



Post Quantum Noise

Yawning Angel*
yawning@oasislabs.com
Oasis Labs
USA, San Francisco

Benjamin Dowling*
b.dowling@sheffield.ac.uk
University of Sheffield
UK, Sheffield

Andreas Hülsing*
andreas@huelising.net
TU Eindhoven
The Netherlands, Eindhoven

Peter Schwabe*
peter@cryptojedi.org
MPI-SP
Germany, Bochum

Florian Weber*[†]
f.j.weber@tue.nl
mail@florianjw.de
TU Eindhoven
The Netherlands, Eindhoven

ABSTRACT

We introduce PQNoise, a post-quantum variant of the Noise framework. We demonstrate that it is possible to replace the Diffie-Hellman key-exchanges in Noise with KEMs in a secure way. A challenge is the inability to combine key pairs of KEMs, which can be resolved by certain forms of randomness-hardening for which we introduce a formal abstraction. We provide a generic recipe to turn classical Noise patterns into PQNoise patterns. We prove that the resulting PQNoise patterns achieve confidentiality and authenticity in the fACCE model. Moreover we show that for those classical Noise-patterns that have been conjectured or proven secure in the fACCE model our matching PQNoise patterns eventually achieve the same security. Our security proof is generic and applies to any valid PQNoise pattern. This is made possible by another abstraction, called a hash-object, which hides the exact workings of how keying material is processed in an abstract stateful object that outputs pseudorandom keys under different corruption patterns. We also show that the hash chains used in Noise are a secure hash-object. Finally, we demonstrate the practicality of PQNoise delivering benchmarks for several base patterns.

CCS CONCEPTS

• Security and privacy → Cryptography; Security protocols.

KEYWORDS

Protocol; Post-Quantum Cryptography; Noise; PQNoise; Provable Security

ACM Reference Format:

Yawning Angel, Benjamin Dowling, Andreas Hülsing, Peter Schwabe, and Florian Weber. 2022. Post Quantum Noise. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security (CCS '22)*.

*Author list in alphabetical order, see: <https://www.ams.org/profession/leaders/culture/CultureStatement04.pdf>

[†]Second mail is *strongly* preferred, first is only listed for institutionally paid submission-fee.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CCS '22, November 7–11, 2022, Los Angeles, CA, USA

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9450-5/22/11.

<https://doi.org/10.1145/3548606.3560577>

November 7–11, 2022, Los Angeles, CA, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3548606.3560577>

1 INTRODUCTION

In 2014, Perrin set out to simplify the process of designing, describing, analyzing, and securely implementing secure-channel protocols through the *Noise Protocol Framework* [32]. The success of this endeavour is illustrated by the long list of users, including WhatsApp, WireGuard, Slack, I2P and the Lightning Network [29]. At the heart of Noise is the idea of using Diffie-Hellman (DH) key exchange [17] as the only asymmetric primitive – forward secrecy is achieved through DH with ephemeral keys, authentication of parties is achieved through static DH keys. Perrin informally described this concept of authenticated key agreement without signatures used by Noise as “*Hash all these DHs together to get a final key*” [33].

What DH operations are performed and what exactly is “hashed together” is expressed in *handshake patterns*, that Noise specifies in a concise and easy-to-parse language. As one example, consider the “KN” pattern (“\” indicates a line break): “-> s \\\... \\\-> e\\<- e, ee, se”

On a high level what this pattern means is that the responder (on the right side of the arrows) is aware of the static public DH key (-> s) of the initiator (on the left side of the arrows) before the online phase of the protocol starts (-> s is before ...). In the online phase the initiator first generates an ephemeral DH key pair and sends the ephemeral public key to the responder (-> e). The responder then also generates an ephemeral key pair (e) and sends the public key to the initiator (<- e). Both parties then combine their respective ephemeral secret keys with the ephemeral public key of the peer to obtain a shared ephemeral-ephemeral DH key (ee) and additionally compute the static-ephemeral DH se using the initiator’s static secret key and the responder’s ephemeral public key on the initiator’s side and the initiator’s static public key and the responder’s ephemeral secret key on the responder’s side. For a more detailed description of how patterns translate into cryptographic operations and protocols messages, and in particular how public and shared keys are absorbed into protocol state, see the Noise Protocol Framework specification [32]; for the cryptographic protocol implementing the KN pattern, see Figure 1.

The KN pattern is an example of a *named pattern* in Noise; a subset of these named patterns are the so-called *fundamental patterns*.

There exist twelve interactive and three non-interactive fundamental patterns. These patterns exist for every combination of each party having 1) No static public key, 2) a static public key **Known** to their peer, or 3) a static public key that has to be transmitted (**X**) during the interaction. For the initiator there is furthermore the possibility that 4) he has a static public key that he is willing to send with the Initial message, even if this may reduce anonymity. For each of the resulting 12 cases, Noise defines a fixed pattern, named by the two letter combination derived from concatenating the letters indicating the case for the initiators key and that for the responders key, giving: NN, NK, NX, KN, KK, KX, XN, XK, XX, IN, IK and IX. E.g., NN deals with the case where neither party has a static public key, whereas IK applies to the case where the initiator's key is not initially known to the responder, but the responder's key is known to the initiator upfront and the initiator is willing to send his public key with the first message.

One interesting feature of secure-channel protocols in Noise is that they do not separate the key-agreement or handshake phase from the data-transmission phase: in fact, Noise allows to send early payload messages together with every handshake message. These early payload messages are encrypted under whatever shared key material has already been established, but they typically do not enjoy the full security properties established by the end of the handshake. This means that we cannot analyze the security of Noise handshakes as standalone, monolithic authenticated-key-agreement protocols in, for example, the CK [13], eCK [28], or CK⁺ [26] model. This issue was addressed by Dowling, Rösler, and Schwenk with the introduction of the fACCE model [19], a multi-stage variant of the ACCE model introduced for the analysis of TLS [24].

The design decision to rely on Diffie-Hellman as the only asymmetric primitive in Noise leads to elegant protocols offering extensive security and privacy properties; the instantiation of DH with X25519 [7] in Noise also leads to efficient implementations of these protocols in multiple programming languages. However, the strong reliance on DH also comes with a downside: Noise does not have any straight-forward migration path to post-quantum cryptography. Indeed, DH offers the functionality of non-interactive key exchange (NIKE) [20], and no efficient post-quantum instantiation of this functionality is known today. The most plausible candidate is CSIDH [14], but unfortunately even at bleeding-edge security levels and with all state-of-the-art optimizations it is about three orders of magnitude slower than X25519 [6]. Also, the concrete security against quantum attackers is still subject of heavy debate [9, 11, 31].

Our Contribution. The closest primitive to DH that *does* have efficient post-quantum instantiations is key encapsulation mechanisms (KEMs). For specific DH-based authenticated key-exchange protocols, KEMs have been used before to replace DH, e.g., in PQWireGuard [23]; in this paper we generalize this approach and investigate what a purely KEM-based, post-quantum Noise framework looks like.

While it is straightforward to replace DH by a KEM in some cases, in others it is not, for a multitude of reasons: First, authentication with KEMs can only be done in an interactive challenge-response fashion, whereas it is possible to view any DH public

key as an already existing challenge, allowing for non-interactive authentication. Secondly, it is possible to combine arbitrary DH keyshares, which is not the case for KEMs as public keys cannot be combined. This causes issues in the cases where Noise combines two static shares. Thirdly, Noise is extremely flexible and offers a huge amount of possible patterns. So far, computational security proofs are given for individual patterns which results in a large number of individual security proofs, and many patterns without computational proofs of security at all, though a number of symbolic analyses of Noise exists [22, 25].

We resolve all of these issues. We provide a recipe to translate a Noise-pattern into a PQNoise pattern that, at the possible cost of additional roundtrips, achieves the same confidentiality and authenticity as the original pattern. In some cases we can do better than applying our generic translation. We provide optimized PQNoise alternatives for all 12 interactive fundamental patterns and for the non-interactive N-pattern (The K- and X-patterns don't have non-interactive equivalents in PQNoise). Our recipes solve the second issue by noting that approaches like the NAXOS trick provide a way to mix a static secret into the randomness effectively guaranteeing that the result is secret as long as either the randomness or the static secrets are uncorrupted. We introduce *static-ephemeral entropy combination* (SEEC) as an abstraction of these approaches, which is suitable for the security analysis of PQNoise, is met by many existing constructions, and allows the implementer to choose a suitable instantiation for their respective target system.

We give a generic proof of security in the computational model resolving issue three. This is enabled by the introduction of another abstraction termed "hash-object", a formalization of the "Hash all these DHs together to get a final key" idea. A hash-object is a stateful object into which values can be fed and from which keys can be extracted. When using this to analyze PQNoise, we require that the outputs of this object are pseudorandom as long as at least one random input was absorbed into the object before that is unknown to the adversary. We provide a formal definition of this primitive and prove that the way Noise hashes key shares into a hash-chain instantiates it. This abstraction allows to remove a lot of pattern-specific complexity from the security proofs, which in turn allows us to write them in a generic manner. We conjecture that this approach is fully applicable to all versions of classical Noise, allowing for a more comprehensive computational analysis than what currently exists, though we leave that for future work. We remark here that our proof does in fact not just apply to the specific PQNoise patterns that we specify, but to every PQNoise protocol, including for example hybrid ones (which we don't specify here).

Our security analysis is performed in the fACCE-model [19], that was already used in the analysis of Noise, though we modify the model in a few places. First, there are some cosmetic changes that we believe make both the model and the resulting statements more accessible, such as renaming confusingly named operations. Second, we provide the resulting security statements as a simple table that maps uncorrupted secrets to achieved security goals in a given stage instead of providing a list of named security goals that are also not necessarily independent. This allows to simplify the freshness conditions significantly.

