



The Challenges of Bringing Cryptography from Research Papers to Products: Results from an Interview Study with Experts

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

<https://www.usenix.org/conference/usenixsecurity24/presentation/fischer>

This paper is included in the Proceedings of the 33rd USENIX Security Symposium.

August 14-16, 2024 • Philadelphia, PA, USA

978-1-939133-44-1

Open access to the Proceedings of the 33rd USENIX Security Symposium is sponsored by USENIX.

The Challenges of Bringing Cryptography from Research Papers to Products: Results from an Interview Study with Experts

KONSTANTIN FISCHER¹, IVANA TRUMMOVÁ², PHILLIP GAJLAND^{1,3},
YASEMIN ACAR⁴, SASCHA FAHL⁵, AND ANGELA SASSE¹

¹*Ruhr University Bochum*

²*Czech Technical University in Prague*

³*Max Planck Institute for Security and Privacy*

⁴*Paderborn University and The George Washington University*

⁵*CISPA - Helmholtz-Center for Information Security*

Abstract

Cryptography serves as the cornerstone of information security and privacy in modern society. While notable progress has been made in the implementation of cryptographic techniques, a substantial portion of research outputs in cryptography, which strive to offer robust security solutions, are either implemented inadequately or not at all. Our study aims to investigate the challenges involved in bringing cryptography innovations from papers to products.

To address this open question, we conducted 21 semi-structured interviews with cryptography experts who possess extensive experience (10+ years) in academia, industry, and nonprofit and governmental organizations. We aimed to gain insights into their experiences with deploying cryptographic research outputs, their perspectives on the process of bringing cryptography to products, and the necessary changes within the cryptography ecosystem to facilitate faster, wider, and more secure adoption.

We identified several challenges including misunderstandings and miscommunication among stakeholders, unclear delineation of responsibilities, misaligned or conflicting incentives, and usability challenges when bringing cryptography from theoretical papers to end user products. Drawing upon our findings, we provide a set of recommendations for cryptography researchers and practitioners. We encourage better supporting cross-disciplinary engagement between cryptographers, standardization organizations, and software developers for increased cryptography adoption.

1 Introduction

Cryptography serves as a fundamental pillar in safeguarding data and information within our modern society. By providing essential elements such as confidentiality, integrity, and authenticity, it plays a pivotal role in securing private data for individuals and organizations. These cryptographic mechanisms are crucial for ensuring secure communications and transactions, defending against a wide range of threats to digital systems.

However, bringing cryptography innovations from papers to products is fraught with numerous challenges. The standardization processes serve as a crucial gateway for these innovations, enabling their integration into cryptographic libraries used by developers. These libraries, in turn, facilitate the implementation of cryptographic products for end-users. However, past research and incidents have highlighted the widespread failures and shortcomings in the successful adoption of cryptographic innovations at various stages of this process.

The widespread use of cryptography in software and hardware products is essential for effective security. While there are increasing threats where cryptographic solutions could help, the full potential of cryptography is often not implemented or deployed. One prominent example is email encryption. PGP [18] and S/MIME [36] have offered end-to-end encryption for email since the 1990s. However, past research identified multiple usability problems [19, 39, 41], limited adoption [1, 42], and vulnerable implementations [32, 35].

Our work is furthermore motivated by the long struggle to get secure TLS deployed across the web [5, 28, 29, 34], and the mission statement of the Real World Crypto conference [44]: That the dialogue between cryptography researchers and developers implementing cryptography needs to be strengthened.

We believe that the cryptography and security community would benefit from a clear understanding of the challenges around bringing cryptographic innovations from research papers to products.

In this research, we develop a map of the cryptography ecosystem, understand the relevant stakeholders, processes, and key blockers involved in cryptography adoption, investigate challenges to the effective adoption of cryptography, and identify potential paths to improve future adoption.

We aim to answer the following research questions:

RQ1. *What steps are involved in adopting cryptography, and who are the relevant stakeholders?* Foundational cryptography research on primitives, algorithms, and protocols inform standardization bodies and cryptography product implementa-

tions and deployments. We are interested in the involved steps, stakeholders, and their interaction in adopting cryptography.

RQ2. *What are the key obstacles hindering the widespread adoption and correct use of cryptography?* Previous research and past incidents illustrated a limited and often incorrect use of cryptography. We are interested in the experiences and views of leading cryptography experts on key obstacles in adopting cryptography and their root causes.

RQ3. *What are potential ways to overcome these obstacles?* Increasing the adoption of and correct use of cryptography can lead to stronger security overall. We aim to identify promising paths to help with the more widespread and correct use of cryptography.

To answer our research questions, we conducted 21 semi-structured interviews with leading cryptography experts from academia and industry. Each participant had at least ten years of experience and heavy involvement in the community. We conducted iterative thematic analysis on the interview data to investigate the challenges of secure cryptography adoption.

With this work, we make the following contributions:

Insights from Experienced Cryptography Experts We conduct 21 semi-structured interviews with experienced cryptography experts from academia and industry to collect and report on their insights, opinions, and learnings about secure cryptography adoption.

The Path of Cryptography Adoption We develop and propose a *map* to help make sense of – and argue about – the complex and ever-changing dynamics of secure cryptography adoption. The map is based on the interview data collected and illustrates relevant actors and artifacts in the cryptography adoption ecosystem.

Challenges to Cryptography Adoption We report on the challenges that actors on the path of cryptography adoption face and investigate root causes for the challenges of cryptography adoption, based on our thematic analysis of the expert interviews we conducted.

Path Forward We outline root causes to tackle and recommend actions to take, on both an individual level and a community level, to foster the secure adoption of cryptography – from papers to product.

The rest of the paper is structured as follows: Section 2 summarizes related work on the topics of cryptography adoption and cryptography breakdowns. Section 3 explains the methods we used to gather and analyze data that allows us to address the research questions. Section 4 presents a resulting overview of the path of cryptography adoption. Section 5 reports on the identified challenges to cryptography adoption. Section 6 discusses the results in the context of our initial research questions, and presents recommendations for fostering the secure adoption of cryptography.

2 Related Work

We focus our discussion of related work on research identifying fundamental challenges using cryptography.

Developer Centered Cryptography. In 1993, Anderson discussed the challenges faced by cryptographic system designers due to limited information on system failures [3]. His work revealed that implementation errors and management failures, rather than technical attacks, were the main causes of fraud in retail banking systems, emphasizing the need for a paradigm shift in computer security.

The work of Georgiev et al. exposed significant flaws in SSL certificate validation across various security-critical applications and libraries [20]. The vulnerabilities arose from poorly designed APIs and configurations, making SSL connections vulnerable to man-in-the-middle attacks. The findings emphasize the need for improved API design, comprehensive testing, and enhanced documentation to ensure the security of SSL connections.

In 2018, Haney et al. conducted 21 in-depth interviews of highly experienced individuals from organizations that include cryptography in their products [23]. Within their sample, they found evidence for strong organizational security culture, careful selection of cryptographic resources, and formal, rigorous development and testing practices. Their findings support past studies that suggest that the usability of cryptographic resources may be deficient [2, 21, 33].

Heninger et al. conducted a comprehensive survey of TLS and SSH servers, uncovering widespread key vulnerabilities due to insufficient entropy during generation [24]. Approximately 0.75% of TLS certificates and 1.70% of SSH host keys were at risk of compromise. The study highlights software behaviors, including a Linux random generator flaw, primarily affecting headless or embedded devices.

Fahl et al. investigate the security risks associated with benign Android apps that utilize SSL/TLS protocols to safeguard data during transmission [16]. By analyzing 13,500 popular free apps, the study identified potential vulnerabilities to Man-in-the-Middle (MITM) attacks, with 8.0% of the apps found to be potentially susceptible. Additionally, the research underscores the significance of addressing user misconceptions and inadequate visual indicators for SSL/TLS usage, necessitating the implementation of effective countermeasures.

Cryptographic Libraries. In 2012 Bernstein et al. present NaCl, a cryptographic library that is intended to be securely usable by non-experts, to prevent “cryptographic disasters” that previous, less usable cryptographic libraries had led to [8]. The 2017 work of Acar et al. presented the first empirical evaluation of cryptographic libraries, examining their impact on code security and functionality [2]. The study found that while simplicity of APIs is important, comprehensive documentation and accessible code examples are crucial

for promoting both secure development and functional correctness. In 2022, Jančár et al. conducted a questionnaire study with 44 developers of cryptography libraries, investigating if and how they ensure that their code is not vulnerable to timing attacks [25]. They found that many developers perceive updating their code to be constant-time as too high of an investment of time and effort to actually tackle it. The authors promote the use and improvement of analysis tools, security-aware compilers, and constant-time cryptographic libraries, which all have the aim of making writing constant-time code easier.

