

Analysing Cryptography in the Wild

A Retrospective

Martin R. Albrecht¹ and Kenneth G. Paterson²

¹ King's College London

² ETH Zurich

Abstract. We reflect on our experiences analysing cryptography deployed “in the wild” and give recommendations to fellow researchers about this process.

1 Introduction

Cryptography is not in Kansas, anymore. From e-commerce, online banking, payment systems, mobile phones, government/military communication systems and authentication systems of old, it has spread to secure messaging, video conferencing, full disc encryption, encrypted cloud storage, cryptographic custody solutions, anonymous authentication tokens, private browser analytics and privacy-preserving contact tracing, to name but a few.

Meanwhile, academic cryptography, too, has moved on and expanded its scope. Additionally, the rigour with which we study cryptographic artefacts, from primitives (such as public-key encryption), schemes (such as modes of encryption) to protocols (such as TLS), has increased.

Yet, in all the excitement about deploying cryptography to secure our brave new digital world, those cryptographic solutions in practice rarely receive that “academic” level of scrutiny, either in private before deployment or in public post deployment. This leads to a glut of cryptographic technologies “out there, in the wild”, often protecting the data of millions of people, where the veracity of their security promises is unclear.

Studying “cryptography in the wild”, then, means to find examples of cryptography being used in standards, products or deployed systems. It means to analyse them by either finding vulnerabilities and reporting them or by building security models and proofs for the cryptographic cores of these systems. The end result is that those who use these systems gain greater assurance about the security of the systems on which they rely.

Cryptography in the wild as it is understood here is *not* about developing new cryptographic primitives and pushing them towards practice. It is not applied cryptography in the sense of applying cryptography to real-world settings. This is an incredibly valuable activity and is the traditional theory-to-practice transfer process. Rather, here we are studying current practice as is. The activities considered here can be defined as applied cryptanalysis, the analysis of cryptographic solutions after their realisation in the wild. In this article, we present

39 our reflections on this activity. These are based on our combined experience of
40 pursuing it over the last couple of decades, as well as our observation of the work
41 of the small, but growing, community of researchers who have chosen to plough
42 the same furrow as us.

43 The nature of this piece – a reflection on our own practice – implies a certain
44 accumulation of references to our own work. We would like to stress that this is
45 not at all reflective of this area, which is shaped by the works of many others
46 such as [HDWH12, ABD⁺15, GGK⁺16, MBP⁺19, DPS21] to name but a few.

47 2 Methodology

48 The methodology consists of selecting an object of study (a protocol, an imple-
49 mentation, a scheme), deciding on what it means for this object to be “secure”
50 by identifying a suitable adversarial model and security goals, and analysing the
51 extent to which those goals are met.³

52 *Object Selection.* The process of selecting an object to study can start from two
53 directions. We can start from the cryptographic technology. Does a product use
54 non-trivial and/or non-standard cryptographic constructions? Do the security
55 claims of the vendor seem to match what we could reasonably expect from the
56 advertised building blocks? Is the cryptographic processing easy to analyse in
57 the form of available source code and whitepapers? If not, is it easy and legal to
58 reverse engineer? What are the first impressions? On the other hand, we may
59 also start from social or societal considerations. How big is the (claimed) “user
60 base”, i.e. people relying on this piece of technology, either in absolute terms or
61 in some area, e.g. a particular country or in a particular group of people? Are
62 there users in sometimes (e.g. protesters) or continuously (e.g. domestic abuse
63 victims) particularly vulnerable situations?

64 Our analysis of Telegram [AMPS22] came about as a result of the second kind
65 of consideration. A previous work [ABJM21] had highlighted (similarly to prior
66 works in other settings) that protesters in Hong Kong were relying on Telegram
67 to coordinate their activities. Crucially, the optional end-to-end encryption in
68 Telegram, which had been considered somewhat in the cryptographic literature,
69 played essentially no role in this setting. Instead, large group chats secured by
70 Telegram’s bespoke MTPProto 2.0 protocol in lieu of TLS were at the centre of
71 the communication infrastructure for these protests. Based on this [AMPS22]
72 studied Telegram’s secure channel protocol, found several vulnerabilities (which
73 were fixed after being reported) and fashioned a formal statement of the security
74 guarantees of the repaired protocol.