Finally we present a proof-of-concept implementation of PQNoise in Go and report on benchmarking results. The results clearly demonstrate the practicability of post-quantum key exchanges in a wide variety of settings. We remark here that providing a post-quantum version of Noise essentially provides a solution for all applications that need key exchanges that are requiring neither backwards-compatibility nor crypto-agility at runtime; naturally the former is not a problem that can be solved generically and the latter is a property whose desirability is getting called increasingly into question as more and more newer protocols don't offer which matches community-surveys [37]. The software is available from: <https://gitlab.com/yawning/nyquist/-/tree/experimental/pqnoise>.

2 PQNOISE

In this section we present our design for PQNoise. We start with a description of PQNoise. Afterwards, we introduce SEEC (Static-Ephemeral Entropy Combination), our abstraction of methods that mix a static key into the randomness source to guarantee security in a bad randomness setting. With this we then present our recipe to translate Noise patterns into PQNoise patterns. We conclude with a discussion of the optimized fundamental-patterns for PQNoise.

2.1 PQNoise

PQNoise aims to be the post-quantum counterpart to Noise and shares many of its characteristics. One of these is the generic approach of providing a large number of possible patterns whose description is similar to that of Noise patterns. However, given that PQNoise uses KEMs for key exchange, some tokens are different. The single-letter tokens (s and e) stand for the sending of public keys, just as before. The four tokens (ee, se, es and ss) representing combination of DH-key-shares are dropped. In their place PQNoise introduces ekem and skem, that indicate the sending of a ciphertext that was encapsulated to the ephemeral/static public key of the receiving party and the mixing of the encapsulated secret into the hash-object (our abstraction of the hash-chains used in Noise) similar to the old two-letter tokens.

On a lower abstraction-level PQNoise intentionally works essentially exactly like classical Noise, with the exceptions that we replace the asymmetric primitives and use SEEC for the entropy of all probabilistic algorithms, except the generation of static keys. Noise starts mixing shared keys into its hash chain as soon as they are available, extracts a session key from it and starts encrypting all further messages, except the ephemeral key shares, using an AEAD scheme. We stick to this approach.

Noise and PQNoise maintain effectively two hash-chains (one of which we will later model as a hash-object): The first one, h , is initialized as the hash of a pattern-label. Whenever a value x needs to be added to it, the party in question computes $H(h, x)$ and replaces h with it. The first thing that is added to h are unspecified associated data that can be chosen freely by the application. Following that all public keys are added as soon as they are transmitted (if they are Known, they get added at the very start). Furthermore all AEAD-ciphertexts are added after they are sent/successfully decrypted. In turn h is used directly (without further hashing) as associated data whenever an AEAD-ciphertext is created and is

intended to be usable as a unique handshake-hash after the completion of the handshake-phase.

The second hash-chain ck is the one from which the protocol derives its encryption-keys. The key-chain ck is initialized by the hash of the pattern-label as well. Afterwards, whenever both parties establish a shared secret k_i (in classical Noise a Diffie-Hellman shared secret, in PQNoise the key that is encapsulated in a KEM-ciphertext), Noise computes a temporary value (which we will refer to as tmp) as $HMAC-HASH(ck, k_i)$ and derives a new value for ck and whatever keys it needs by computing $HMAC-HASH(tmp, ctr)$, where ctr is set to 0 for the new value of ck and to 1 for the derived key. There is one exception to this with the last addition of a shared secret, where the two produced values are not used as a new value for ck and a session-key, but instead as the initiator's and responder's session keys for the remaining session. For the purposes of our analysis we model this as hash-object and refer to Section 4.1 for more details.

The actual encryption in PQNoise is done via an AEAD-scheme, where the key is the session-key derived from ck , h is used as associated data and the nonce is a simple counter, that is initially set to zero, increases by 1 with every use and is reset to zero once a new session-key is established. To send an ephemeral key (e), the sender creates a new ephemeral keypair pk_e, sk_e using the key-generation-algorithm with the output of SEEC as entropy and adds pk_e to the current payload and h . To send a static key (s), the sender adds their static public key to the current payload and h .

Sending of KEM-ciphertexts (ekem, skem) is where the largest differences between Noise and PQNoise are: Firstly we differentiate between the ephemeral (EKEM), the initiator's (IKEM) and the responder's (RKEM) KEM. This allows the use of different KEMs in the same protocol in a similar manner to PQWireguard [23] which can allow for more efficient protocols and enable a "poor man's hybrid encryption", where even a catastrophic break of one scheme preserves confidentiality if there is no additional corruption.) As depicted in Algorithm 1, during Send the sender encapsulates a key k_x to the receiver's public key pk_x using hardened randomness (see Section 2.2). If the KEM in question is not ephemeral (for compatibility with Noise) and there is already a shared key k_i (which by the requirements of Noise has to be at least partially derived from EKEM) the resulting ciphertext ct_x (together with possible further payload pl that doesn't further affect the KEM-operation, see the full version [3]) is encrypted with the AEAD-scheme under k_i using the current nonce n and the current handshake-hash h as associated data and the resulting ciphertext is added to the send-buffer. Otherwise ct_x is added directly to the send-buffer. In either case h is updated by hashing the previous value with whatever was added to the send-buffer and k_x is added to the keychain by calling $ck.in(k_x)$, producing the next secret key k_{i+1} . Lastly the sender sets the nonce n to 0.

The actions by the receiver during Recv mirror those of the sender: After either decrypting or receiving ct_x he adds what he received to h , decapsulates it with his secret key sk_x and inputs the resulting key into the key-chain ck producing k_{i+1} and resets the nonce n to 0.

We refrain from providing detailed pseudocode for the other operations here as they are essentially identical to classical Noise and

Algorithm 1: Transmission of KEM-ciphertexts.

```

1 Function Send:
2   ...
3    $r \leftarrow \text{SEEC.GenRand}(\text{seec\_sk})$ 
4    $ct_X, kk_X := \text{XKEM.encaps}(pk_X, r)$ 
5   if  $\text{XKEM} \neq \text{EKEM} \wedge k_i \neq \perp$ :
6      $c_i := \text{AEAD.enc}(k_i, n, h, pl)$ 
7      $h := H(h, c_i)$ 
8   else:
9      $h := H(h, ct_X)$ 
10   $k_{i+1} := ck.in(kk_X), n = 0$ 
11  ...
12 Function Recv:
13  ...
14  if  $\text{XKEM} \neq \text{EKEM} \wedge k_i \neq \perp$ :
15     $... || ct_X := \text{AEAD.dec}(k_i, n, h, c_i)$ 
16     $h := H(h, c_i)$ 
17  else:
18     $h := H(h, ct_X)$ 
19   $kk_X := \text{XKEM.decaps}(sk_X, ct_X)$ 
20   $k_{i+1} := ck.in(kk_X), n = 0$ 
21  ...

```

refer to the full version [3] instead. We note however that we implemented a compiler that transforms any PQNoise-pattern into such detailed pseudocode while also performing some basic soundness-checks (for example that there is no use of keys that are not yet known) on the input and full type-checking on the produced code (though the types are not displayed as part of the LaTeX-output). We provide the pseudocode resulting from the thirteen fundamental PQNoise-patterns (see below) in the full version [3] and the compiler at <https://florianjw.de/diverses/pqnoise-codegen.tar.bz2>.

To give an illustrative example of how PQNoise and Noise differ we refer to Figure 1, which displays the KN-pattern of classical Noise and its PQNoise-counterpart. The main differences can be seen around the use of KEMs: Since KEM-keys cannot be combined, PQNoise requires the sending of additional ciphertexts (ct_e and ct_j instead of just g^b) which also have to be encrypted (c_0) and added to h . That (and the use of SEEC) aside, the protocols are however remarkably similar. Overall these similarities and differences are representative for the other patterns.

2.2 SEEC

Bad random number generators are a real-world issue. And it does not matter for this whether they are intentionally broken by malicious governments [10] or accidentally by well-meaning individuals [16]. Hence, this is covered in modern definitions of security for protocols, introducing the corruption of ephemeral secrets as a valid attack.

The Noise-framework itself considers this an issue that should be solved on system level instead of per-protocol and does not include any countermeasures for this case. Nonetheless the KK- and the IK-patterns derive their key among other sources from




Figure 1: The KN patterns of classical Noise (left) and PQNoise (right). For reasons of space we use the following conventions here: **highlighted** actions are performed by both parties as soon as they receive all necessary values to perform the computations in question. If any decryption- or decapsulation-algorithm returns $\perp$, the party in question aborts the protocol.

a static-static Diffie-Hellman exchange. The intention behind this was purely to achieve initiator-authenticity earlier than otherwise possible. However, later academic analysis [19] came to rely upon it to achieve protection from so called Maximal Exposure (or MEX-) attacks [26], where the adversary can learn the randomness of parties.

Removing this protection from PQNoise would therefore weaken the patterns compared to the security that published analysis promises for their classical counterparts, even if those properties were never promised by the designers of these patterns. As we outlined above, there is no direct replacement for the Static-Static exchange when using KEMs. Nevertheless, similar security properties can be achieved when combining a static secret with the random coins used in the encapsulation algorithm.

The first time something like this was proposed was as part of the NAXOS-protocol [28]. Later Fujioka, Suzuki, Xagawa and Yoneyama [21] used “twisted PRFs” to achieve a similar result. Later still Akhmetzyanova, Cremers, Garratt, Smyshlyayev and Sullivan [1] proposed to use hashed signatures of random messages, arguing that the secret keys for signature-schemes often reside in special protected hardware to begin with, making this a practical match. This was then standardized by the IETF as RFC 8937 [15].

Since the exact choice of such a system should be transparent for all peers, we consider it an implementation detail and refrain from specifying any concrete technique. Instead we introduce the notion of Static-Ephemeral Entropy Combination (SEEC) as an abstraction of all of these and similar approaches and base our analysis on this abstract notion. This allows us to generically analyze

PQNoise without forcing implementers to use any specific system. Indeed, SEEC also covers cases where the mixing is done on system level, matching well with the philosophy of Noise while formally describing the requirements to achieve security also under MEX attacks.