End-to-End Encrypted Email. The seminal paper “Why Johnny can’t encrypt” [45] spawned a string of research on the usable security and adoption of email encryption. Even though the problem of ensuring end-to-end-secured email communications seemed to be solved on a technical level by existing implementations of both PGP and S/MIME, the bad usability of these solutions was a major blocker for adoption. In 2019, Ruoti et al. summarize these almost 20 years of “Johnny”-papers on the usability of secure email communication [39]. They recommend tight integration of security tools with users’ existing ways of communication, context-sensitive tutorials, and trustworthy design and call for more research into the – to this day – unsolved challenge of usable secure key management for private end users. Stransky et al. analyzed 27 years of email data from a large university and found that only 5.46% of users used S/MIME or PGP, resulting in 0.06% encrypted and 2.8% signed emails [42]. The research reports that key management issues and the use of multiple email clients negatively impact encryption adoption.

Clark et al. investigated email encryption by identifying stakeholders in the current ecosystem of email communication. They infer that the current, less-than-ideal state of end-to-end encryption in email communication stems from the evolution of fragmented secure email solutions created by industry, academia, and independent developers [13]. There are now competing solutions that address different interests of seven stakeholders they identify as Email Service Providers, Enterprise Organisations, Privacy Enthusiasts, Vulnerable Users, Secure Mailbox Providers, and Typical Users. While some enterprise organizations are legally required to be capable of exceptional access to their employees’ communications, such exceptional access would be considered a plain backdoor by privacy enthusiasts and vulnerable users. On the other hand, typical users might highly prefer systems that allow server-side processing of their emails by Email Service Providers for, e.g., spam filtering or for reliably pre-sorting email messages into categories.

Standardization Processes. Paterson and van der Merwe present how, for the development of TLS 1.3, the IETF TLS working group was able to move from a *design-release-break-patch* cycle to a *design-break-fix-release* cycle [34]. This means that they were able to involve academia and industry heavily during the design stage of the protocol, instead of

relying on the historical way of releasing a standard and then releasing patches after vulnerabilities are found through analysis by academics or through use in the real world. They state that better tools and greater academic community engagement enabled this move. They postulate that a requirements *analysis-design-prove-release* process might have been even better. In 2020, Halpin et al. investigated why attempts to update the OpenPGP standard to a modern security level have failed at the IETF [22]. They find the core reason to be a missing simple AEAD interface, which in turn requires a decentralized public key infrastructure – that does currently not exist.

Also in 2020, Bernstein surveys standardization procedures of past cryptography competitions, and finds performance pressures and limited time for security analysis as sources of security risks, but also possible NSA interference and incentives in academic publishing [7].

3 Methodology

We conducted semi-structured interviews with 21 cryptography experts from academia, industry, and nonprofit and governmental organizations. Our aim was to elicit their personal experiences and reflections on the process of bringing cryptography research output “from paper into practice”—including potentially contentious ones—and to be able to ask follow-up questions. All interviewees had at least ten years of experience researching, designing, standardizing, or implementing cryptography and had high standing and visibility in the cryptography community.

This section provides an overview of our methodology, describing the process of developing the semi-structured interview guide, participant recruitment, interview procedure, the qualitative coding process, and ethical considerations and limitations in this section.

Initial Recruitment and Instrument Development. We initially recruited three participants from our professional networks—researchers with a strong publication record in cryptography and security who had created a cryptographic protocol or application that is widely used today—to scope the problem space. These were in-depth (80–110 min) interviewee-led interviews on their involvement in and experiences with a successful deployment, centered around the question: *What were the obstacles (blockers) they encountered on the adoption path? What did they have to do to overcome them?* We shared the transcripts of those interviews with the interviewees, and asked a number of clarification and follow-up questions. Based on the interviews and their answers, we developed an interview guide for the remaining expert interviews.

Our semi-structured interviews were interviewee-led: Interviews centered around questions that participants answered in-depth, as well as topics and questions they thought would

add to our line of inquiry, which both led to deep insights and helped prioritize and value participants' time. Since we were looking for insights beyond published literature, we asked for their personal assessment and opinions of the causes of problems. Our first three participants had no reservations identifying them, but also qualified that not everyone would agree. We therefore selected such "strong statement" quotes and discussed them at the end of subsequent interviews (see Section 3.1). Expert-led interviews and "feeding forward" quotes for responses by other participants was inspired by the Delphi method [4, 38].

Further Recruitment. In the three initial interviews, interviewees mentioned other academic cryptography researchers, industry-based researchers, or policy experts that would be helpful for investigating the different challenges for effective cryptography.

We wanted to recruit an experienced set of cryptography experts from different parts of the cryptography ecosystem: academic researchers, industry researchers and practitioners, and those working in nonprofit and governmental organisations, and who had been involved with more than one aspect of cryptography research or implementation: experiences ranging from cryptography theory to standardisation to implementing cryptography for expert and non-expert users, and policy work. While we focused on recruiting experts with past successes in getting cryptography adopted, we also aimed for multidisciplinary interviewees to cover a broad set of expertise, experiences, and opinions.

From the recommendations, we invited interviewees who met the eligibility criteria, starting with those mentioned multiple times. We issued 20 invitations in total. 18 accepted; 2 invitees did not respond. See Table 1 for a summary of interviewees' backgrounds and experience. In total, we recruited three interviewees from our professional network and 18 interviewees through snowballing; however, many of the experts suggested by interviewees were also known to us prior to the interviews.

Interview Procedure. All interviews were conducted by one interviewer and one to two additional interviewers, except for P21, where no additional interviewers could be present. The lead interviewer could fully focus on asking questions. The additional interviewers ensured that no questions were left out, could ask follow-up questions that emerged, or take over in case of internet connection issues. We conducted all interviews remotely using a self-hosted Big Blue Button Instance, Zoom, or Google Meet, depending on the interviewees' preferences. We expected the interviews to take 45–90 minutes and scheduled one or two-hour appointments with all interviewees, yielding about 30 hours of interview data. Interviews lasted between 45 and 155 minutes; the median duration was 84 minutes. We based the interviews around non-leading, open questions, allowing interviewees to elaborate on their thoughts and answers.

Positionality. The researchers who carried out the interviews and analysis had a range of disciplinary backgrounds, including psychology, cryptography, and computer science, and a broad range of academic research experience. We made all decisions (study design, interview guide, and results reporting) with reference to published research practices. Nevertheless, decisions on which particular lines of inquiry to pursue, or what to include in the results, are likely to be influenced by the perspectives present and the dynamics between the researchers. Other researchers analyzing our data may have focused on different aspects (though we are confident that on the aspects we are reporting on, they would not have come to radically different results).

3.1 Interview Guide

We describe the semi-structured interviews below and in Figure 1.

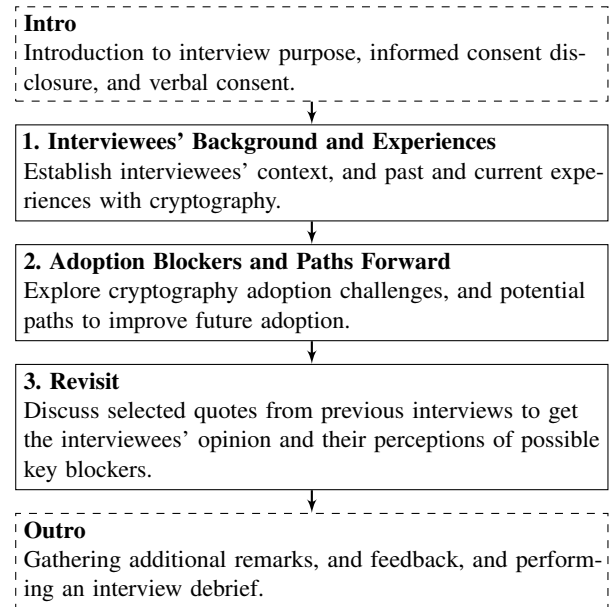


Figure 1: Overview of the interview flow and topics. We followed up the introduction of each section with specific questions (if not already covered). The semi-structured interviews allowed interviewees to diverge from this flow at any time.