75 The latter is otherwise known as “a proof of security” in the literature.
76 However, here we deliberately shun this language to avoid the impression of a

³ The wider literature often uses the term “threat model” to also mean adversarial models as understood here. We avoid this term here because it only has a rather loose meaning and thus can give a semblance of rigour where it is lacking.

77 binary “yes/no” answer to the question of security. Rather, the relevant theorems
78 establish certain security guarantees under certain conditions.

79 This conflation – a security proof is a binary statement rather than a statement
80 of a particular security guarantee in a given model – is also one reason why object
81 selection can be contentious at times. Producing a formal security statement for,
82 say, a brittle protocol or a protocol in a product from a controversial company
83 might be considered as providing a “seal of approval” for shoddy practices.⁴ We
84 argue that often this ship has sailed. If a considerable number of people rely on
85 a given piece of technology, its security guarantees need to be understood, no
86 matter how we might view the particular parties involved.

87 It is worth noting that an object having survived one or more security audits
88 should not deter researchers from studying it. These audits are typically performed
89 under a strict time limit and so tend to focus on more surface-level security
90 issues than those uncovered by an in-depth cryptographic investigation of the
91 type discussed here. Our own prior works referenced throughout this article all
92 took several person-months to complete. Multiplying that time by the assumed
93 day-rate of an experienced auditor results in an amount that is well beyond what
94 most vendors are willing to pay. The flip-side of this observation is that we do
95 not consider the existence of positive security audits as providing meaningful
96 evidence of the soundness of a cryptographic solution. Going further, we strongly
97 suggest that consumers of cryptography demand formal theorems and proofs to
98 at least rule out large classes of attacks.

99 *Adversarial Model.* When studying cryptographic solutions in the wild, we have
100 to settle on what adversarial model we select. In the case of [AMPS22] this meant
101 not considering the Telegram server as the adversary, i.e. not targeting end-to-end
102 encryption. This decision was based on prior works suggesting this was not the
103 most pressing security concern. Often, defining such a model is straightforward:
104 if end-to-end encryption is offered then it is fair to consider the service provider
105 as an adversary. Sometimes, the marketing material used to promote a given
106 technology will provide an indication of what the developers themselves see as
107 being a suitable adversarial model, though non-standard terminology may be
108 used to make these claims.⁵

109 However, the decision to choose a particular adversarial model can be non-
110 trivial and, again, lead to controversy. Just because end-to-end encryption is
111 built-in and even advertised as present, this feature may not be of central
112 importance to the developers. Yet, some groups of people may rely on these
113 guarantees. For example, Bridgefy advertised itself as a solution for reaching

⁴ Here we do not mean misrepresenting the security guarantees, e.g. by omitting non-standard assumptions or the brittleness of the design. We simply mean formalising the security guarantees provided and the conditions under which these guarantees hold.

⁵ For example, the term “zero-knowledge encryption” has become popular in the realm of cloud storage, meaning roughly that not even the storage provider should be able to access user data. However, this term does not match any standard technical definition in the research domain.

114 target groups in disaster areas thanks to its mesh networking technology. On the
115 other hand, media reports claimed adoption of Bridgefy’s flagship messenger in
116 some protest and conflict settings, claims which Bridgefy were happy to amplify on
117 their website and social media. However, while the Bridgefy developers advertised
118 their product as being secure, this was not tended to with sufficient care: the
119 protocol was broken, fixed, and broken again with practical attacks, all in the
120 span of two years.