Intuitively a SEEC-scheme consists of a pair of algorithms *GenKey* and *GenRand*. *GenKey* is a probabilistic algorithm that returns a long-term key *sk*. *GenRand* then takes *sk* and some random coins and returns a pseudorandom value *r*, where *r* is indistinguishable from a true random value if either *sk* is uncorrupted and the random coins fresh (but possibly known to the adversary) or if the random coins are uncorrupted. Additionally we allow but do not require *GenRand* to modify *sk*. The reason is that this is necessary to allow SEEC-schemes to implement pre- and post-compromise security. This is a weaker notion than one could strive for, but most existing schemes would not instantiate a stronger notion which would therefore undermine our goal of allowing the implementer to choose freely which one to use. We provide a more formal definition and “PRP-SEEC” as a simple, exemplary instantiation in the full version [3].

2.3 Translating Patterns

Given the general description of PQNoise, it remains to be shown how we move from a Noise pattern to a PQNoise pattern, generically. While the translation of most steps is straightforward, there are two non-trivial cases that have to be handled with care. The first one is any instance of a Static-Ephemeral or Ephemeral-Static exchange, where the sending-party is the owner of the static key. In the DH case they immediately prove the identity of the sender (assuming the static key is uncorrupted), establishing authenticity right away. This does not work with KEMs, as the owner of the static public key cannot combine their secret key with their peer’s ephemeral public key. The obvious workaround of having their peer create and send the ciphertext (as in the case where the sending party is owner of the ephemeral key) is not equivalent as it does not yet confirm the authenticity of the owner of the static key. Instead, the owner has to send an additional key-confirmation message (i.e., if this was the last message, another AEAD ciphertext from the static key owner using a key derived from the encapsulated value is necessary). In effect this adds up to one roundtrip.

The second one is replacing a Static-Static exchange, as there is no direct equivalent for it. However, using and extending the technique from the previous paragraph we can create a workaround that achieves similar security. In Noise the Static-Static exchange establishes authenticity for both parties and confidentiality (assuming uncorrupted static keys). Sending encapsulations for both static KEMs would almost establish the same properties as long as we secure the coins used by the encapsulating party using SEEC with a static secret. The resulting shared secrets are unknown to the adversary as long as the static keys are uncompromised. Afterwards, a key confirmation is necessary by the initial sender. Given that this adds a full roundtrip and that *ss* is usually used to get an early shared secret before a roundtrip, it may sometimes be more reasonable to drop this combination entirely, using *se* and *es* for authenticity.

Everything else can usually stay as it is, with the exception that the transmission of the responder’s ephemeral public key can be dropped entirely. (This is under the assumption that the initiator sends an ephemeral public key before the responder does; otherwise the initiator’s ephemeral public key gets dropped. As it would delay the arrival at forward secrecy, we see little reason to deviate from that convention and are unaware of any proposals to use such patterns.)

This gives us the following recipe to translate patterns in a manner that we conjecture to preserve the security (“conjecture” because while we prove the security of the PQNoise patterns, we lack a generic proof of classical Noise to compare the results to):

Ephemeral-Ephemeral exchanges (*ee*) can be directly replaced by sending a ciphertext for the ephemeral KEM (*ekem*).

Ephemeral-Static exchanges (*es*) sent by the initiator and Static-Ephemeral exchanges (*se*) sent by the responder can be directly replaced by sending a ciphertext for the receiving parties KEM (*skem*).

Ephemeral-Static exchanges (*es*) sent by the responder and Static-Ephemeral exchanges (*se*) sent by the initiator are more complicated: When the initiator in Noise sends *se*, we replace this by the initiator finishing his current turn, following to which the responder sends *skem*, the initiator replies with a key-confirmation that may also contain the remaining operations given in the line of the original pattern. The same approach is used for *es* sent by the responder, with reversed roles.

Static-Static exchanges (*ss*) where the initiator is the original sender, are replaced as follows. The initiator sends *skem*, computing her coins using SEEC with a static secret and ends her turn. The responder responds with *skem*, the coins also obtained using SEEC with a static secret, and ends his term. Lastly the initiator has to send another message for key confirmation. If the responder is the original sender, roles are reversed.

After removing duplicate actions (usually multiple uses of *skem*), we conjecture the resulting pattern to achieve the same confidentiality, authenticity, integrity, anonymity and deniability as the original one; While we don’t further analyze the later three goals, the lack of a generic analysis of classical noise (which is out of scope for this work) prevents us from proving the first two. Our generic analysis does however show that PQNoise matches (or exceeds) the conjectured / proven security [19] for those original Noise-patterns, that have been analyzed in the *fACCE*-model. This may come at the disadvantage of only achieving that security at a later point in the interaction, due to the additional roundtrips.

One property that cannot be preserved by this translation-approach is that the ephemeral key-share of a party is used with both shares of their peer; Because of this the peer could be certain that the derived keys belong to the same ephemeral key using DH. This is no longer the case with KEMs since there is in general no way to derive useful information about the used ephemeral entropy from the ciphertext and reusing entropy may even introduce vulnerabilities. In our formal analysis this does not cause problems for confidentiality and authenticity. However, protocol designers who rely upon the dual-use of the ephemeral keys in Noise for other purposes need to be aware of this.

2.4 Fundamental Patterns

With the above recipe we can convert any Noise pattern into a PQNoise-pattern. The result of this transformation is however not always optimal in terms of roundtrips. For this reason we hand-picked a PQNoise-pattern to match each of the twelve fundamental classical Noise-patterns ($\{I, N, K, X\} \times \{N, K, X\}$). These PQNoise patterns are designed to not only achieve at least the same amount of confidentiality and authenticity, put to also do that as efficient as possible. (The equivalent security follows from the use of the same KEMs, for which our proofs show that each KEM will introduce some degree of security independent of the order of their use, but possibly at different protocol-stages.)

All of them ended up having a direct equivalent in classical Noise when looking beyond its own fundamental patterns from which they result as part of the generic translation: The IK and the KK patterns are equivalent to IKnoss and KKnoss [34]. All patterns that involve (non-early) transmitted keys are equivalent to the deferred patterns where the transmitted key sees deferred use (every “X” becomes “X1”), for example IX is equivalent to IX1 and XX is equivalent to X1X1. All other patterns are equivalent to their namesakes.

While some of these patterns require more roundtrips than their classical counterparts and may achieve certain degrees of security at a slightly later point, they all eventually end up achieving the same degree of security that was conjectured or proven [19] for their classical counterparts.

Noise also provides three non-interactive patterns ($\{N, K, X\}$). The authenticated K - and X -patterns cannot be translated into non-interactive versions of PQNoise, as the initiator cannot prove his identity in a non-interactive way using only KEMs. The unauthenticated N -pattern can however be translated trivially and essentially results in the standard KEM/DEM-construction. We note that our analysis applies to the N -pattern as well and therefore include it in the list of the thirteen fundamental PQNoise patterns.

We depict the interactive ones of these in Figure 2 and provide more detailed descriptions and a comparison with their Noise counterparts in the full version [3].

3 OVERVIEW OF THE FLEXIBLE ACCE FRAMEWORK

We analyze the security of PQNoise in the flexible authenticated and confidential channel establishment (fACCE) framework [19] which was developed for the analysis of Noise. Here we give a high-level overview of the cryptographic primitive fACCE, and define fACCE security, highlighting areas that we have modified for our specific setting of post-quantum channel-establishment protocols.

The main divergence between the original fACCE model and our version is how we structure and represent *freshness conditions*. These allow the protocol analyser to determine in which settings an attack is valid, i.e. after what set of compromises or adversary actions is the adversary considered to win the game. This is largely determined by a definition of when a protocol is supposed to achieve a certain security goal. In the original fACCE model, this was represented as a series of freshness *counters*, which captured confidentiality or authentication under certain types of attacks e.g. au^ρ defines when the party with role ρ authenticates itself, and thus

pqNN:	pqNK:	pqNX:
-> e	<- s	-> e
<- ekem	...	<- ekem, s
	-> skem, e	-> skem
	<- ekem	
pqKN:	pqKK:	pqKX:
-> s	-> s	-> s
...	<- s	...
-> e	...	-> e
<- ekem, skem	-> skem, e	<- ekem, skem, s
	<- ekem, skem	-> skem
pqXN:	pqXK:	pqXX:
-> e	<- s	-> e
<- ekem	...	<- ekem, s
-> s	-> skem, e	-> skem, s
<- skem	<- ekem	<- skem
	-> s	
	<- skem	
pqIN:	pqIK:	pqIX:
-> e, s	<- s	-> e, s
<- ekem, skem	...	<- ekem, skem, s
	-> skem, e, s	-> skem
	<- ekem, skem	

Figure 2: The interactive fundamental PQNoise patterns.

when it is considered a non-trivial attack that the adversary can inject or modify messages from the ρ -party. This resulted (in their full model) in ten counters, each capturing a specific type of attack and compromise paradigm.

We instead represent all combinations of secrets (long-term and ephemeral) for each session as rows in a *security table* (ST), with *authentication* and *confidentiality* columns. For each combination of secrets, we indicate in which stage(s) authentication and confidentiality hold if the adversary *has not* corrupted those secrets. This results in a simpler, more intuitive representation as it focuses on the natural question: *under any given compromise strategy, when does the protocol achieve (if at all) confidentiality and authenticity?* instead of requiring the reader (or protocol designer) to understand and interpret cryptographic history (e.g., the eck counter describes an adversary that can compromise either session’s ephemeral randomness). We also use copies of ST (which we denote the *freshness table* or FT) as a tool within our formalism, serving to simplify our freshness conditions: each session begins with a full FT, and whenever an adversary compromises a particular type of secret, the rows with that secret are removed from FT. An adversary that attempts to break the security of stages that are not associated with some combination of secrets are considered invalid as the result of trivial attacks. In addition to this, we make a small number of mostly aesthetic changes:

- We rename Enc and Dec as Send and Recv. We note that this better matches their semantics, as Send and Recv also transmit channel-establishment material, and potentially do

not perform encryption and decryption at all, depending on the protocol.

- We require that each Send operation increments the stage counter of the channel. The original fACCE model only incremented the stage counter when new (and increased) security properties are reached. This change ties the stage of the channel to its flow in the channel communication. As a result, we modify the definition of fACCE protocols to no longer output the stage counter ς when sending or receiving messages, as it is sufficient to count the messages between communicating parties.