We structured the interviews into five main sections, including an introduction, questions about the participant's background and experiences with cryptography, their viewpoints on cryptography adoption blockers and potential paths forward, a discussion of quotes from previous interviews, and a debrief. Each section included opening questions and corresponding follow-up questions and prompts. The full interview guide can be found in Appendix A.1.

At the beginning of each interview, we introduced ourselves, the research project and its goals and provided details

Table 1: Detailed overview of the cryptography experts we interviewed.

Alias	Duration	YoE ¹	Background ²				Experience in Cryptography ²				
	h:mm	yrs	Academic	Industry	Non-Profit	Government	Design & Analysis	Standards	Impl. for Devs	Impl. for End Users	Policy
P1	1:37	15+	●	○	●		●			●	
P2	1:48	25+	●				●	●	●	●	●
P3	1:24	30+	●		●		●	●			
P4	1:22	30+		●	●		●				
P5	1:55	15+	●				●	●	●		
P6	1:08	25+	●	○			●		●		
P7	0:48	30+	●	○	●		●		●		
P8	1:14	15+	●	○	●		●			●	
P9	1:13	20+		●			●	●	●		●
P10	1:37	20+	●	○		●		●			●
P11	1:33	25+	●	○	●		●		●	●	
P12	1:24	30+	○	○	●	●	●				
P13	1:30	25+		●	●	●	●				
P14	1:11	10+			●	●					●
P15	1:22	20+	●	○			●	●			
P16	1:19	20+	●				●		●		●
P17	1:30	20+	●	○			●	●			
P18	1:18	15+	●				●		●	●	
P19	1:28	10+	●			●	●			●	
P20	0:57	30+	○	○	●		●	●			
P21	1:34	30+	●	○	●		●	●			●

¹ Years of experience in relevant fields. ² Based on self-reported information and internet research.
 Column Academic: ●: 7+ years post-PhD in a tenure-track academic position. ○: 1–6 years post-PhD teaching at University
 Column Industry: ●: 7+ years employment in major tech firms. ○: 1–6 years employment and/or part-time role, e.g. co-founding a start-up.

of the interview process. In particular, we explained that interview participation was voluntary, that interviewees could skip any question and that we would not judge their answers. We guaranteed full de-identification of quotes we might use in a publication and offered them to send a preprint of the paper before publication.

Interviewees could ask questions and we asked them to provide additional verbal consent for interview recording and data processing. We started the recording and began the interviews with the structure below:

Background and Experiences. We first asked interviewees to describe the work or research they primarily do and characterize the cryptography field they work in. This section aimed to connect with the interviewees, ease potentially nervous participants, and establish an initial context about their work.

Adoption Blockers and Paths Forward. We asked interviewees to report on their experiences with blockers for the adoption of cryptography. Specifically, we were interested in experiences related to their work. If available, we asked them to describe some of the work they had hoped would be adopted but which did not meet their expectations. In the second half of this section, we asked them for the most important factors or steps that led to the adoption of cryptography. At the end of this section, we were interested in our interviewees’ views on things that have to improve or change for a better adoption of cryptography.

Revisit. In the final section, we showed participants quotes from previous interviews. We selected quotes based on their significance to be an adoption blocker, to provoke discussion,

or to judge expert agreement about a blocker better. For each interview, we selected 5–10 quotes that we knew the respective interviewee to have expertise on. In total, 13 unique quotes were shown across all expert interviews. A list of all discussion quotes is provided in the extended version of this paper.

We did not provide financial compensation to our participants, which, based on our experience with similar expert populations, are often declined; participants supported our research out of the intrinsic motivation to improve the adoption of cryptography.

3.2 Data Analysis

We recorded the audio of all interviews, removed identifying information from the recordings, transcribed them internally and reviewed the transcripts for potential mistakes. We used an iterative semi-open coding approach [12, 15, 43] to perform thematic analysis [14] for all interview transcripts. We stopped interviewing after 21 interviewees, reaching saturation [17].

The main author conducted open coding to develop an initial codebook on all interview transcripts. We additionally followed a deductive approach to code all areas in the cryptographic landscape that participants discussed. In a second coding step, the main author coded all blockers for adopting cryptography using inductive coding. Another co-author independently coded all transcripts. The main author and two other co-authors reviewed both codings, discussed im-

portant insights and identified core topics in multiple affinity diagramming [9] sessions on a whiteboard.

Our approach does not require the reporting of inter-coder agreement. We resolved each conflict when it emerged following established practices in the HCI community [31].

3.3 Ethics and Data Protection

Our institutions did not require formal ethics approval for this type of study. However, we modeled the interview study after the ethical principles for human subjects research involving information and communication technologies outlined in the Menlo report [27]. 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 GDPR. We provided all interviewees with information about the study procedure and data handling before signing 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 sent interviewees a preprint of this paper before publication so they could request changes or correct misunderstandings; following their feedback, we made small changes to Table 1. We did not offer any compensation to our interviewees as they were all highly successful individuals motivated to work on cryptography to make the digital world safer. Personally identifiable information was stored securely and encrypted at rest and in transit, compliant with the GDPR. We removed parts of interviews participants flagged as too sensitive to transcribe and de-identified participants using the identifiers P01–P21 and removed any information that would easily identify our participants from the transcripts. After checking transcripts for correctness, we deleted all audio recordings.

3.4 Limitations

Our research is affected by limitations common to interview studies, including potential over- and under-reporting, self-reporting, recall, social-desirability biases, sampling bias, and limited generalizability [30]. While our sample is a convenience sample that may not represent all cryptography experts, we tried to account for the above biases by interviewing a diverse sample of cryptography experts fitting our recruitment criteria. Our sample includes various academic and industry settings, from leading experts in foundational cryptography research to industry-leading developers of cryptographic products. The interview data imply that our sample is broad and diverse. However, we refrain from making quantitative statements due to the qualitative nature of our research methodology. We conducted 20 interviews in English and one in German.

4 The Path of Cryptography Adoption

In answer to RQ1 “*What steps are involved in bringing cryptography from papers to products, and who are the relevant stakeholders?*”, our participants referred to stakeholders and processes that are part of what we have called “the cryptography adoption path” which is embedded in the cryptography ecosystem. Figure 2 shows a map of that ecosystem, containing entities (actors), activities, and artifacts (products) in the current cryptography ecosystem. Our goal was to create a *simplified* map to help us to make sense of—and to argue about—the cryptography ecosystem as a whole. We are certain that for specific end products, the map can be extended with additional relevant actors and artifacts.

This map helped us to structure the actors and processes involved in bringing cryptography from papers to products. We grouped different stakeholders involved into entities that perform different roles or jobs in the cryptography ecosystem, and which may be performed by a single person or groups of persons. Both the actor entities and the processes were explicitly mentioned by interviewees 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 (see Section 7)

Our results reporting follows the path of bringing cryptography theory from papers to end-user products, as supported by our interview data.

Disclaimer. While the cryptography adoption path we illustrate in the map emerged from our interviews, it is a simplification and specific cases might deviate from it.

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 (simplified) 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 hard 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.