121 So the study of cryptography in the wild may require defining an adversarial
122 model in the face of opposition from the providers of the technology in question;
123 adversarial adversary definitions, so to speak. We consider it a responsibility of
124 academic cryptographic research to not allow vendors to define adversarial models
125 for their own technologies. As we will argue below, we see our responsibility as
126 not to the technology producers but to those who (may be forced to) rely on
127 it. Discharging this responsibility starts with defining what we consider to be
128 a suitable adversarial model based on the actual use of the technology, i.e. the
129 place of that technology in the world, not the way in which its designers prefer
130 to think about it or intend it to be used.⁶

131 *Ingesting.* The next step is to digest the available information, a process that
132 may take months to complete. This information might come from whitepapers,
133 design documents, API documentation, security audit reports or source code. It
134 may also involve reverse engineering readable source code from binary blobs or
135 minified archives, a process that we caution may itself take weeks if not months
136 to conduct.

137 A natural next step is then to build pseudo-code models describing the
138 cryptographic “core” of the object under study. Defining this core, too, can be
139 difficult. For example, cryptographic group management is typically considered
140 out of scope/outside the model when formally studying (group) messaging, yet
141 this is a routine source of vulnerabilities as any kind of secure group management
142 is often simply absent (as for example is the case with Matrix and WhatsApp).

143 Another difficulty here is to find the right level of abstraction and to unify the
144 possibly conflicting implementations in different implementations (for example
145 mobile and desktop clients) into that one pseudo-code model. Continuing with our
146 example of analysing Telegram from above, we looked at three different official
147 Telegram clients and found that they all performed slightly different checks on
148 messages to decide on whether to accept them or not, each choice having different
149 security consequences. Even official Telegram clients deviated from Telegram’s
150 own implementation advice for developers. This left open the question of what
151 the “correct” choice should be in a model.

⁶ Living up to this responsibility can lead to highly complex questions more adequately answered by the social sciences rather than computer science: how do we understand actual use and how do we cryptographers translate that into a suitable adversarial model? A technology might have many places in the world, how do we decide on an adversarial model, which “place” is taken into consideration? Or do we consider different adversarial models for a piece of technology depending on its role in different social settings?

152 The choice of model can have a dramatic effect on whether relevant and
153 practical attacks are discovered or not. Again, picking an example from our own
154 work, our first work breaking (Open)SSH [APW09] succeeded despite it enjoying
155 a proof of security [BKN02]. This was possible because our attack exploited a
156 property of the protocol abstracted away in the security model. Among other
157 works, our later work [ADHP16] produced a security theorem in a refined model.
158 This, too, was later shown to be incorrect in a work presenting a practical attack
159 on SSH [BBS23]. This attack, again, succeeded because a simplifying assumption
160 – an abstraction – made in our security model turned out to be both false and
161 significant.

162 *Attacks.* With pseudo-code models in hand, we can start to reason about potential
163 avenues of attack (or proofs). This might entail considering known attacks from the
164 literature: ECB mode, exotic or home-made encryption modes, lack of integrity
165 mechanisms, improper use of integrity (e.g. Mac-then-Encrypt Encrypt-and-
166 Mac), the presence of padding oracle vulnerabilities, nonce reuse, lack of proper
167 key separation/key reuse, lack of domain separation, bad interactions between
168 different protocols, use of weak PRNGs or home-brew randomness generation
169 methods, compression combined with encryption, etc, etc.

170 Of course, a more rewarding direction is to come up with novel (variants of)
171 attacks, possibly chaining together multiple attacks to achieve a given aim. This
172 is a process sometimes endearingly referred to as “stunt cryptography” (a phrase
173 we attribute to Thomas Ptacek⁷). This may be needed if the object under study
174 does not allow us to apply known attacks as is, yet the object is – by inspection
175 – insecure. Turning that inspection into a practical attack might then require
176 significant new ideas.

177 *Proof of Concepts.* It is much easier to convince third parties of the seriousness
178 of a vulnerability by providing a working proof-of-concept exploit. Indeed, most
179 bug bounty programs require this of submitters. Yet, developing such exploits
180 can be a laborious process. The obstacle to pulling off that stunt might be a
181 computation taking 2^{60} steps and several person-months to carefully implement,
182 cf. the construction of a special key pair needed to complete an attack in [PST23].
183 Even if a vulnerability should be exploitable without consuming too many
184 resources (queries, time, memory), developing proof-of-concept exploits can be
185 time and resource intensive. Quite often, the complexity arises from wrestling with
186 the state-machine or parser of the considered protocol and not the cryptographic,
187 and thus interesting, core.