We now turn to describing the fACCE primitive and security framework on a high-level, and give some additional insight into the changes made to the freshness conditions. The full model can be found in the full version [3].

fACCE Primitive Description. On a high-level, fACCE is a cryptographic protocol that both establishes a secure channel and provides authenticated and confidential communication between two parties. Eschewing a modular approach, channel establishment and payload transmission are handled by the same algorithms – where Send sends channel establishment information and (potentially encrypted) payload data, and Recv receives. These functions may also update the internal state of the sessions.

Definition 3.1 (Flexible ACCE). A flexible ACCE protocol fACCE is a tuple of four algorithms KGen, Init, Send, Recv associated with a long-term secret key space $\mathcal{LSK}$, a long-term public key space $\mathcal{LPK}$, an ephemeral secret key space $\mathcal{ESK}$ an ephemeral public key space $\mathcal{EPK}$, and a state space $\mathcal{ST}$. The definition of fACCE algorithms are as follows:

KGen $\rightarrow_{\S} (sk, pk)$ generates long-term keys where $sk \in \mathcal{LSK}$, $pk \in \mathcal{LPK}$. Note that this captures both long-term asymmetric key pairs, as well as potential long-term symmetric secrets (which we consider a part of sk).

Init(sk, ppk, ρ, ad) $\rightarrow_{\S} st$ initializes a session to begin communication, where sk (optionally) are the initiator's long-term secret keys, ppk (optionally) is the long-term public key of the intended session partner, $\rho \in \{i, r\}$ is the session's role (i.e., initiator or responder), ad is data associated with this session, and $sk \in \mathcal{LSK} \cup \{\perp\}$, $ppk \in \mathcal{LPK} \cup \{\perp\}$, $ad \in \{0, 1\}^*$, $st \in \mathcal{ST}$.

Send(sk, st, m) $\rightarrow_{\S} (st', c)$ continues the protocol execution in a session and takes message m to output new state st' , and messages c^1 , where $sk \in \mathcal{LSK} \cup \{\perp\}$, $st, st' \in \mathcal{ST}$, $m, c \in \{0, 1\}^*$. Note that Send may generate additional ephemeral key pairs $(epk, esk) \in \mathcal{EPK} \times \mathcal{ESK}^2$.

Recv(sk, st, c) $\rightarrow_{\S} (st', m)$ processes the protocol execution in a session triggered by c and outputs new state st' , and message m , where $sk \in \mathcal{LSK} \cup \{\perp\}$, $st \in \mathcal{ST}$, $st' \in \mathcal{ST} \cup \{\perp\}$, $m, c \in \{0, 1\}^*$. If $st' = \perp$ is output, then this denotes a rejection of this ciphertext.

We assume messages sent in fACCE are sent in a ping-pong fashion, i.e., the initiator sends a message to the responder, who replies

¹Note that messages here may consist of channel establishment data (such as keying material), encrypted payload data, or even plaintext payload data. In what follows, we refer to these generically as “ciphertexts”, even when sending plaintext data.

²In the security experiment, these are stored within state st

to the initiator, and so on. Multiple messages in a single flow are thus extensions of a single message. Each message monotonically increases the stage of the protocol, i.e., the first message sent from initiator to responder is stage one, the first message sent from responder to initiator is stage two, etc. This differs from the original fACCE, which only increments stages when achieving new security properties.

We define the correctness of an fACCE protocol in the full version [3]. Intuitively an fACCE protocol is correct if messages sent from the established channel were equally accepted by their partner.

Execution Environment. Here we describe (on a high-level) the execution environment for our fACCE security experiment. We consider a set of n_P parties each (potentially) maintaining a long-term key pair $\{(sk_1, pk_1), \dots, (sk_{n_P}, pk_{n_P})\}$, $(sk_i, pk_i) \in \mathcal{LSK} \times \mathcal{LPK}$. Each party can participate in up to n_S sessions, with each session potentially lasting n_T stages. Each session samples per-session randomness $rand$ used throughout the protocol execution. We denote both the set of variables that are specific for a session s of party i as well as the identifier of this session as π_i^s . Further details on the session state can be found in the full version [3].

Honest partnering is defined over the transcript sent between two sessions. Intuitively, a session has an honest partner if all ciphertexts the honest partner received were sent by the session (without modification) and vice versa, and at least one party received a ciphertext at least once. The full definition of honest partner can be found in the full version [3].

The fACCE model can capture authentication and confidentiality under various compromise paradigms, similar to the *levels* of authentication and confidentiality encoded by the original fACCE's various counters. We also highlight that this approach aligns with the typical structures of proofs of fACCE protocols – when one of the right-hand columns is not ∞ , this represents a case distinction in the proof. This proof structure is common in the analysis of authenticated key exchange protocols, especially those in the extended-Canetti-Krawczyk (eCK) model [28], such as the proofs of WireGuard [18] and PQWireGuard [23].

To facilitate the security game, the challenger maintains for each session π_i^s a set $\mathcal{S}_{\pi_i^s}$ that contains labels of all secrets that each session (and its honest partner) maintains – the long-term secret values sk_i, sk_j (both asymmetric and symmetric), all ephemeral secret values sampled during the n_T stages of the protocol execution $esk_s^1, esk_t^1, \dots, esk_s^{n_T}, esk_t^{n_T}$ and the state maintained during the protocol executions at each stage $st_s^1, st_t^1, \dots, st_s^{n_T}, st_t^{n_T}$. Thus $\mathcal{S}_{\pi_i^s} = (sk_i, sk_j, esk_s^1, esk_t^1, \dots, esk_s^{n_T}, esk_t^{n_T}, st_s^1, st_t^1, \dots, st_s^{n_T}, st_t^{n_T})$.

Each session in an fACCE experiment is associated with a four column freshness table FT (a copy of the original ST), with each element of the powerset of $\mathcal{S}_{\pi_i^s}$ (labels for each secret for itself and its honest partner) contained in the left column, and stage counters / tuples in the *Confidentiality*, *Authenticity of Initiator*, and *Authenticity of Responder* columns. The intuition here is that the table declares at which stages confidentiality and authenticity (for each role) are achieved under the assumption that the associated combinations of secrets have *not* been compromised by an attacker.

Consider the NK Noise Pattern ST displayed in Table 1 (see the full version [3] for the full protocol). In the table we denote

Table 1: The NK Noise Pattern (<-s \\. . . \\.>e, es \\.<-e, ee) and associated fACCE security table.

Secrets	Conf	Auth - i	Auth - r
$s_{\mathcal{R}}$	∞	∞	∞
$s_{\mathcal{R}}, e_{\mathcal{J}}$	1	∞	2
$s_{\mathcal{R}}, e_{\mathcal{R}}$	∞	∞	∞
$e_{\mathcal{J}}, e_{\mathcal{R}}$	2	∞	∞
$s_{\mathcal{R}}, e_{\mathcal{J}}, e_{\mathcal{R}}$	1	∞	2

the ephemeral Diffie-Hellman secret value that the initiator samples as $e_{\mathcal{J}}$ and the responder samples as $e_{\mathcal{R}}$, and the long-term Diffie-Hellman secret value that the responder maintains (B) as $s_{\mathcal{R}}$. If (at least) the long-term key of the responder $s_{\mathcal{R}}$ and the ephemeral key of the initiator $e_{\mathcal{J}}$ remain uncompromised the NK Pattern achieves responder authentication in stage $\varsigma = 2$, and does not achieve initiator authentication. If $e_{\mathcal{R}}, e_{\mathcal{J}}$ remain uncompromised, NK achieves confidentiality in stage $\varsigma = 2$, and if $s_{\mathcal{R}}$ and $e_{\mathcal{J}}$ remain uncompromised then NK achieves confidentiality of messages in stage $\varsigma = 1$. The intuition on an attacker's winning condition is that if the adversary breaks security in any stages associated with a particular combination of secrets that have not been compromised, the adversary wins.

Adversarial Model. In order to model active attacks in our environment, the security experiment provides the OInit , OSend , ORecv oracles to an adversary $\mathcal{A}$, who can use them to control communication among sessions, together with the oracles OCorrupt , OReveal and ORevealRandomness .

Following the direction of the original fACCE work, we treat the authentication and confidentiality properties similarly to the original AEAD notion of Rogaway [35]: the game maintains a win flag (to indicate whether the adversary broke authenticity or integrity of ciphertexts) and changes encryption behaviour based on randomly sampled challenge bits (to model indistinguishability of ciphertexts). In order to win the security game, adversary $\mathcal{A}$ either has to trigger win $\leftarrow 1$ or output the correct challenge bit $\pi_i^s.b_{\varsigma}$ of a specific session stage ς at the end of the game.

In addition, the challenger maintains a set of freshness flags $\pi_i^s.fr_{\varsigma}$ for each stage ς of each session π_i^s . When $\mathcal{A}$ makes a query to OCorrupt , OReveal or ORevealRandomness , then $\mathcal{C}$ deletes all rows in the freshness table FT that contain the secret revealed to $\mathcal{A}$. All stages for all sessions that are not an element of the right-hand columns are now considered un-fresh, and the corresponding freshness flags are set to 0. When $\mathcal{A}$ terminates and outputs a session π_i^s and a stage counter ς such that the freshness flag associated with $\pi_i^s.\varsigma$ is 0, then $\mathcal{C}$ simply outputs a random bit b^* instead of $\pi_i^s.b_{\varsigma} = b'$.

We describe the function of each oracle below. The details on excluding trivial attacks as the result of these oracles can be found in the full version [3].

$\text{OInit}(i, pk_j, \rho, ad)$ initializes a new session π_i^s (if not yet initialized) of party i to be partnered with party j , invoking $\text{fACCE.Init}(sk_i, pk_j, \rho, ad) \rightarrow_{[\pi_i^s.rand]} \pi_i^s.st$ using randomness $\pi_i^s.rand$ and returning the index of the session s .