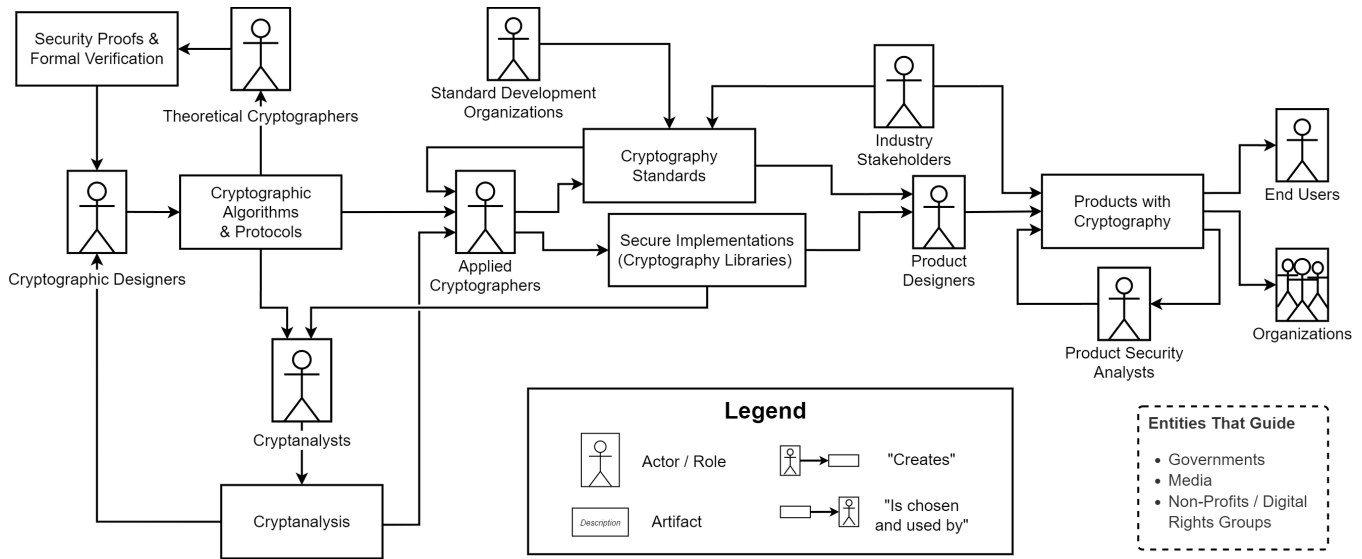


Figure 2: A visualization of the cryptography ecosystem and the path of bringing cryptography from paper to product. This map emerged from our interviews and serves as an abstracted illustration of the complex and ever-changing dynamics of cryptography adoption.

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. Standard Development Organizations (SDOs) 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 IETF, or in the form of competitions, as often organized by NIST for cryptographic primitives. Product vendors can choose to put resources into standards development for one of mainly three reasons: They want to be able to interoperate with their or competitor’s systems, implement a standard for compliance reasons, or use it as an argument in marketing their product.

Secure Implementation. Cryptographic research papers partially provide proof-of-concept implementations; according to our participants, these are generally not robust enough to be used in a product. Applied cryptographers take the cryptographic research output and turn it into implementations that can be widely used by non-cryptography experts, e.g., 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 feed back the results to improve robustness.

Product Development. Hardware or software vendors and developers may want to protect the data their product is handling. To do that, they 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 path of cryptography adoption path are everyday end users and organizations, who choose and use specific cryptographic products. They can choose products that are inherently insecure, or use products in unintended ways and thus potentially risk their privacy or security.

Entities That Guide. The following additional entities are noteworthy. They can influence cryptography adoption in multiple ways.

Governments can create legislation and regulations that impact the funding of research and critical internet infrastructure, impact 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.

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*¹ and the *Let’s Encrypt*² initiative were driving forces for widespread adoption of HTTPS.

¹<https://certbot.eff.org/>
²<https://letsencrypt.org/>

5 Challenges of Cryptography Adoption

Our interviews revealed a myriad of challenges that can occur when bringing cryptography from research papers to products. We report on these in the categories 1) misaligned or conflicting incentives in academia, 2) challenges in standardization and 3) challenges in reference implementations, 4) communication gaps and unclear responsibilities, and 5) usability issues, and provide examples from our interviews.

5.1 Misaligned Incentives in Academia

Incentives hold significant sway over the trajectory of human endeavors, both in their presence and absence. When examining the path of cryptography adoption, a central aspect are actors' incentives, as they often do not align perfectly with our overarching goal of secure cryptography adoption.

A frequent pattern we identified is many actors not directly working towards the overarching goal of cryptography adoption because they are not directly incentivized. Additionally, some interviewees illustrated cases where actors partaking in the ecosystem have conflicting incentives and thus actively work against the adoption of secure cryptography.

Many pointed out that the creation and maintenance of production-quality cryptographic code, i.e., code that can be safely used by software developers, is not a primary task of cryptography researchers.

Research is mostly funded through grants. Most interviewees agreed that for many grants, one of the most important metrics is the number of published top-tier research papers. For many researchers, there is not much to gain from putting rigorous amounts of work into usable and secure implementations of their cryptography: *“practical impact does not get you papers at CRYPTO”* (P9). Some interviewees mentioned that some cryptography conferences are starting to reward the “engineering side” of research more. Multiple participants highlighted the effort of the Real World Crypto (RWC) Symposium in this direction. *“Real World Crypto is actually a wonderful place where industry and academia come together. [...] The community is growing and a lot of papers that analyse a crypto standard will now actually appear at the security conferences.”* (P3)

But few participants also argued a different way, implying that RWC is not as inclusive as some would want it to be: *“RWC, even by it's name, it conveys what the message is: 'Don't bring your theoretical nonsense here. We don't want to hear about it!'”* (P13). Either way, our participants agree that when designing new cryptography, conferences at most reward proof-of-concept implementations, which may sometimes be labeled as ‘reference implementations’, but are far from production-ready code. We identify this as a major challenge for cryptography adoption, since almost all our interviewees underlined the importance of providing usable code when working towards adoption of new cryptography.

“If you really want people to use it, you have to have code that they can use. To play with. And it has to replicate more than what just the paper does. I think that is what academics don't do enough. It has to be a piece of code that is genuine enough that it can be used to do stuff, that is not just what you have said in the paper, but more general use.”
— P16

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, and thus offers an opportunity for ample amount of cryptography research. [6]

One participant mentioned a controversial incentive that comes with the concept of boring cryptography: A community incentive for cryptographers themselves to not 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 we persuaded researchers working on “interesting cryptography” to pay more attention to “boring cryptography”, most 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.”* (P9)

Some participants shared sentiments that can be summarized as a statement that academia should maintain its ability to explore uncharted territories that may be detached from immediate real-world applications. A participant with industry background noted that it is unreasonable to expect academics, who are not directly compensated by industry, to undertake tasks that primarily benefit industry. *“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)

Academics possess a diverse skill set that extends beyond their narrow focus; they lack incentives to engage in various activities critical to the overall cryptographic ecosystem. Some participants noted that contributing to standardization efforts and attending standard meetings are essential tasks, but they often receive limited recognition and reward.

“Standardization [...] is very, very painful. There are people who do that, I mentioned some names in the IETF, they are doing it because they feel it's a socially important thing... out of some kind of duty. Social duty. But it's difficult to find those folks.”
— P15

Another participant expanded on this, explaining that *“people won't get that much academic credit for spending like three years flying to IETF meetings and fine-tuning a standard. So, maybe there's an academic incentive problem.”* (P8)

However, some interviewees mentioned that incentives for academics to partake in standardization efforts appear to be growing: “*I would have agreed with this more five years ago, so I think it’s going in the right direction. Ten years ago 99% of the community were not interested in standards at all. And this is getting better.*” (P17)

5.2 Standardization

Some interviewees with industry backgrounds mentioned that the IETF standard drafting process has changed over the past 20 years, becoming more cumbersome: “*In practice the IETF-process doesn’t work as well as it used to. [...] It discourages participation. We did end up trying to play the IETF-game for [widely adopted cryptographic protocol]. But it was so painful, we all stopped doing it before it was standardized.*” (P4)

Some interviewees suggested to be critical of standardization process participants: “*You’d think the goal is to bring out efficient and secure standards, but of course, companies are also there to protect their own interests. And we know from the Snowden documents that governments send employees there to boycott the process or to make sure that standards are not usable. [...] I think you should understand the game and you should play the game, and you should accept the outcome could never be perfect. [...] It’s about everybody protecting their interests and minimizing their losses.*” (P7)

Participants provided examples including the financial industry’s resistance to forward secrecy in TLS 1.3, where, near the end of the TLS1.3 drafting process, BITS [10], an organization that at that time represented about 100 of the top 150 US-based financial services, fought to re-add RSA static keys to the standard. Security experts had removed static RSA keys from the draft because static RSA keys break forward secrecy and are perceived as a potential backdoor. BITS argued the need for static RSA keys to be compliant and improve security by monitoring the content of TLS connections cheaply [40].

The above example illustrates that, even in open and transparent cryptography standardization processes, the actors’ motivations might not support the most secure outcome possible. They can be about reducing engineering effort and minimizing the cost of updating systems to accommodate new protocols:

“I have no idea what to do about the unfortunate influence of industry [in cryptography standards], like: ‘We have this system from twenty years ago that is deployed. Our interest is to make sure that it continues to be officially sanctioned.’ – It is tough.”
— P9

Different stakeholders with diverging goals. Different actors in standards processes have different goals and incentives.