188 *Proofs.* Finally, the cryptographic object under study might be amenable to a
189 more formal analysis, i.e. the establishment of a formal theorem characterising its
190 security in the previously defined adversarial model. This rules out large classes
191 of attacks and thus gives greater confidence in the actual security properties of
192 the studied object. The difficulties here often stem from needing to first develop

⁷ <https://news.ycombinator.com/item?id=31829130>

193 models of security, the studied protocol attempting to use novel cryptographic
194 functionalities, or the studied protocol using unconventional approaches to solve
195 cryptographic problems with known solutions, e.g. by hashing a key together with
196 some message in lieu of a MAC construction. For these bespoke constructions, we
197 might be able to reduce their security to some standard and well-studied assump-
198 tions, or we might need to simply state the required (likely novel) assumptions on
199 the underlying primitives needed for the proof to go through. In the latter case,
200 opportunities are created for follow-up research to either support or invalidate
201 these assumptions.

202 3 Responsibilities

203 A standard step in this line of research, if significant vulnerabilities are discovered,
204 is to consider how to disclose these vulnerabilities in a way that minimises harm,
205 defined in some way. The standard approach here is “coordinated vulnerability
206 disclosure”, previously known as “responsible disclosure”. This involves privately
207 disclosing vulnerabilities to vendors, typically with a 90-day deadline attached. If
208 by the end of the disclosure period no remedy is made available, the vulnerability
209 will be publicly announced regardless. Prior to this deadline, the invitation is
210 for the vendor and the researchers to coordinate the public disclosure. In our
211 experience, it may happen for the initial deadline to slip to 120 days but not
212 longer. We strongly recommend that the decision on timing should be maintained
213 under the complete control of the research team.

214 It is worth, though, mentioning that for some (classes of) vulnerabilities a
215 90 day disclosure vulnerability is *a priori* unrealistic. Examples that come to
216 mind are hardware/micro-architecture level vulnerabilities. Put differently, the
217 “90 days” rule is somewhat arbitrary but works fairly well for software-based
218 vulnerabilities when the vendor has an established update process.

219 A common misconception here is that the research team has a responsibility
220 to the vendor. Yet, no such particular responsibility exists, as typically there is
221 no pre-established legal or commercial relationship between the parties. Rather,
222 insofar we want to or can speak of responsibility, we see it as being to those
223 who (have to) rely on the technology provided by the vendor in question.⁸ The
224 users of an insecure-by-design protocol may be better protected by them being
225 warned against relying on it rather than attempts to patch it by an inexperienced
226 team of developers. Yet, avoiding a given protocol may simply be impossible for
227 many. This suggests cooperation with vendors as being a solid strategy. This may
228 involve informally advising vendors on remediation strategies and reviewing their
229 patches. We recommend that research teams discuss amongst themselves (or
230 potentially with trusted mentors) to whom they consider themselves responsible
231 before disclosing to vendors.

232 Researchers do have a responsibility to be realistic about the impact of any
233 vulnerabilities discovered when communicating publicly. Insinuating the existence

⁸ While we, the authors, certainly consider this our responsibility, we are mindful that we are in no position to prescribe such a responsibility for others.

234 of serious vulnerabilities in a robust system may deter some from using it, possibly
235 encouraging them to pivot to insecure alternatives.

236 4 Reception

237 The style of works discussed here is received by three rather distinct audiences.

238 *Vendors.* Disclosing vulnerabilities to vendors is often a time-consuming and
239 frustrating process. Some vendors may simply never have handled a vulnerability
240 disclosure before and thus lack any processes for dealing with one when it arrives
241 out of the blue. Other vendors enthusiastically start to make patches in public
242 code repositories, forgetting everything that the word “coordinated” implies, and
243 potentially allowing the patches to be reverse-engineered and exploited by others
244 before new product versions can be distributed. Yet other vendors do not ever
245 tell their users that their product is being updated for security reasons. More
246 rarely, disclosure is a butter-smooth and pleasurable experience, in which case
247 the vendor concerned should be publicly credited.