$\text{OSend}(i, s, m_0, m_1)$ triggers the encryption of the message m_b where $b = \pi_i^s.b_{\varsigma}$ by invoking $\text{Send}(sk_i, \pi_i^s.st, m_b) \rightarrow_{[\pi_i^s.rand]} (st', c)$ for an initialized π_i^s if $|m_0| = |m_1|$. Note that c contains both the explicit ciphertext encryption of the message m_b and any channel establishment messages that are sent in this stage. Finally c is appended to $\pi_i^s.T_s$.

$\text{ORecv}(i, s, c)$ triggers invocation of $\text{Recv}(sk_i, \pi_i^s.st, c) \rightarrow_{[\pi_i^s.rand]} (st', m)$ for an initialized π_i^s and returns (m, ς) only if π_i^s has no honest partner, and returns ς if an honest partner exists. If an honest partner exists, and the session is currently fresh, then outputting the plaintext message m would leak the challenge bit, so we must prevent this leakage. The adversary breaks authentication (and thereby win $\leftarrow 1$ is set) if the received ciphertext was not sent by a session of the intended partner but was successfully received (i.e., there exists no honest partner and the output state is $st' \neq \perp$), and $\mathcal{A}$ has not issued queries that trivially break authentication in this stage. Finally c is appended to $\pi_i^s.T_r$ if decryption succeeds.

$\text{ORevealRandomness}(i, s) \rightarrow rand$ outputs the ephemeral randomness $rand$ sampled by session π_i^s . The freshness table FT and freshness flags are updated by the challenger.

$\text{OCorrupt}(i) \rightarrow sk_i$ outputs the long-term secret key sk_i of party i and updates the freshness table FT and freshness flags.

$\text{OReveal}(i, s) \rightarrow \pi_i^s.st$ outputs the current session state $\pi_i^s.st$, and updates the freshness table FT and freshness flags.

Finally, we formalise the security of an fACCE primitive in the full version [3]. A flexible ACCE protocol fACCE is post-quantum secure if it is correct and $\text{Adv}_Q^{\text{fACCE}}$ is negligible for all quantum algorithms $\mathcal{Q}$ running in polynomial-time.

4 ANALYSIS

In this section we present our security analysis of PQNoise. To begin, we model Noise's use of key-derivation-functions as a "hash-object". This allows us to separate the analysis of Noise into the analysis of the hash-object, which focuses on the local key derivation activities of a user, and the analysis of the key exchange executed between users.

We note that besides the key-derivation-chain ("key-chain") Noise also computes a second hash-chain h to create a handshake-hash; the modelling and analysis in the following section do not apply to that chain.

4.1 Hash-Object

Noise has a somewhat convoluted key derivation process as it derives fresh symmetric keys every time it computes a new shared key. Towards this end, Noise makes use of a key-chain into which all shared secrets are absorbed and from which all session-keys are extracted. This chain effectively is a PRF chain in which a previous chaining value is used as key, and any new input is used as input. The output is split into an output and a new chaining value. In an analysis this can be treated as a series of independent pseudorandom function calls. However, the proofs that result from this approach tend to have a long sequence of game hops applying the dual-PRF assumption to replace PRF outputs by random values based on the chaining value or the input being pseudorandom.

These are shared by many proof-steps and distract from the core part of the different proofs. Because of this, we introduce an abstraction that allows us to treat such chains as a single object with new security properties that allow to prove security of protocols like (PQ)Noise. We call the new object a hash-object, provide a definition of pseudorandomness for such objects, and prove that the construction of a hash-object used in Noise achieves this property.

Noise usually creates multiple outputs whenever it inputs a new value into its hash-chain, the first of which is usually used as a form of a state that we model as the state s of our hash-object. At the end of the handshake-phase Noise uses the first result directly as output and forgoes the creation of a new state. To model this we introduced a function `finalize` that mostly behaves like the regular input-function, except that it does not return a new state.

Definition 4.1 (Hash-Object). Formally a hash-object is a tuple of three deterministic algorithms: `create`, `input`, `finalize`, and an integer-constant n .

`create`(1^λ) $\rightarrow s$ takes a security-parameter λ and returns a state s .

`input`(s, m) $\rightarrow s', h$ takes a state s and message $m \in \{0, 1\}^*$ and returns a new state s' and a list $h \in (\{0, 1\}^\lambda)^n$ of hashes of length n .

`finalize`(s, m) $\rightarrow h$ works like `input`, except that it does not return a state.

For convenience sake we will use class-style notation (i.e. $h := s.\text{input}(m)$ instead of $s, h := \text{input}(s, m)$).

Definition 4.2 (Pseudorandom Hash-Object). We say that a hash-object HO is a pseudorandom hash-object if and only if $\forall \mathcal{A} \in \text{QPT}, \lambda \in \mathbb{N}$:

$$\left| \Pr \left[\text{Exp}_{\text{HO}, \mathcal{A}, 0}^{\text{PRHO}}(1^\lambda) = 1 \right] - \Pr \left[\text{Exp}_{\text{HO}, \mathcal{A}, 1}^{\text{PRHO}}(1^\lambda) = 1 \right] \right| =: \text{Adv}_{\text{HO}, \mathcal{A}}^{\text{PRHO}}(1^\lambda) \leq \text{negl}(\lambda) \text{ where } \text{Exp}_{\text{HO}, \mathcal{A}}^{\text{PRHO}} \text{ is defined as in Experiment 1.}$$

The core idea behind this definition is that the adversary receives oracle-access to an arbitrary number of hash-objects into which he can feed whatever values he likes. At any point in time he can request to add the random secret r that is sampled once at the start of the game to any oracle by invoking `Rand`. From that point onwards all the outputs from the randomized hash-object will either be true random values or real, depending on the challenge-bit. Everything else in this definition is just there to prevent trivial attacks: *history* keeps track of the exact queries performed on each hash-object. *queries* is a dictionary that saves the set of queries that were previously performed on hash-objects with a given history to prevent running both `In` and `Fin` on objects in the same state, as the later would reveal the resulting state of the former. *cache* is a dictionary that is used to ensure that two hash-objects with the same history always return the same results even if they have been randomized and return truly random values.

With this we define the Noise Hash Object as depicted in Algorithm 2. This is more or less a direct recreation of how Noise defines HKDF, except that it distinguishes the case where the first argument is then used as state from the case where no state is maintained and everything is returned as output.

THEOREM 4.3. A Noise Hash Object NHO is a secure pseudo-random Hash-Object if HMAC-HASH is a dual-prf with: $\text{Adv}_{\text{NHO}, \mathcal{A}, q_i}^{\text{PRHO}}(1^\lambda) \leq \left(\text{Adv}_{\text{HMAC-HASH}, \mathcal{A}'}^{\text{CollRes}}(1^\lambda) + \text{Adv}_{\text{HMAC-HASH}, \mathcal{A}'}^{\text{PRF-SWAP}}(1^\lambda) + (2 \cdot q) \cdot \text{Adv}_{\text{HMAC-HASH}, \mathcal{A}'}^{\text{PRF}}(1^\lambda) \right)$ where q refers to the total number of oracle-queries.

We refer to the full version [3] for a proof. Intuitively the collision-resistance of HMAC-HASH implies that only identical histories result in equal states and the HMAC-HASH being a dual-PRF (the full version [3] for a definition) ensures that once r has been added to a chain, its first state becomes pseudorandom which is retained upon subsequent calls.

4.2 PQNoise

At this point we can now start the analysis of PQNoise itself. We consider PQNoise with and without the use of SEEC. The reason for analyzing both is that Noise has traditionally considered bad RNGs a problem of the operating system which combined with the fact that the use of SEEC is (if there is no corruption) unobservable from the outside, suggests that the Noise-project may refuse to specify the use of SEEC and leave it as an implementation-detail.

Let Π be a PQNoise-protocol and Π' be the same protocol without the use of SEEC. Let $\#I$, $\#R$ and $\#E$ refer to the stage of Π/Π' during which the KEM-ciphertexts for the initiator's/ responder's/ ephemeral public keys are sent and ∞ if they are not sent. Let n_P be the number of parties participating in a protocol, n_S be the maximum number of sessions a party participates in, and n_K the total number of session-keys that a party uses. We are using standard-definitions for AEAD, PRFs, PRF-SWAPs, and KEMs (see the full version [3]). On top of that we use the definition of pseudo-random hash-object (PRHO) from above.

Intuitively the following four theorems can be summarized like this: In PQ-Noise, authenticity for a party $\mathcal{P}$ is established once it sends a valid reply to a message that was encrypted with uncorrupted randomness under $\mathcal{P}$'s uncorrupted public key. This is because $\mathcal{P}$'s peer $\mathcal{U}$ is by the definition and requirements of this case an honest peer whose KEM-ciphertext is fresh and can only be decrypted by $\mathcal{P}$. $\mathcal{P}$'s response contains an AEAD-ciphertext whose key is derived from the shared secret that only $\mathcal{P}$ and $\mathcal{U}$ have access too. As AEAD-ciphertexts cannot be forged without the key, and $\mathcal{U}$ knows that the reply was not created by her, she can, by the corruption-setting, conclude that she is talking to $\mathcal{P}$.