Academic Researchers According to our interviewees, most researchers will struggle to find funding for putting

work into standardization efforts. Standardization work was described as ‘tiring’, tedious’, ‘deathly boring’, and ‘unrewarding’ work, that demands long-term commitment—which requires a magnitude of resources that does not fit well into common funding structures in academic research. Our interviewees’ collective sentiment towards open standardization processes, like the IETF’s Internet Standard Drafting Process [11] and open cryptography competitions run by NIST [7], was positive when compared to more closed standardization processes, like those of the International Organization for Standardization (ISO) or the European Telecommunications Standards Institute (ETSI). At the same time, some interviewees said that NIST processes could be improved, questioning the competition requirements set by the organizers, or being concerned about NIST’s collaboration with the NSA.

Industry Stakeholders Some interviewees argued that companies are missing incentives to put resources towards standard development. When creating new standards, companies generally do not gain much from putting more resources into standard development than the bare minimum needed to make it work for themselves. Any resources they spend on making the standard specs and documentation particularly usable (see Section 5.5) reduces the extent of resources their competitors have to invest to implement the standard.

Additionally, our interviewees reported on industry players pushing to have their self-developed protocols standardized for financial gain:

“An ISO working group was trying to standardise some smart card based authentication protocol, and some [nationality redacted] guy had invented [a fitting protocol]. He had managed to get it turned into a national standard. And now he was a one-man consulting company and he was trying to get it turned into an ISO standard. And it was complete junk, we broke the protocol, in several different ways. But he still was really, really pushing for this thing to be standardized because, you know, most didn’t know what was going on.” — P3

Misunderstandings in Standardization. A particular challenge for academics and engineers in standardization processes is that, since they have different backgrounds, they have different terminologies, expectations, and ways of looking at things. As one participant highlighted, “*There are a lot of cases where the same words are used to mean different things and different words used to mean the same things*” (P12), making effective communication difficult. Another interviewee cautioned that “*There’s not many people sitting in the middle to translate between these worlds. I think there’s like half a dozen people who are able to do that*” (P3). And while most participants agreed, some added the correction that there are more people who are able to do that—but don’t do it due to conflicting incentives:

“There are other people who can do it—but don’t do it. [...] They get highly paid to work inside Microsoft or inside Amazon. And they probably look at what happens in the standards world to make sure it’s not too crazy, but they don’t want to spend too much time. It’s very time-consuming, being that person.” — P15

(see Section 5.4 “Communication Gaps and Unclear Responsibilities”). The absence of these “translators” can lead to misunderstandings and miscommunication, and result in unsuccessful standards or even security flaws. One of our interviewees reported interactions with the European Telecommunications Standards Institute (ETSI), where academics have not been involved in an early drafting process of a standard as *“I was at some ETSI meetings and [...] we were sitting there with 2 or 3 academics in the back row and our job was just to call ‘bullshit!’ from time to time when things really clearly went in the wrong direction”* (P5)

According to (P4), *“the people who are involved in standards are inevitably becoming more theory-based”*, suggesting that a better understanding of theoretical cryptography is growing among standard organizations. The same participant went on to remark that *“the people who are involved in actually implementing things are getting less involved in standards because they just can’t be bothered with the process”* (P4). This illustrates the division between those involved in setting standards and those implementing them.

Adoption of Standards. Some interviewees reported that many cryptography standards are not adopted because they fail to identify their users and use cases in a meaningful way. Several interviewees agreed that standards go wrong if they misunderstand real-world use cases: *“Often, the standardization process is done in a vacuum, where they don’t talk to the potential customers in any meaningful way [...] and 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). Our interviewees suggest that this tends to happen more often in closed standardization processes, like those of ISO or ETSI, rather than in open standardization processes.

5.3 Reference Implementations

Our interviewees mentioned the important role of reference implementations. Reference implementations occur on different levels: 1) reference implementations of research ideas as part of a research paper; 2) reference implementations of standards; 3) reference implementations for the use of cryptographic library APIs.

Reference implementations provide practical examples that support developers implementing cryptography. They furthermore allow developers to verify the interoperability of their implemented product.

Our participants expressed consensus on the variability of reference implementation quality. *“My reference implementations were never production-ready, they were reference implementations. But this is a problem with code libraries, [...] people just grab them and use them without knowing what they are.”* (P20)

One participant stated, *“reference implementations have a very big role to play in making sure interoperability is there.”* (P17) A distinction was made between reference implementations and optimized implementations. *“You have your reference implementation which is written in the plain scene, no tricks, no optimizations. And then you have the ‘weird and whacky’ implementation, which gives me my performances.”* (P16) However, there is a tendency for optimized implementations to be considered reference implementations, making them challenging to work with for those lacking specific additional skills. As one participant pointed out, *“There are companies that are selling very large numbers of different [IoT] objects. They’re just using reference implementation code off the shelf, putting it into devices, and then we discover a huge number of vulnerabilities. And they just don’t have the skill-set in their engineers to check whether the reference that they’re using is sound.”* (P15). Other examples of misuse of reference implementations provided by some participants are the copying of hard-coded credentials and insecure configurations.

Additionally, low-quality reference implementations can hinder the adoption of cryptography by rendering them simply unusable. One participant shared their experience with software written by academics:

“There is no data abstraction. There is no separation of interface from implementation. Half of them, we looked at the code and we go: ‘We want to include that in our library.’ Then we read it and we go: ‘Eww, no!—We’ll write it ourselves.’ It might be super fast, but it is super fast because it breaks every rule under the sun.” — P16

To address these challenges, some participants made various suggestions. One suggestion involved labeling proof of concepts to distinguish them from production-ready code. However, many participants were quick to point out the very limited effects that this labeling has: Developers may not fully comprehend or pay attention to the potential risks of copying code. As one participant stated about a university class assignment, where students had to implement an encrypted chat protocol:

“People don’t really realise the danger of what they’re putting out there. I developed a class assignment with my students, where they implemented an encrypted chat protocol. The code is on GitHub, so I tried to put a warning on the front page of the repo: ‘Don’t actually use this for anything! This is

just a classroom example.’—But I’ve thought about that a lot. Any crypto code you put on the web might be used by somebody. A lot of people put stuff up just as a pedagogical example or proof of concept. And then it ends up getting used in real products.” — P8

In the context of the NIST Post Quantum competition one participant made a suggestion to reduce these risks: *“The reference implementations of post-quantum are of variable quality. That is a worry. It means the implementation ecosystem is not seeded with the very highest quality things. [...] I think the answer is that [...] we need to flood the space with high-quality implementations. And if that happens, we will be ok. And if it doesn’t, it will be more problematic.”* (P9) By providing high-quality reference implementations, the chances of developers picking up the bad reference implementations are decreased.

Some interviewees pointed out that the skills required for creating high-quality reference implementations differ from those needed to write academic papers. *“The quality is very variable because the skill set of making an implementation is different from the skill set of writing a paper that is accepted at [a crypto conference].”* (P11)

5.4 Communication Gaps and Unclear Responsibilities

Many interviewees reported misunderstandings and miscommunication and unclear responsibilities between the different stakeholders involved in the cryptography ecosystem.

As mentioned by some participants the shortage of “translators” poses a significant challenge in bringing cryptography from research papers to products and can lead to output not being implemented correctly, or not being implemented at all: *“[Engineers] have a system and they want to make it secure. And so you indeed have to translate your scheme and explain them what you want to do, what you want to achieve and why these properties are important.”* (P7)

Cryptography research is evolving rapidly, spans many specialized subjects, and uses highly diverse terminology and different language. For example, experts in symmetric and asymmetric cryptography already speak a very different language.

It is not unusual for theoreticians to view applied aspects of their craft as too difficult, messy, and complicated: One of our interviewees reported a theoretical researcher telling them *“No! I don’t want to understand the problem with the application. That’s your job! My job is just the design and mathematics!”* (P10)

Fortunately, there is a growing recognition of the value of collaboration and cross-disciplinary learning, despite the inherent difficulties or human cognitive constraints that come with working beyond one area of specialization.

“I think there is now a) appreciation from both sides and these people are talking. This is fairly new. And b) I think that there is also even an appreciation in the crypto conferences now.” (P17)

By breaking down barriers and engaging in effective communication, cryptographers can gain a better understanding of the ideas and motivations of their colleagues from different backgrounds leading to more productive and collaborative research.

5.5 Usability Challenges

We frequently spotted usability issues when interviewees reported on challenges to secure the adoption of cryptography. This includes the usability of end-user products, but also the usability of cryptography libraries, reference implementations, standard specifications, and documentation.