248 A worrying trend in this area is that vendors “outsource” this process to
249 “bug bounty” service providers. These services are based around the idea of
250 vulnerability disclosure being financially rewarded, often in return for some form of
251 agreement on non-disclosure. This process is unsuited to the activities considered
252 here; both with regard to their motivation – improving security independently of
253 financial reward – and with regard to their content – cryptographic vulnerabilities
254 are rarely disclosed and thus the industry around vulnerability disclosures and
255 bug bounties is inexperienced in dealing with them.

256 We recommend avoiding any intermediaries in the disclosure process, since
257 they may impose their own policies with negative consequences. For example,
258 our disclosure of the previously mentioned plaintext recovery attack against
259 (Open)SSH [APW09] was hampered by us going via the UK’s Centre for the
260 Protection of National Infrastructure, resulting in the following notification on
261 the OpenSSH project webpages:

262 *The OpenSSH team has been made aware of an attack against the SSH*
263 *protocol version 2 by researchers at the University of London. Unfortu-*
264 *nately, due to the report lacking any detailed technical description of the*
265 *attack and CPNI’s unwillingness to share necessary information, we are*
266 *unable to properly assess its impact.*⁹

267 Vendors who ship cryptography may also overestimate their own understand-
268 ing of it. Vendors might not be impressed with an attack on the IND-CPA security
269 of their scheme but may expect key or plaintext recovery. Similarly, they might
270 dismiss an attack taking 2^{60} steps as too expensive, underestimating the tendency
271 of “attacks only getting better” with time. As an extreme illustration of this,
272 consider the rapid evolution of attacks against MEGA, an end-to-end encrypted

⁹ See <https://www.openssh.com/txt/cbc.adv>.

273 cloud storage service. The first analysis required a user to make 512 logins in
274 order to be able to recover their private key [BHP23]. This was followed within
275 a couple of months by a second, more sophisticated analysis needing only six
276 logins [HR23]. Then just a few months later, a third analysis requiring only two
277 logins was published [AHMP23]. This evolution somewhat undermined MEGA’s
278 claim that their users were secure because, according to their logs, none of them
279 had logged in as many as 512 times. Here we must also recognise MEGA, however,
280 as being one of those vendors who made disclosure and remediation a positive
281 experience.

282 Moreover, vendors have an incentive to “downplay” the significance of vul-
283 nerabilities, fearing negative market perception. A shift towards downplaying is
284 often seen as the agreed-upon disclosure date approaches and the conversation
285 moves from the security or development team to a more senior party, such as a
286 CEO. The flip side is that the disclosing researchers have an interest to “play up”
287 the vulnerability for obvious prestige reasons.

288 *Scientific Community.* Reception in the scientific community is typically both
289 rather negative and exceptionally positive, with the two reactions occurring
290 simultaneously. First, strong attacks or detailed proofs of high-profile targets
291 routinely win best paper awards at security research venues, highlighting the
292 community’s appreciation for this sort of work. At the same time, it is rare for
293 these works to find a home in more traditional cryptographic venues, e.g. the
294 flagship venues of the International Association for Cryptologic Research (IACR).
295 This is because many reviewers perceive this endeavour as lacking scientific
296 significance, as expressed in reviews asking for comparative studies, lessons
297 learned for the field, a higher degree of novelty, attacks with greater technical
298 depth, new primitives, new security models, etc.

299 In an ideal world, every system relying on cryptography would be thoroughly
300 analysed before release, drawing on the wealth of scientific literature available,
301 and there would be no need for the study of cryptography in the wild. This
302 does happen sometimes, for example, with TLS 1.3 or more recently when Apple
303 engaged two teams of academic researchers to review their new post-quantum
304 version of the iMessage protocol prior to its release. However, this is rarely the
305 case. Our contention is that we, the scientific community, can gain a lot from
306 observing cryptography in the wild. Not only do we obtain valuable examples
307 for use in the classroom (answering in concrete ways questions like “Why is
308 key separation really needed?”), but we may also begin to understand why
309 cryptography is seemingly so hard to get right in practice, why developers so
310 often get it wrong, and how we can design new cryptographic primitives and
311 protocols so that they can be more safely consumed.