THEOREM 4.4. Π achieves initiator-authenticity in stage $\#I + 1$ if the initiator's static key and either of the responders keys have not been corrupted with:

$$\text{Adv}_{\Pi, \mathcal{A}}^{\text{fACCE.auth}_j}(1^\lambda) \leq \text{Adv}_{\text{H}, \mathcal{A}'}^{\text{CollRes}}(1^\lambda) + \frac{4(n_P \cdot n_S)^2}{2^\lambda} + n_P^2 \cdot n_S \cdot (\text{Adv}_{\Sigma, \mathcal{A}'}^{\text{SEEC}}(1^\lambda) + \text{Adv}_{\text{IKEM}, \mathcal{A}', n_S}^{\text{IND-CCA}}(1^\lambda) + \text{Adv}_{\text{H}, \mathcal{A}'}^{\text{PRHO}}(1^\lambda) + n_K \cdot \text{Adv}_{\text{AEAD}, \mathcal{A}'}^{\text{auth}}(1^\lambda))$$

THEOREM 4.5. Π' achieves initiator-authenticity in stage $\#I + 1$ if the initiator's static key and the responders ephemeral key have not been corrupted with:

$$\text{Adv}_{\Pi, \mathcal{A}}^{\text{fACCE.auth}_j}(1^\lambda) \leq \text{Adv}_{\text{H}, \mathcal{A}'}^{\text{CollRes}}(1^\lambda) + \frac{4(n_P \cdot n_S)^2}{2^\lambda} + n_P^2 \cdot n_S \cdot (\text{Adv}_{\text{IKEM}, \mathcal{A}', n_S}^{\text{IND-CCA}}(1^\lambda) + \text{Adv}_{\text{H}, \mathcal{A}'}^{\text{PRHO}}(1^\lambda) + n_K \cdot \text{Adv}_{\text{AEAD}, \mathcal{A}'}^{\text{auth}}(1^\lambda))$$

Experiment 1: $\text{Exp}_{\text{HO}, \mathcal{A}, b}^{\text{PRHO}}$ the pseudo-randomness experiment for a hash-object HO.

```

1   $r \xleftarrow{\$} \{0, 1\}^\lambda$ 
2   $\text{hashes} := [], \text{history} := [], j := 0$ 
3   $\text{randomized} := [0, \dots, 0]$ 
4   $\text{finalized} := [0, \dots, 0]$ 
5   $\text{queries} := \emptyset$ 
6   $\text{cache} := \text{dict}()$ 
7  Oracle Create:
8  |  $i, j := j, j + 1$ 
9  |  $\text{hashes}[i] := \text{create}(1^\lambda)$ 
10 |  $\text{history}[i] := []$ 
11 | return  $i$ 
12 Oracle Rand( $i, \text{finalize}$ ):
13 |  $\text{randomized}[i] := 1$ 
14 | if  $\text{finalize}$ :
15 | | return  $\text{Fin}(i, r)$ 
16 | else:
17 | | return  $\text{In}(i, r)$ 

18 Oracle In( $i, m$ ):
19 | abort_if( $\text{finalized}[i]$ 
20 | |  $\vee (\text{history}[i] \parallel ("Fin", m)) \in \text{queries}$ )
21 |  $\text{history}[i].\text{append}("In", m)$ 
22 |  $\text{queries} \cup = \text{history}[i]$ 
23 | if  $\text{randomized}[i]$ :
24 | | if  $\text{cache}[\text{history}[i]] \neq \perp$ :
25 | | | return  $\text{cache}[\text{history}[i]]$ 
26 | |  $h_0 := \text{hashes}[i].\text{input}(m)$ 
27 | |  $h_1 \xleftarrow{\$} (\{0, 1\}^\lambda)^n$ 
28 | |  $\text{cache}[\text{history}[i]] := h_b$ 
29 | | return  $h_b$ 
30 | else:
31 | | return  $\text{hashes}[i].\text{input}(m)$ 

31 Oracle Fin( $i, m$ ):
32 | abort_if( $\text{finalized}[i]$ 
33 | |  $\vee (\text{history}[i] \parallel ("In", m)) \in \text{queries}$ )
34 |  $\text{finalized}[i] := 1$ 
35 |  $\text{history}[i].\text{append}("Fin", m)$ 
36 |  $\text{queries} \cup = \text{history}[i]$ 
37 | if  $\text{randomized}[i]$ :
38 | | if  $\text{cache}[\text{history}[i]] \neq \perp$ :
39 | | | return  $\text{cache}[\text{history}[i]]$ 
40 | |  $h_0 := \text{hashes}[i].\text{finalize}(m)$ 
41 | |  $h_1 \xleftarrow{\$} (\{0, 1\}^\lambda)^n$ 
42 | |  $\text{cache}[\text{history}[i]] := h_b$ 
43 | | return  $h_b$ 
44 | else:
45 | | return  $\text{hashes}[i].\text{finalize}(m)$ 

45 return  $\mathcal{A}^{\text{Create}, \text{In}, \text{Fin}, \text{Rand}}(1^\lambda)$ 

```

Algorithm 2: Noise-compatible instantiation for the pseudo-random hash-object.

```

1 Function  $\text{create}(1^\lambda)$ :
2 | return “”
3 Function  $\text{finalize}(\text{state}, m)$ :
4 |  $(h_0, [h_1, \dots, h_n]) := \text{input}(\text{state}, m)$ 
5 | return  $[h_0, \dots, h_{n-1}]$ 
6 Function  $\text{input}(\text{state}, m)$ :
7 |  $\text{tmp} := \text{HMAC-HASH}(\text{state}, m)$ 
8 |  $\text{last} := \text{“”}$ 
9 | for  $i \in \{0, \dots, n\}$ :
10 | |  $h_i := \text{HMAC-HASH}(\text{tmp}, \text{last} \parallel \text{byte}(i))$ 
11 | |  $\text{last} := h_i$ 
12 | return  $h_0, [h_1, \dots, h_n]$ 

```

THEOREM 4.6. Π achieves responder-authenticity in stage $\#R+1$ if the responders static key and either of the initiator’s keys have not been corrupted with:

$$\text{Adv}_{\Pi, \mathcal{A}}^{\text{fACCE.auth}_{\mathcal{R}}} (1^\lambda) \leq \text{Adv}_{\text{H}, \mathcal{A}'}^{\text{CollRes}} (1^\lambda) + \frac{4(n_P \cdot n_S)^2}{2^\lambda} + n_P^2 \cdot n_S \cdot (\text{Adv}_{\Sigma, \mathcal{A}'}^{\text{SEEC}} (1^\lambda) + \text{Adv}_{\text{RKEM}, \mathcal{A}', n_S}^{\text{IND-CCA}} (1^\lambda) + \text{Adv}_{\text{H}, \mathcal{A}'}^{\text{PRHO}} (1^\lambda) + n_K \cdot \text{Adv}_{\text{AEAD}, \mathcal{A}'}^{\text{auth}} (1^\lambda))$$

THEOREM 4.7. Π' achieves responder-authenticity in stage $\#R+1$ if the responders static key and the initiator’s ephemeral key have not been corrupted with:

$$\text{Adv}_{\Pi, \mathcal{A}}^{\text{fACCE.auth}_{\mathcal{R}}} (1^\lambda) \leq \text{Adv}_{\text{H}, \mathcal{A}'}^{\text{CollRes}} (1^\lambda) + \frac{4(n_P \cdot n_S)^2}{2^\lambda} + n_P^2 \cdot n_S \cdot (\text{Adv}_{\text{RKEM}, \mathcal{A}', n_S}^{\text{IND-CCA}} (1^\lambda) + \text{Adv}_{\text{H}, \mathcal{A}'}^{\text{PRHO}} (1^\lambda) + n_K \cdot \text{Adv}_{\text{AEAD}, \mathcal{A}'}^{\text{auth}} (1^\lambda))$$

We provide a detailed proof in the full version [3]. Intuitively the security is a consequence of the ciphertext for the initiator’s/responder’s static public key being generated with good randomness for an uncorrupted key. The resulting shared secret is then fed into the hash-object whose outputs can be treated as random from then on. Eventually the adversary would therefore have to break the authenticity of the AEAD-scheme in order to create a message as the key is essentially random at that point. The relatively low tightness is a consequence of having to guess the attacked session and parties.

Intuitively the next six theorems state that confidentiality is achieved once a KEM-ciphertext for an uncorrupted keypair is sent. As the stage in which these are sent depends on the pattern, the actual stage at which messages are confidential depends on it too and on the corruption in question. To get the first confidential stage, one has to pick the lowest stage of the applicable results given below (if the conditions for a result are unmet because of unacceptable corruption or non-use of the associated KEM, that stage is ∞). The first four of these theorems deal with uncorrupted static keys.

THEOREM 4.8. Π achieves confidentiality in stage $\#I$ if either of the responders and the static key of the initiator is uncorrupted and the responder is an honest party with: $\text{Adv}_{\Pi, \mathcal{A}}^{\text{fACCE.conf}_J} (1^\lambda) \leq \frac{4(n_P \cdot n_S)^2}{2^\lambda} + \text{Adv}_{\text{H}, \mathcal{A}'}^{\text{CollRes}} (1^\lambda) + n_P^2 \cdot n_S^2 \cdot (\text{Adv}_{\Sigma, \mathcal{A}'}^{\text{SEEC}} (1^\lambda) + \text{Adv}_{\text{IKEM}, \mathcal{A}', n_S}^{\text{IND-CCA}} (1^\lambda) + \text{Adv}_{\text{H}, \mathcal{A}'}^{\text{PRHO}} (1^\lambda) + (n_S - 1) \cdot \text{Adv}_{\text{AEAD}, \mathcal{A}'}^{\text{auth}} (1^\lambda) + n_K \cdot \text{Adv}_{\text{AEAD}, \mathcal{A}'}^{\text{INDs-CPA}} (1^\lambda))$

Where n_K is the total number of session-keys used in a session and n'_K refers to the total number of session keys that have been used before the key encapsulated in IKEM has been put into the hash-object.

THEOREM 4.9. Π' achieves confidentiality in stage #I if the responders ephemeral key and the static key of the initiator are uncorrupted and the responder is an honest party with: $\text{Adv}_{\Pi, \mathcal{A}}^{fACCE.conf_I}(1^\lambda) \leq \frac{4(n_P \cdot n_S)^2}{2^\lambda} + \text{Adv}_{H, \mathcal{A}'}^{CollRes}(1^\lambda) + n_P^2 \cdot n_S^2 \cdot (\text{Adv}_{IKEM, \mathcal{A}', n_S}^{IND-CCA}(1^\lambda) + \text{Adv}_{H, \mathcal{A}'}^{PRHO}(1^\lambda) + (n_S - 1) \cdot \text{Adv}_{AEAD, \mathcal{A}'}^{auth}(1^\lambda) + n_K \cdot \text{Adv}_{AEAD, \mathcal{A}'}^{INDS-CPA}(1^\lambda))$

Where n_K is the total number of session-keys used in a session and n'_K refers to the total number of session keys that have been used before the key encapsulated in IKEM has been put into the hash-object.