Some interviewees reported usability issues (in addition to low quality) with reference implementations. Reference implementations can be difficult to understand because they might include tricks that are unfamiliar to developers who are not well-versed in the particular hardware or software being used. *“[...] One trick, for example, is if you got anything that relies on floating point and people want to speed it up, they change it to fixed point. And there are people who don’t know about fixed point arithmetic.”* (P15) Helpful reference implementations should prioritize code readability over performance optimizations and adhere to established coding practices.

This is further underlined by an interviewee describing the implementations of new cryptography proposals in a sub-community of cryptography research, *“Academics, when they write papers, care about speed too much. So they remove everything that does safety: All those array boundary checkings, all that making sure stuff interfaces correctly together and decodes [...] Their paper is sold on how fast it is. Therefore they only care for how fast it is. And therefore, all the stuff you need in a proper system is thrown out of the window.”* (P16), referring to safe programming, documentation, and good engineering practices.

Some participants reported usability issues with standard specifications in terms of structure and readability. Our interviewees expressed concerns about the current state of specifications: *“Having specs be large, complicated, and reference implementations that can not be executed is, I think, an issue.”* (P9) The same interviewee mentioned that the expectations of the level of understanding of the specifications being too high:

“The contributors are all experts, who know what the interpretations of this wording are. They are going to be fine. The spec only needs to be roughly approximate for them because they, basically, have an understanding in their head already. Whereas for

the world as a whole, maybe you want something like UX researchers?” — P9

Many participants commented on the low usability of non-established cryptographic libraries and tools. They agreed that for new ideas to get adopted, their creators have to provide working and usable code that others can play around with. Our participants reported on usability issues they encountered when setting up such tools:

“I at some point tried to install Project Everest. And [researcher] told me “Here’s a script! You download that script. You run it. And either that just installs everything and everything works, or you’re in big, big trouble.” I was in big, big trouble so I spent a whole weekend actually not installing Project Everest. I didn’t manage in a whole weekend.” — P5

Our participants agree that, for end users, usability is paramount for adoption. *“Identity is another really good example where there is lots of good cryptography to help [...], but almost none deployed. [...] the UX is the hard bit.”* (P4). Our participants explain, the adoption of cryptography, such as end-to-end encryption in email, may impede users’ desired functionality, leading to justified resistance.

“Deploying crypto takes away features that people really want. It takes away easy recovery, backups, text search or fancy AI features that people want to do. I’ve seen that especially with encrypted mail. People don’t want encryption, they want easy backups and search over their mail history.” — P8

Other interviewees agreed that to get cryptography more widely adopted by end users, it is imperative to make sure the proposed system does not worsen the user experience by decreasing performance or sacrificing features in the name of security. Our interviewees note that there is existing research towards implementing features like server-side spam filtering in privacy preserving ways, like in the realm of homomorphic encryption, but that more research is needed to arrive at acceptable solutions: *“There are all kinds of fancy cryptographic proposals for being able to do spam-filtering without being able to see plaintext. They are not deployed. They are not totally reasonable at this time.”* (P11).

6 Discussion

We discuss our results in the context of our research questions and make recommendations to improve the process of bringing cryptography from research paper to product.

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 (c.f. Section 4). One important insight is that research papers do not make it into practice not because cryptographers do not care, but because they are assessed by how many top-tier scientific articles they produce. Phil Rogaway [37] argued that cryptography researchers should focus more on real-world impact, but 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.

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.

There is also a shortage of actors who understand cryptographic research output in sufficient depth to anticipate how it will work at the system level and few cryptography researchers who understand the practices and pitfalls of secure software engineering at a level that would allow them to produce secure and usable crypto libraries by themselves. Due to a shortage of double experts, one way to increase adoption is by further encouraging cross-disciplinary engagement between cryptographers and expert software designers. We find implementers who are either missing the skills or incentives to put the needed amount of rigor, effort, and time into secure cryptography implementations.

Developers are pressured by impending deadlines and are equipped with usually only a little knowledge about security and cryptography. They need to be supported with high-quality crypto APIs to prevent security flaws. We find end users caring about many things, but cryptography is seldom one of them, especially when it introduces friction to their everyday life or takes features away: *“In general users don’t care very much: I mean good cryptography is cryptography that users don’t see, right?”* (P7).

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 (see Section 5.4). 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 as 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” (see Section 5.4).

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.

On the other hand, we saw that the cryptography ecosystem is evolving and most of the interviewees reflect that. To give an example of work that we were not aware of at the time of this study, in 2022, Kannwischer et al. [26] created the PQClean project to improve the quality of reference implementations submitted to the NIST Post Quantum Cryptography competition. They find that properly implemented, a set of guidelines together with a testing framework can increase code quality, reduce efforts, and thus benefit submitters, the community, and the standardization body itself.

6.1 Recommendations

Based on our results we provide recommendations for the academic community, industry and standardization organizations. Table 2 provides an overview of the main challenges and our recommendations in this paper.

Academic community. Section 5.1 describes 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 aim to support the adoption of secure cryptography, they might consider funding and rewarding the participation of academics in standardization bodies or secure implementation of cryptography.

There is a lot to gain from funding academics to attend non-free standardization meetings, and there are examples

of this funding being beneficial: “*Within my grant there’s 10000€ for ETSI membership. What’s interesting is, that puts in the whole of [university name] as an ETSI member. (...) So now anyone can go to their meetings and call bullshit. So I think it’s something that maybe more universities should be doing and then have just academics go there and follow what’s going on and see how these standards are being made. It is a very interesting procedure. You learn stuff when you are at these meetings.*” (P5)

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 real-world use cases and requirements.

Our results imply communication challenges between 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 between cryptography experts and engineers and recommend developing communication skills that support cross-disciplinary collaboration.

When building reference implementations, cryptographers should mind the difference between proof-of-concept implementations that support results in academic papers, and ones that support re-use by software engineers. We recommend cryptographers clearly mark proof-of-concept implementations on the 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 reach out to experts from the field of software engineering when implementing protocols and algorithms with the aim of improved implementation quality.

Industry. Our interviews imply that wide cryptography adoption is hindered when 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, or Google opting to encrypt all data in transit after the Snowden revelations. Additionally, we encourage companies and organizations to consider investing in core infrastructure maintenance projects like the Open Source Security Foundation (OSSF). We encourage implementers and users of standards to reach out upstream, communicating problems, and needs, or actively contribute

Table 2: Summary of identified challenges and recommendations.

Challenge	Recommendations
Academic Incentives	
Missing incentives to create usable reference implementations	Conferences could try to not only judge contributions by their performance and theoretical security levels, but also by usability.
Missing incentives to participate in Standardization efforts	Standardization Organizations could improve outreach towards academia. Grant givers could reward standards participation more explicitly. Raise community awareness of the importance of standards work.
Missing incentives to create production ready code	Promote collaboration between cryptography experts and software engineering experts.
Standardization	
Misunderstandings between academics and engineers. Things get lost in translation	Standards could be UX tested. Have translators between theoretical researchers and engineers. Make standards machine readable or executable.
Different actors with different interests participate in standardization efforts	Open Standardization process; transparent competition requirements; raising awareness of possible ulterior motives; proper requirement analysis in the beginning.
Fail to identify a standard's users and use cases.	Standardization organisation's working groups need to do proper requirement analysis in the beginning of drafting a standard.
Reference Implementations	
Low quality of reference implementations	Support cross-disciplinary outreach between cryptography experts and software engineers; In competitions, provide practical support via ,e.g., CI, automatic testing, . . .
Proofs of Concept mislabeled as "Reference Implementation"	Make sure that PoCs are not mislabelled. Conferences might set formal rules for what may be submitted as a "reference implementation".
Misunderstandings	
Communication challenges	Try to establish common terminology and language to help cryptographers communicate better. Promote collaboration between cryptography experts and software engineering experts.
Usability	
Low usability of Reference Implementations	UX testing, UX research, support cross-disciplinary outreach between cryptography experts and software engineers, In competitions, provide practical support via e.g. CI, automatic testing, . . .
Low usability of Standards	UX testing, UX research, have machine redable and executable standards with proper documentation.
Low usability of Crypto Libraries	UX testing, UX research. Grant givers can allocate resources for maintaining infrastructure code.
Low usability of End-User Products	UX testing, UX research. Engage in theoretical research on unsolved root causes like key management.

to open standard development, instead of developing and standardizing cryptographic solutions behind closed doors.