312 Indeed, if we consider science to be the process by which we get to understand
313 the world, then a security analysis of a significant object in the world – say
314 a protocol relied upon by millions of people – is then in and of itself a valid
315 scientific result. The situation here is not at all dissimilar to many other sciences
316 which feature branches labelled “theoretical”, branches labelled “applied” and
317 branches labelled “experimental”, “empirical” or “observed”. We consider the

318 study of cryptography as it exists in the world as “observed cryptography” in
319 that sense.

320 *“The Public”*. Eventually, the results of the analysis are made public. If those
321 results are “only” a formal security statement then news of this barely reaches
322 even IT-security-focused practitioner circles. However, if vulnerabilities are also
323 disclosed and if the object is sufficiently high-profile then this may well produce
324 a short burst of interest with news reporting and social media threads. In our
325 own practice, we typically reach out to a journalist before disclosure to give
326 them a chance to get a more accurate technical account of what is and what is
327 not broken and what that means. This aids with communicating the findings
328 clearly. In addition, we may set up a special website containing high-level and
329 less-technical discussions of the vulnerabilities, the disclosure process and patch
330 status. This, of course, creates an additional burden on the research team. We
331 regard the production of logos and stickers as highly optional.

332 The typical social media response to vulnerabilities is then some variant of
333 “don’t roll your own crypto”. While this truism is, well, true – people not trained
334 in the development of cryptography should not develop it, just like people who
335 are not trained to design bridges should avoid that activity – it glosses over
336 what allows someone to “roll their own”: careful modelling and cryptographic
337 design, formal security analysis, and an appreciation for secure cryptographic
338 implementation. Instead of this simple scientific message, we are typically left
339 with a simple measurement contest – “who is better at it” – which does not help
340 to improve the general state of cryptographic solutions.

341 Moreover, despite the heavy moralising typically found on social media about
342 the failings of vendors to roll cryptography, as far as we can tell, this “name and
343 shame” approach has little lasting effect. In our own experience, such approaches
344 do not harm the popularity of the relevant vendor.

345 5 Conclusion

346 We have described some of our experiences with and approaches to studying
347 cryptography in the wild. We encourage the further development and refinement
348 of this folklore methodology. As the scope of application of cryptography continues
349 to broaden, it becomes ever more deeply embedded as a foundation of privacy,
350 trust and security in our digital society. So we anticipate the topic of cryptography
351 in the wild to be of perennial interest and value.