THEOREM 4.10. Π achieves confidentiality in stage #R if either of the initiators keys and the static key of the responder is uncorrupted and the initiator is an honest party with: $\text{Adv}_{\Pi, \mathcal{A}}^{fACCE.conf_R}(1^\lambda) \leq \frac{4(n_P \cdot n_S)^2}{2^\lambda} + \text{Adv}_{H, \mathcal{A}'}^{CollRes}(1^\lambda) + n_P^2 \cdot n_S^2 \cdot (\text{Adv}_{\Sigma, \mathcal{A}'}^{SEEC}(1^\lambda) + \text{Adv}_{RKEM, \mathcal{A}', n_S}^{IND-CCA}(1^\lambda) + \text{Adv}_{H, \mathcal{A}'}^{PRHO}(1^\lambda) + (n_S - 1) \cdot \text{Adv}_{AEAD, \mathcal{A}'}^{auth}(1^\lambda) + n_K \cdot \text{Adv}_{AEAD, \mathcal{A}'}^{INDS-CPA}(1^\lambda))$

Where n_K is the total number of session-keys used in a session and n'_K refers to the total number of session keys that have been used before the key encapsulated in RKEM has been put into the hash-object.

THEOREM 4.11. Π' achieves confidentiality in stage #R if the initiators ephemeral key and the static key of the responder are uncorrupted and the initiator is an honest party with: $\text{Adv}_{\Pi, \mathcal{A}}^{fACCE.conf_R}(1^\lambda) \leq \frac{4(n_P \cdot n_S)^2}{2^\lambda} + \text{Adv}_{H, \mathcal{A}'}^{CollRes}(1^\lambda) + n_P^2 \cdot n_S^2 \cdot (\text{Adv}_{RKEM, \mathcal{A}', n_S}^{IND-CCA}(1^\lambda) + \frac{4(n_P \cdot n_S)^2}{2^\lambda} + \text{Adv}_{H, \mathcal{A}'}^{PRHO}(1^\lambda) + (n_S - 1) \cdot \text{Adv}_{AEAD, \mathcal{A}'}^{auth}(1^\lambda) + n_K \cdot \text{Adv}_{AEAD, \mathcal{A}'}^{INDS-CPA}(1^\lambda))$

Where n_K is the total number of session-keys used in a session and n'_K refers to the total number of session keys that have been used before the key encapsulated in RKEM has been put into the hash-object.

The proofs are largely similar to the previous ones, the main-difference being that instead of relying on the unforgeability they now need to rely on the confidentiality (IND\$-CPA) of the AEAD-scheme. The increased loss in tightness is a result of having to guess the peer's session and the attacked AEAD-key on top of what the previous proofs had to guess. We provide them in the full version [3].

The last two of the six confidentiality theorems deal with uncorrupted ephemeral keys and are largely analogous to the previous four besides that.

THEOREM 4.12. Π achieves confidentiality in stage #E if both the initiator and the responder have at least one uncorrupted key and both are honest partners with:

$$\text{Adv}_{\Pi, \mathcal{A}}^{fACCE.conf_E}(1^\lambda) \leq \frac{4(n_P \cdot n_S)^2}{2^\lambda} + \text{Adv}_{H, \mathcal{A}'}^{CollRes}(1^\lambda) + n_P^2 \cdot n_S^2 \cdot (2 \cdot \text{Adv}_{\Sigma, \mathcal{A}'}^{SEEC}(1^\lambda) + \text{Adv}_{EKEM, \mathcal{A}', 1}^{IND-CCA}(1^\lambda) + \text{Adv}_{H, \mathcal{A}'}^{PRHO}(1^\lambda) + n_K \cdot \text{Adv}_{AEAD, \mathcal{A}'}^{INDS-CPA}(1^\lambda))$$

Where n_K is the total number of session-keys used in a session and n'_K refers to the total number of session keys that have been used before the key encapsulated in EKEM has been put into the hash-object.

THEOREM 4.13. Π' achieves confidentiality in stage #E if neither the initiators nor the responders ephemeral keys are uncorrupted and both are honest partners with:

$$\text{Adv}_{\Pi, \mathcal{A}}^{fACCE.conf_E}(1^\lambda) \leq \frac{4(n_P \cdot n_S)^2}{2^\lambda} + \text{Adv}_{H, \mathcal{A}'}^{CollRes}(1^\lambda) + n_P^2 \cdot n_S^2 \cdot (\text{Adv}_{EKEM, \mathcal{A}', 1}^{IND-CCA}(1^\lambda) + \text{Adv}_{H, \mathcal{A}'}^{PRHO}(1^\lambda) + n_K \cdot \text{Adv}_{AEAD, \mathcal{A}'}^{INDS-CPA}(1^\lambda))$$

Where n_K is the total number of session-keys used in a session and n'_K refers to the total number of session keys that have been used before the key encapsulated in EKEM has been put into the hash-object.

The proofs for these theorems are largely analogous to the one before, except that they require two applications of SEEC (one on either peer). We provide them in the full version [3].

Using these results we can then easily set up the fACCE-table for any given PQNoise protocol: The authenticity-results can be taken as they are. For confidentiality the lowest value that achieves security is the relevant one. For the fundamental PQNoise patterns this gives the results presented in Table 2. When compared with the partially conjectured results for Noise [19] (see the full version[3]) we see that PQNoise ends up with the same security as classical Noise eventually.

5 IMPLEMENTATION

We implement PQNoise as an extension of the “nyquist” implementation of Noise by Angel in the Go programming language [2]. As underlying instantiation of all KEMs we use Kyber-768 [4, 12]; specifically the highly optimized Go implementation in Cloudflare’s Circl library [27]. As Circl implements other post-quantum KEMs, including all parameter sets of Kyber, but also SIKE [5], it would be easy to change to a different instantiation of the KEMs.

When comparing performance between PQNoise handshakes and corresponding Noise handshakes, we notice that *computationally*, i.e., in terms of CPU cycles, PQNoise is more efficient. This is not surprising, because Kyber-768 is considerably faster than X25519-based DH key exchange in Noise. For example, on an Intel Xeon E-2124 (Coffee Lake) CPU, eBACS [8] reports 125 303 cycles for X25519 key generation and 135 390 cycles for X25519 shared-key computation; on the same CPU eBACS reports only 39 881 cycles for Kyber-768 key generation, 53 841 cycles for encapsulation, and 42 281 cycles for decapsulation. On other recent 64-bit CPUs the absolute numbers differ, but the big picture is similar: Kyber-768 outperforms X25519 in terms of cycle counts.

However, this advantage in computational performance does not mean that handshake times for PQNoise are faster than in Noise. In fact all cryptography used in Noise or in (our instantiation of) PQNoise is so fast that handshake times are largely determined by data transmission, and this is where two disadvantages of PQNoise kick in: first, post-quantum KEMs have much larger public keys and ciphertexts than (pre-quantum) ECDH. Second, in some scenarios KEM-based AKE requires more round trips to achieve the same security. In order to investigate how these two factors influence real-world handshake performance, we consider the KK and the XX handshake patterns from Noise together with their pqKK and pqXX counterparts from PQNoise. While both patterns eventually achieve mutual authentication, for the KK and pqKK patterns the amount of round trips is the same, while for the pqXX pattern an additional message from the responder is required compared to XX. We run benchmarks on a machine with two Intel Xeon Gold 6230 CPUs and 196 GB of RAM. The experimental setup is using the Linux kernel’s network-emulation features and

Table 2: Security of the fundamental PQNoise patterns. Values of the form x/∞ mean that the security is achieved in stage x if the party/parties that doesn't/don't use a static KEM still uses a static SEEC-key and never if that is not the case. We don't provide separate rows for authenticity without SEEC, as the $s_J, e_{\mathcal{R}}/e_J, s_{\mathcal{R}}$ -cases are identical, and the $s_J, s_{\mathcal{R}}$ -cases are trivially insecure and have those rows dropped entirely if SEEC is not used.

Security	Uncorr.	N	NN	NK	NX	KN	KK	KX	XN	XK	XX	IN	IK	IX
Confidentiality	$e_J, e_{\mathcal{R}}$	∞	2	2	2	2	2	2	2	2	2	2	2	2
	$e_J, s_{\mathcal{R}}$	1	2/ ∞	1	2	2/ ∞	1	2	2/ ∞	1	2	2/ ∞	1	2
	$s_J, e_{\mathcal{R}}$	∞	2/ ∞	2/ ∞	2/ ∞	2	2	2	2	2	2	2	2	2
	$s_J, s_{\mathcal{R}}$	1/ ∞	2/ ∞	1/ ∞	2/ ∞	2/ ∞	1	2	2/ ∞	1	2	2/ ∞	1	2
Confidentiality (Without SEEC)	$e_J, e_{\mathcal{R}}$	∞	2	2	2	2	2	2	2	2	2	2	2	2
	$e_J, s_{\mathcal{R}}$	1	∞	1	3	∞	1	3	∞	1	3	∞	1	3
	$s_J, e_{\mathcal{R}}$	∞	∞	∞	∞	2	2	2	4	4	4	2	2	2
Authenticity ($\mathcal{J}$)	$s_J, e_{\mathcal{R}}$	∞	∞	∞	∞	3	3	3	5	5	5	3	3	3
	$s_J, s_{\mathcal{R}}$	∞	∞	∞	∞	3/ ∞	3	3	5/ ∞	5	5	3/ ∞	3	3
Authenticity ($\mathcal{R}$)	$e_J, s_{\mathcal{R}}$	∞	∞	2	4	∞	2	4	∞	2	4	∞	2	4
	$s_J, s_{\mathcal{R}}$	∞	∞	2/ ∞	4/ ∞	∞	2	4	∞	2	4	∞	2	4

Table 3: Median handshake times in ms of KK and XX Noise patterns and their pqKK and pqXX counterparts in PQNoise.