Standardization Organizations. Our results imply issues with the complexity, readability, and actionability of standards. Our experts recommend working towards machine-readable and possibly executable, standard specifications, to support automated security proofs. The inspiration can come from unit tests—made for standards. Multiple interviewees suggested the idea or expressed interest in this area, they also

see the way forward in automating code audits, and assigning properties to cryptographic functions. After developers specify what security needs they have, there should be an analysis of how the developer “plumbed” the crypto functions together. Some interviewees criticized the proprietary nature of some standardization organizations. For example, ISO standards are expensive and exclude interested cryptanalysts and developers. We recommend the cryptography community focuses on open standards and standardization organizations to be

more inclusive and support easy access to their specifications. Standardization organizations should make sure they are trustworthy by further engaging with the academic community, emphasizing open communication, and open competitions. A “seal of approval” of a trusted standardization organization is a major driver for cryptography adoption.

7 Conclusion

We investigated key challenges in bringing cryptography research from papers to cryptographic products. Therefore, we conducted 21 semi-structured interviews with cryptography experts with high visibility and standing in the community. Based on our interviews we developed a map of the cryptography ecosystem and illustrated involved actors and stakeholders. The map serves as the foundation to highlight challenges and pain points in the ecosystem and report and discuss further results. We identify five major challenges that hinder the adoption of cryptography and provide opportunities to improve future adoption. Misaligned or conflicting incentives of actors in the ecosystem are a big challenge. For example, cryptography researchers are not incentivized for implementations beyond research publications and may not have strong motivations to write cryptographic code that goes beyond proof-of-concept implementations. We find that cryptography standardization is challenging and that some of the current standardization processes do not support the secure and wide adoption of cryptography. For example, involved stakeholders might have diverging goals, and standard specifications tend to be hard to read or incomplete. Reference implementations are crucial for the adoption of cryptography and often contribute to their limited or insecure adoption. Current reference implementations tend to be buggy, hard to read, or do not focus on meaningful use cases for software developers. We also find communication challenges between actors and unclear responsibilities to be major challenges. For example, cryptographers and software developers use different terminologies and languages, and cryptography researchers tend to not feel responsible for anything more than providing theory background. Hence, developers tend to be overwhelmed and misuse standard specifications or provided implementations. We conclude the paper by discussing our results and making recommendations for academic research, industry, and standards organizations. Overall, we recommend being more transparent and open to better support the cross-disciplinary correct and secure use of cryptography. We hope this can help to bring more cryptography research from academic papers to cryptographic products and improve overall security.

Acknowledgements

We thank our interviewees for their valuable time, their helpful insights, and their openness toward our research methods.

We thank Leonie Schaewitz and Nikol Rummel (Ruhr University Bochum), who contributed to the design of the interview guide and the initial coding scheme.

This work is funded by the Deutsche Forschungsgemeinschaft (DFG, 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 No. SGS23/211/OHK3/3T/18 funded by the MEYS of the Czech Republic. This work is also supported in part by the United States National Science Foundation under Grant Number 2206865. Any findings and opinions expressed are those of the authors and do not necessarily reflect the views of the funding agencies.

Availability

We make our interview guide and our codebook publicly available in the appendix of this paper. We make the quotes used in the interview part “revisit”, as well as more background information, definitions, and further results available in an extended version of this paper, at <https://doi.org/10.5281/zenodo.8404611>.

References

- [1] Ruba Abu-Salma, M. Angela Sasse, Joseph Bonneau, Anastasia Danilova, Alena Naiakshina, and Matthew Smith. Obstacles to the adoption of secure communication tools. In *2017 IEEE Symposium on Security and Privacy (SP)*, 2017.
- [2] Yasemin Acar, Michael Backes, Sascha Fahl, Simson L. Garfinkel, Doowon Kim, Michelle L. Mazurek, and Christian Stransky. Comparing the usability of cryptographic APIs. In *2017 IEEE Symposium on Security and Privacy (SP)*, 2017.
- [3] Ross Anderson. Why cryptosystems fail. In *Proceedings of the 1st ACM Conference on Computer and Communications Security (CCS '93)*, 1993.
- [4] Ian Belton, Alice MacDonald, George Wright, and Iain Hamlin. Improving the practical application of the delphi method in group-based judgment: A six-step prescription for a well-founded and defensible process. *Technological Forecasting and Social Change*, 147:72–82, 2019.
- [5] Matthew Bernhard, Jonathan Sharman, Claudia Ziegler Acemyan, Philip Kortum, Dan S. Wallach, and J. Alex Halderman. On the usability of HTTPS deployment. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*, 2019.

- [6] Daniel J. Bernstein. Boring crypto (slides), 2015. <https://cr.y.p.to/talks.html#2015.10.05> (accessed 09/11/2023).
- [7] Daniel J. Bernstein. Cryptographic competitions. Cryptology ePrint Archive, Report 2020/1608, 2020. <https://eprint.iacr.org/2020/1608>.
- [8] Daniel J. Bernstein, Tanja Lange, and Peter Schwabe. The security impact of a new cryptographic library. In *Progress in Cryptology – LATINCRYPT '12*, 2012.
- [9] Hugh Beyer and Karen Holtzblatt. *Contextual Design: Defining Customer-Centered Systems*. Morgan Kaufmann Publishers Inc., 1997.
- [10] BITS. Website: About BITS - Bank Policy Institute. <https://bpi.com/bits/> (accessed 09/11/2023).
- [11] Scott O. Bradner. The Internet Standards Process – Revision 3. RFC 2026, 1996.
- [12] Kathy Charmaz. *Constructing Grounded Theory*. Sage, 2014.
- [13] Jeremy Clark, P. C. van Oorschot, Scott Ruoti, Kent Seamons, and Daniel Zappala. SoK: Securing email—a stakeholder-based analysis. In *Financial Cryptography and Data Security (FC'21)*, 2021.
- [14] Victoria Clarke, Virginia Braun, and Nikki Hayfield. Thematic analysis. *Qualitative psychology: A practical guide to research methods*, 3:222–248, 2015.
- [15] Juliet Corbin and Anselm Strauss. Grounded theory research: Procedures, canons and evaluative criteria. *Zeitschrift für Soziologie*, 19(6):418–427, 1990.
- [16] Sascha Fahl, Marian Harbach, Thomas Muders, Lars Baumgärtner, Bernd Freisleben, and Matthew Smith. Why eve and mallory love android: an analysis of android SSL (in)security. In *Proceedings of the 2012 Conference on Computer and Communications Security (CCS '12)*, 2012.
- [17] Patricia Fusch and Lawrence Ness. Are we there yet? data saturation in qualitative research. *Qualitative Report*, 20:1408–1416, 2015.
- [18] Simson Garfinkel. *PGP: pretty good privacy*. O'Reilly Media, Inc., 1995.
- [19] Simson L. Garfinkel and Robert C. Miller. Johnny 2: a user test of key continuity management with S/MIME and outlook express. In *Proceedings of the 2005 symposium on Usable privacy and security (SOUPS '05)*, 2005.
- [20] Martin Georgiev, Subodh Iyengar, Suman Jana, Rishita Anubhai, Dan Boneh, and Vitaly Shmatikov. The most dangerous code in the world: Validating SSL certificates in non-browser software. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security (CCS '12)*, 2012.
- [21] Matthew Green and Matthew Smith. Developers are not the enemy!: The need for usable security APIs. *IEEE Security & Privacy*, 14:40–46, 2016.
- [22] Harry Halpin. SoK: Why johnny can't fix PGP standardization. In *Proceedings of the 15th International Conference on Availability, Reliability and Security (ARES '20')*, 2020.
- [23] Julie M. Haney, Mary Theofanos, Yasemin Acar, and Sandra Spickard Prettyman. "We make it a big deal in the company": Security mindsets in organizations that develop cryptographic products. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS '18)*, 2018.
- [24] Nadia Heninger, Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. Mining your ps and qs: Detection of widespread weak keys in network devices. In *21st USENIX Security Symposium (USENIX Security 12)*, 2012.
- [25] Jan Jancar, Marcel Fourné, Daniel De Almeida Braga, Mohamed Sabt, Peter Schwabe, Gilles Barthe, Pierre-Alain Fouque, and Yasemin Acar. "They're not that hard to mitigate": What cryptographic library developers think about timing attacks. In *43rd IEEE Symposium on Security and Privacy (SP)*, 2022.
- [26] Matthias J. Kannwischer, Peter Schwabe, Douglas Stebila, and Thom Wiggers. Improving software quality in cryptography standardization projects. In *2022 IEEE European Symposium on Security and Privacy (EuroS&P)*, 2022.
- [27] Erin Kenneally and David Dittrich. The Menlo report: Ethical principles guiding information and communication technology research. *SSRN Electronic Journal*, 2012.
- [28] Michael Kranch and Joseph Bonneau. Upgrading HTTPS in mid-air: An empirical study of strict transport security and key pinning. In *Network and Distributed System Security Symposium (NDSS)*, 2015.
- [29] Katharina Krombholz, Wilfried Mayer, Martin Schmiedecker, and Edgar Weippl. "I have no idea what i'm doing" - on the usability of deploying HTTPS. In *26th USENIX Security Symposium (USENIX Security 17)*, 2017.