352 Acknowledgements

353 We thank Rikke Bjerg Jensen for helpful discussions.

354 **References**

- 355 ABD⁺15. David Adrian, Karthikeyan Bhargavan, Zakir Durumeric, Pierrick Gaudry,
356 Matthew Green, J. Alex Halderman, Nadia Heninger, Drew Springall, Em-
357 manuel Thomé, Luke Valenta, Benjamin VanderSloot, Eric Wustrow, Santi-
358 ago Zanella-Béguelin, and Paul Zimmermann. Imperfect forward secrecy:
359 How Diffie-Hellman fails in practice. In Indrajit Ray, Ninghui Li, and
360 Christopher Kruegel, editors, *ACM CCS 2015*, pages 5–17. ACM Press,
361 October 2015. 1
- 362 ABJM21. Martin R. Albrecht, Jorge Blasco, Rikke Bjerg Jensen, and Lenka Mareková.
363 Collective information security in large-scale urban protests: the case of
364 hong kong. In Michael Bailey and Rachel Greenstadt, editors, *USENIX*
365 *Security 2021*, pages 3363–3380. USENIX Association, August 2021. 2
- 366 ADHP16. Martin R. Albrecht, Jean Paul Degabriele, Torben Brandt Hansen, and
367 Kenneth G. Paterson. A surfeit of SSH cipher suites. In Edgar R. Weippl,
368 Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai
369 Halevi, editors, *ACM CCS 2016*, pages 1480–1491. ACM Press, October
370 2016. 2
- 371 AHMP23. Martin R. Albrecht, Miro Haller, Lenka Mareková, and Kenneth G. Paterson.
372 Caveat implementor! Key recovery attacks on MEGA. In Carmit Hazay
373 and Martijn Stam, editors, *EUROCRYPT 2023, Part V*, volume 14008 of
374 *LNCS*, pages 190–218. Springer, Heidelberg, April 2023. 4
- 375 AMPS22. Martin R. Albrecht, Lenka Mareková, Kenneth G. Paterson, and Igors
376 Stepanovs. Four attacks and a proof for Telegram. In *2022 IEEE Symposium*
377 *on Security and Privacy*, pages 87–106. IEEE Computer Society Press, May
378 2022. 2, 2
- 379 APW09. Martin R. Albrecht, Kenneth G. Paterson, and Gaven J. Watson. Plaintext
380 recovery attacks against SSH. In *2009 IEEE Symposium on Security and*
381 *Privacy*, pages 16–26. IEEE Computer Society Press, May 2009. 2, 4
- 382 BBS23. Fabian Bäumler, Marcus Brinkmann, and Jörg Schwenk. Terrapin attack:
383 Breaking SSH channel integrity by sequence number manipulation. *CoRR*,
384 abs/2312.12422, 2023. 2
- 385 BHP23. Matilda Backendal, Miro Haller, and Kenneth G. Paterson. MEGA: Mal-
386 leable encryption goes awry. In *2023 IEEE Symposium on Security and*
387 *Privacy*, pages 146–163. IEEE Computer Society Press, May 2023. 4
- 388 BKN02. Mihir Bellare, Tadayoshi Kohno, and Chanathip Namprempre. Authenti-
389 cated encryption in SSH: Provably fixing the SSH binary packet protocol.
390 In Vijayalakshmi Atluri, editor, *ACM CCS 2002*, pages 1–11. ACM Press,
391 November 2002. 2
- 392 DPS21. Luca De Feo, Bertram Poettering, and Alessandro Sorniotti. On the
393 (in)security of ElGamal in OpenPGP. In Giovanni Vigna and Elaine Shi,
394 editors, *ACM CCS 2021*, pages 2066–2080. ACM Press, November 2021. 1
- 395 GGK⁺16. Christina Garman, Matthew Green, Gabriel Kaptchuk, Ian Miers, and
396 Michael Rushanan. Dancing on the lip of the volcano: Chosen ciphertext
397 attacks on apple iMessage. In Thorsten Holz and Stefan Savage, editors,
398 *USENIX Security 2016*, pages 655–672. USENIX Association, August 2016.
399 1
- 400 HDWH12. Nadia Heninger, Zakir Durumeric, Eric Wustrow, and J. Alex Halderman.
401 Mining your ps and qs: Detection of widespread weak keys in network
402 devices. In Tadayoshi Kohno, editor, *USENIX Security 2012*, pages 205–220.
403 USENIX Association, August 2012. 1

- 404 HR23. Nadia Heninger and Keegan Ryan. The hidden number problem with small
405 unknown multipliers: Cryptanalyzing MEGA in six queries and other applica-
406 tions. In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *PKC 2023*,
407 *Part I*, volume 13940 of *LNCS*, pages 147–176. Springer, Heidelberg, May
408 2023. 4
- 409 MBP⁺19. Jens Müller, Marcus Brinkmann, Damian Poddebniak, Hanno Böck, Sebas-
410 tian Schinzel, Juraj Somorovsky, and Jörg Schwenk. “Johnny, you are fired!”
411 - Spoofing OpenPGP and S/MIME signatures in emails. In Nadia Heninger
412 and Patrick Traynor, editors, *USENIX Security 2019*, pages 1011–1028.
413 USENIX Association, August 2019. 1
- 414 PST23. Kenneth G. Paterson, Matteo Scarlata, and Kien T. Truong. Three lessons
415 from Threema: Analysis of a secure messenger, 2023. 2