	Fast network		Slow network	
	Init.	Resp.	Init.	Resp.
KK	16.35	0.42	98.73	0.41
pqKK	16.07	0.25	100.28	0.27
XX	16.02	16.1	98.47	98.6
pqXX	31.83	16.1	199.31	100.36

is largely following the setup used in [30] and [36]. For each of the patterns we take 1000 measurements of the time it takes to perform a handshake, independently for initiator and responder. We perform this measurements once over a fast network (1000 Mbit throughput, 31.1ms round-trip latency) and once over a slow network (10 Mbit throughput, 195.6ms round-trip latency). The results are listed in Table 3. The KK and pqKK responder times do not include any network communication – after receiving the first handshake message from the initiator, the responder can perform all computations without having to wait for a further message. These times thus show the computational advantage of Kyber-768 over X25519. We also see that increased message sizes in PQNoise do not have a major influence on performance in this TCP/IP-based scenario. What *does* matter is the additional protocol message in pqXX compared to XX: as expected, the initiator times are slower by pretty exactly the half-roundtrip network latency.

6 ACKNOWLEDGEMENT

We thank Trevor Perrin and Denisa Greconici for many helpful discussions.

This work has been supported by the Dutch Research Council (NWO) through VIDI grant No. VI.Vidi.193.066, by the European Research Council through Starting Grant No. 805031 (EPOQUE), and by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) as part of the Excellence Strategy of the German Federal and State Governments – EXC 2092 CASA - 390781972.

REFERENCES

- [1] Liliya R. Akhmetzyanova, Cas Cremers, Luke Garratt, Stanislav Smyslyayev, and Nick Sullivan. Limiting the impact of unreliable randomness in deployed security protocols. In Limin Jia and Ralf Küsters, editors, *CSF 2020 Computer Security Foundations Symposium*, pages 277–287. IEEE Computer Society Press, 2020.
- [2] Yawning Angel. nyquist - a Noise protocol framework implementation. <https://github.com/Yawning/nyquist>.
- [3] Yawning Angel, Benjamin Dowling, Andreas Hülsing, Peter Schwabe, and Florian Weber. Post Quantum Noise. Cryptology ePrint Archive, Paper 2022/539, 2022. <https://eprint.iacr.org/2022/539>.
- [4] Roberto Avanzi, Joppe Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, Gregor Seiler, and Damien Stehlé. Crystals-kyber (version 3.02) – submission to round 3 of the nist post-quantum project, 2021. <https://pq-crystals.org/kyber/data/kyber-specification-round3-20210804.pdf>.
- [5] Reza Azarderakhsh, Matthew Campagna, Craig Costello, Luca De Feo, Basil Hess, Aaron Hutchinson, Amir Jalali, Koray Karabina, Brian Kozziel, Brian LaMacchia, Patrick Longa, Michael Naehrig, Geovandro Pereira, Joost Renes, Vladimir Soukharev, and David Urbanik. Supersingular isogeny key encapsulation. Round-3 submission to the NIST PQC project, 2020. <https://siqe.org/#specification>.
- [6] Gustavo Banegas, Daniel J. Bernstein, Fabio Campos, Tung Chou, Tanja Lange, Michael Meyer, Benjamin Smith, and Jana Sotaková. CtIdh: faster constant-time csidh. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, (4):351–387, 2021. <https://tches.iacr.org/index.php/TCHES/article/view/9069>.
- [7] Daniel J. Bernstein. Curve25519: New Diffie-Hellman speed records. In Moti Yung, Yevgeniy Dodis, Aggelos Kiayias, and Tal Malkin, editors, *PKC 2006*, volume 3958 of *LNCS*, pages 207–228. Springer, Heidelberg, April 2006.
- [8] Daniel J. Bernstein and Tanja Lange. eBACS: ECRYPT benchmarking of cryptographic systems. <https://bench.cr.yp.to> (accessed 29 Sep 2021).
- [9] Daniel J. Bernstein, Tanja Lange, Chloe Martindale, and Lorenz Panny. Quantum circuits for the CSIDH: Optimizing quantum evaluation of isogenies. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part II*, volume 11477 of *LNCS*, pages 409–441. Springer, Heidelberg, May 2019.
- [10] Daniel J. Bernstein, Tanja Lange, and Ruben Niederhagen. Dual ec: A standardized back door. In *The New Codebreakers*, pages 256–281. Springer, 2016.
- [11] Xavier Bonnetain and André Schrottenloher. Quantum security analysis of CSIDH. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 493–522. Springer, Heidelberg, May 2020.
- [12] Joppe Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, and Damien Stehlé. CRYSTALS – Kyber: a CCA-secure module-lattice-based KEM. In *2018 IEEE European Symposium on Security and Privacy, EuroS&P 2018*, pages 353–367. IEEE, 2018. <https://cryptojedi.org/papers/#kyber>.
- [13] Ran Canetti and Hugo Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 453–474. Springer, Heidelberg, May 2001.
- [14] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. CSIDH: An efficient post-quantum commutative group action. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part III*, volume 11274 of *LNCS*,

- pages 395–427. Springer, Heidelberg, December 2018.
- [15] Cas Cremers, Luke Garratt, Stanislav V. Smyshlyaev, Nick Sullivan, and Christopher A. Wood. Randomness Improvements for Security Protocols. RFC 8937, October 2020.
 - [16] The Debian-Project. Debian Security Advisory – DSA-1571-1 openssl – predictable random number generator, May 2008. <https://www.debian.org/security/2008/dsa-1571>.
 - [17] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976.
 - [18] Benjamin Dowling and Kenneth G. Paterson. A cryptographic analysis of the WireGuard protocol. In Bart Preneel and Frederik Vercauteren, editors, *ACNS 18*, volume 10892 of *LNCS*, pages 3–21. Springer, Heidelberg, July 2018.
 - [19] Benjamin Dowling, Paul Rösler, and Jörg Schwenk. Flexible authenticated and confidential channel establishment (fACCE): Analyzing the noise protocol framework. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part I*, volume 12110 of *LNCS*, pages 341–373. Springer, Heidelberg, May 2020.
 - [20] Eduarda S.V. Freire, Dennis Hofheinz, Eike Kiltz, and Kenneth G. Paterson. Non-interactive key exchange. Cryptology ePrint Archive, Report 2012/732, 2012. <https://eprint.iacr.org/2012/732>.
 - [21] Atsushi Fujioka, Koutarou Suzuki, Keita Xagawa, and Kazuki Yoneyama. Strongly secure authenticated key exchange from factoring, codes, and lattices. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 467–484. Springer, Heidelberg, May 2012.
 - [22] Guillaume Girol, Lucca Hirschi, Ralf Sasse, Dennis Jackson, Cas Cremers, and David Basin. A spectral analysis of noise: a comprehensive, automated, formal analysis of diffie-hellman protocols. In *29th {USENIX} Security Symposium ({USENIX} Security 20)*, pages 1857–1874, 2020.
 - [23] Andreas Hülsing, Kai-Chun Ning, Peter Schwabe, Florian Weber, and Philip R. Zimmermann. Post-quantum WireGuard. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 304–321. IEEE Computer Society, 2021. <https://cryptojedi.org/papers/#pqwireguard>.
 - [24] Tibor Jäger, Florian Kohlar, Sven Schäge, and Jörg Schwenk. Authenticated confidential channel establishment and the security of TLS-DHE. *Journal of Cryptology*, 30(4):1276–1324, October 2017.
 - [25] Nadim Kobeissi, Georgio Nicolas, and Karthikeyan Bhargavan. Noise explorer: Fully automated modeling and verification for arbitrary noise protocols. In *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 356–370. IEEE, 2019.
 - [26] Hugo Krawczyk. HMQV: A high-performance secure Diffie-Hellman protocol. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 546–566. Springer, Heidelberg, August 2005.
 - [27] Kris Kwiatkowski and Armando Faz-Hernández. Introducing circl: An advanced cryptographic library. Posting in the Cloudflare Blog, 2019. <https://blog.cloudflare.com/introducing-circl/>.
 - [28] Brian A. LaMacchia, Kristin Lauter, and Anton Mityagin. Stronger security of authenticated key exchange. In Willy Susilo, Joseph K. Liu, and Yi Mu, editors, *ProvSec 2007*, volume 4784 of *LNCS*, pages 1–16. Springer, Heidelberg, November 2007.
 - [29] Nick Mooney. An Introduction to the Noise Protocol Framework, March 2020. <https://duo.com/labs/tech-notes/noise-protocol-framework-intro>.
 - [30] Christian Paquin, Douglas Stebila, and Goutam Tamvada. Benchmarking post-quantum cryptography in TLS. In Jintai Ding and Jean-Pierre Tillich, editors, *Post-Quantum Cryptography - 11th International Conference, PQCrypto 2020*, pages 72–91. Springer, Heidelberg, 2020.
 - [31] Chris Peikert. He gives C-sieves on the CSIDH. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 463–492. Springer, Heidelberg, May 2020.
 - [32] Trevor Perrin. Noise protocol framework. <https://noiseprotocol.org/noise.pdf> (Revision 34 vom 2018-07-11).
 - [33] Trevor Perrin. The Noise Protocol Framework, December 2017. https://media.ccc.de/v/34c3-9222-the_noise_protocol_framework.
 - [34] Trevor Perrin and Justin Cormack. Static-Static Pattern Modifiers for Noise, 2018. Revision 1, 2018-11-18, unofficial/unstable, https://github.com/noiseprotocol/noise_ss_spec.
 - [35] Phillip Rogaway. Authenticated-encryption with associated-data. In Vijayalakshmi Atluri, editor, *ACM CCS 2002*, pages 98–107. ACM Press, November 2002.
 - [36] Peter Schwabe, Douglas Stebila, and Thom Wiggers. Post-quantum TLS without handshake signatures. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security, CCS'20*, pages 1461–1480. ACM, 2020. <https://cryptojedi.org/papers/#kemtls>.
 - [37] Filippo Valsorda. Twitter-Survey on Crypto-Agility, April 2021. <https://twitter.com/FiloSottile/status/1386751406758105089>.