Faculty of Information Technology, Czech Technical University in Prague, Prague 160 00 Praha 6, Czechia. Her research interests include human factors in cryptography and postquantum cryptography.

Trummová received a master's degree in information security and cryptography from Czech Technical University Prague. Contact her at ivana.trummova@fit.cvut.cz.

Phillip Gajland is a Ph.D. student at the Max Planck Institute for Security & Privacy and Ruhr University Bochum, Bochum 44801, Germany, and is currently visiting IBM Research Europe, 8803 Zurich, Switzerland. His research interests include postquantum cryptography. Gajland received a master's degree in theoretical computer science from the Swiss Federal Institute of Technology Lausanne, Lausanne, Switzerland. Contact him at phillip.gajland@rub.de.

Yasemin Acar is a professor in computer science at Paderborn University, 33098 Paderborn, Germany, and a research assistant professor at The George Washington University, Washington, DC 20052 USA. Her research interests include the area of human factors in security and privacy, with a major focus on secure software development, and diversifying security and privacy research, implementing cryptography securely, and researching trust in the open source community. Acar recieved a Ph.D. in computer science from Philipps University Marburg. Contact her at yasemin.acar@uni-paderborn.de.

Sascha Fahl is a faculty member at CISPA–Helmholtz-Center for Information Security, 6123 Saarbrücken, Germany. His research interests include the intersection of computer security and privacy with human factors, particularly the investigation of end users, operators, developers, and designers of computer systems and their interdependencies with computer security and privacy mechanisms. Fahl received a Ph.D. in computer science from Leibniz University Hannover, Germany. Contact him at sascha.fahl@cispa.de.

M. Angela Sasse is a professor of human-centered security at Ruhr University Bochum, 44801 Bochum, Germany, and a part-time professor of human-centered technology in the Department of Computer Science, University College London, WC1E 6EA London, U.K. Her research interests include showing why many IT security mechanisms are "impossible for humans" and how to create less complex and low-effort security and privacy. M. Angela received a Ph.D. in computer science from the University of Birmingham, England. Contact her at martina.sasse@ruhr-uni-bochum.de.