- [30] Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. *Research methods in human-computer interaction*. Morgan Kaufmann, 2017.
- [31] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. Reliability and inter-rater reliability in qualitative research: Norms and guidelines for CSCW and HCI practice. *Proc. ACM Hum.-Comput. Interact. (CSCW)*, 2019.
- [32] Jens Müller, Marcus Brinkmann, Damian Poddebniak, Hanno Böck, Sebastian Schinzel, Juraj Somorovsky, and Jörg Schwenk. “Johnny, you are fired!”—Spoofing OpenPGP and S/MIME Signatures in Emails. In *28th USENIX Security Symposium (USENIX Security 19)*, 2019.
- [33] Sarah Nadi, Stefan Krüger, Mira Mezini, and Eric Bodden. Jumping through hoops: Why do java developers struggle with cryptography APIs? In *Proceedings of the 38th International Conference on Software Engineering (ICSE ’16)*, 2016.
- [34] Kenneth G Paterson and Thyla van der Merwe. Reactive and proactive standardisation of TLS. In *Security Standardisation Research: Third International Conference (SSR 2016)*, 2016.
- [35] Damian Poddebniak, Christian Dresen, Jens Müller, Fabian Ising, Sebastian Schinzel, Simon Friedberger, Juraj Somorovsky, and Jörg Schwenk. Efail: Breaking S/MIME and OpenPGP email encryption using exfiltration channels. In *27th USENIX Security Symposium (USENIX Security 18)*, 2018.
- [36] Blake C. Ramsdell. S/MIME Version 3 Message Specification. RFC 2633, 1999.
- [37] Phillip Rogaway. The moral character of cryptographic work. Cryptology ePrint Archive, Report 2015/1162, 2015. <https://eprint.iacr.org/2015/1162>.
- [38] Gene Rowe and George Wright. Expert opinions in forecasting: The role of the delphi technique. *International Series in Operations Research and Management Science*, 2001.
- [39] Scott Ruoti and Kent Seamons. Johnny’s journey toward usable secure email. *IEEE Security & Privacy*, 2019.
- [40] BITS Security. [TLS] Industry Concerns about TLS 1.3, 2016. Email on IETF TLS Mailing List, available at <https://mailarchive.ietf.org/arch/msg/tls/KQIyNhPk8K6jOoe2ScdPZ8E08RE/>.
- [41] Steve Sheng, Levi Broderick, Colleen Alison Koranda, and Jeremy J Hyland. Why johnny still can’t encrypt: evaluating the usability of email encryption software. In *Proceedings of the 2006 symposium on Usable privacy and security (SOUPS ’06)*, 2006.
- [42] Christian Stransky, Oliver Wiese, Volker Roth, Yasemin Acar, and Sascha Fahl. 27 years and 81 million opportunities later: Investigating the use of email encryption for an entire university. In *2022 IEEE Symposium on Security and Privacy (SP)*, 2022.
- [43] Anselm Strauss and Juliet M Corbin. *Grounded Theory in Practice*. Sage, 1997.
- [44] Real World Crypto Symposium Website. <https://rwc.iacr.org/> (accessed 09/11/2023).
- [45] Alma Whitten and J. D. Tygar. Why johnny can’t encrypt: A usability evaluation of PGP 5.0. In *8th USENIX Security Symposium (USENIX Security 99)*, 1999.

A Appendix

We make our interview guide and our codebook publicly available in this appendix. We make the quotes used in the interview part “revisit”, as well as more background information, definitions, and further results available in an extended version of this paper, at <https://doi.org/10.5281/zenodo.8404611>.

A.1 Interview Guide

This is the interview guide we used for interviews P4–P21.

Start. Hello, thank you very much for allowing us to do this interview.! Thank you for signing the consent form! With your permission we will start the recording now. Are you okay with this? [Start Recording]

Intro to Interview. We work in an interdisciplinary project that is interested in better understanding the ecosystem of cryptography, including the research community, the development of cryptographic algorithms, protocols, standards, as well as the deployment of cryptography in software and the adoption of end user products. Specifically, we want to better understand the blockers and barriers, but also the enablers to the adoption and correct use of cryptographic solutions. What we try to do is to build a map of the ecosystem of cryptography with all the actors and all the processes involved in the adoption and use of cryptography. By this, we try to identify where in this ecosystem there are problems and potential blockers that hinder broad adoption or correct use of cryptography. We try to map all the different stages, from the cryptographic primitives to the use of cryptography in products for end users and their adoption.

Therefore, we conduct interviews with experts from the crypto community.

Warm-Up and Background. To start, we first of all would like to ask you to describe the research/work that you primarily do. What would you say is the major focus of your research/work, how do you characterise your field of work?

Follow-up: Would you say that your research/work is more theoretical or applied?

Do you have an interest in the things you work on being adopted at a large scale in the real world? If so, in which area do you see the greatest potential for the adoption of your work?

Identifying adoption problems. In your opinion, what keeps things you work on regarding cryptography from being adopted in the real world? Can you give an example of a research idea or project that you expected or hoped to be adopted, but which did not meet that expectation?

In general, when you think about the research/work that is done in the crypto community, what would you say are the most central problems when it comes to getting research into practice/application?

In your opinion, what are the greatest challenges that you and the crypto community can do something about?

Identifying possible enablers. What would you say were the most important factors or steps that led to your research/work being successfully adopted in the real world?

What are the main things I should consider if I want my cryptography research/work to be adopted in the real world?

In general, what would have to improve or change so that more research/work from the crypto community is applied and adopted in the real world?

Revisit. In the end, we would like to discuss 3-4 potentially controversial statements about the ecosystem of cryptography with you. We are interested in your opinion on these issues.

[Show + discuss quotes subsequently; read the statement to the person]. How do you understand this statement? What do you think about this statement? Do you agree/disagree with the statement?

End. We are getting to the end of the interview. We have two final questions. Do you have any suggestions for people we should interview for this study? Is there anything you want to add that we have not addressed?

Thank you very much for taking part in this interview!

B Code System

We present our resulting codebook resulting from the merging process described in Section 3.2:

Merged Codebook

1 Patterns that Lead to Adoption Challenges

- 1.1 Misunderstandings / Terminology / Things unsaid
- 1.2 Crypto is hard
- 1.3 Not My Job / Responsibility unclear
- 1.4 Usability/Reference implementations
- 1.5 Conflicting Incentives/ High effort, low reward
- 1.6 Examples - Good Examples from the real world

2 Areas and Actors

2.1 Area on Landscape

- 2.1.1 Algorithm / Protocol Development / Analysis
- 2.1.2 Standardization
- 2.1.3 Crypto Libraries / APIs
- 2.1.4 Software Development
- 2.1.5 Adoption, Deployment, and Use of Software w/ Crypto
- 2.1.6 Policy
- 2.1.7 Misc
- 2.2 Actors
 - 2.2.1 Cryptographers
 - 2.2.2 Standard Development Organizations (SDOs)
 - 2.2.3 Law Enforcement / Secret Services
 - 2.2.4 Government / Lawmakers
 - 2.2.5 Governmental Organizations
 - 2.2.6 Software Developers
 - 2.2.7 Internet Infrastructure Companies and Browser Vendors
 - 2.2.8 Commercial Companies
 - 2.2.9 Messenger Companies
 - 2.2.10 Media / Marketing
 - 2.2.11 End Users
 - 2.2.12 Misc
- 3.0 Good Quotes