Many a Mickle Makes a Muckle: A Framework for Provably Quantum-Secure Hybrid Key Exchange

Benjamin Dowling[1], Torben Brandt Hansen[2], Kenneth G. Paterson[1]

[1] Department of Computer Science, ETH Zurich.
[2] Information Security Group, Royal Holloway, University of London.
benjamin.dowling@inf.ethz.ch, Torben.Hansen.2015@rhul.ac.uk,
kenny.paterson@inf.ethz.ch

Abstract. Hybrid Authenticated Key Exchange (AKE) protocols combine keying material from different sources (post-quantum, classical, and quantum key distribution (QKD)) to build protocols that are resilient to catastrophic failures of the different components. These failures may be due to advances in quantum computing, implementation vulnerabilities, or our evolving understanding of the quantum (and even classical) security of supposedly quantum-secure primitives. This hybrid approach is a prime candidate for initial deployment of post-quantum-secure cryptographic primitives because it hedges against undiscovered weaknesses. We propose a general framework HAKE for analysing the security of such hybrid AKE protocols. HAKE extends the classical Bellare-Rogaway model for AKE security to encompass forward security, post-compromise security, fine-grained compromise of different cryptographic components, and more. We use the framework to provide a security analysis of a new hybrid AKE protocol named Muckle. This protocol operates in one round trip and leverages the pre-established symmetric keys that are inherent to current QKD designs to provide message authentication, avoiding the need to use expensive post-quantum signature schemes. We provide an implementation of our Muckle protocol, instantiating our generic construction with classical and post-quantum Diffie-Hellman-based algorithmic choices. Finally, we report on benchmarking exercises against our implementation, examining its performance in terms of clock cycles, elapsed wall-time, and additional latency in both LAN and WAN settings.

Keywords: Authenticated key exchange, hybrid key exchange, provable security, protocol analysis, quantum key distribution, post-compromise security

1 Introduction

NIST's Post Quantum Cryptography (PQC) process has triggered significant effort into the design of new post-quantum public key algorithms that can eventually be used to replace existing algorithms in protocols such as IPsec and TLS. Indeed, NIST's 2017 call received 69 complete submissions in various categories.

However, much less attention has been paid on how to securely integrate these new algorithms into applications, and to assessing the impact they will have on the performance of real-world network protocols. A key issue is that the new algorithms are relatively immature, and our understanding of their security is still evolving. NIST lacked confidence in 13 of the original submissions [23]; meanwhile Albrecht et al. [4] highlight how poor our current understanding is of how to assess the cost of lattice attacks. During the cryptographic interregnum, sensitive data is still at risk from attackers who are willing to record and store network traffic for later cryptanalysis. One response to this uncertainty is to quickly roll out post-quantum secure algorithms in protocols like TLS. For example, in 2016 Google carried out an experiment in which they deployed the NewHope lattice-based scheme [5] in Chrome and in Google servers [14], and in 2019 Cloudflare and Google jointly carried out similar experiments deploying both lattice and supersingular isogeny-based schemes[21]. These tests adopted hybrid approaches, combining post-quantum schemes with forward-secure key exchange mechanisms, namely Elliptic Curve Diffie Hellman Ephemeral (ECDHE). Adopting a hybrid approach hedges against security vulnerabilities in the post-quantum algorithm (fundamental as well as implementation-related) whilst providing security against quantum adversaries. While discussions have started [28], at this point no formal standardisation has begun integrating post-quantum algorithms into secure Internet protocols, a few unadopted IETF drafts notwithstanding [26, 30]. Standardisation will inevitably be needed, and we anticipate that a hybrid approach will be used. But first the community needs to research a) how to build and analyse hybrid protocols, and b) how to assure the security of their post-quantum components. The former is the main focus of this work, while the latter falls under the aegis of the NIST PQC process.

Quantum Key Distribution (QKD) is often promoted as an alternate solution to the threat posed by large-scale quantum computers, and has some attractive features: when well-implemented, it can offer unconditional security, it is also increasingly well-integrated with standard optical communications and electronics systems, with small package sizes and high raw bit rates, cf. [27]. However, the achievable bit rate does not yet practically allow the use of QKD keying material in a one-time-pad encryption system, so while the keying material may be unconditionally secure, no practical overall secure communications system relying on QKD is (to date). Moreover, QKD is fundamentally range limited (in the absence of quantum repeaters) and so cannot offer true end-to-end security in wide-area networks. Furthermore the technology is still quite immature, and vulnerable to various implementation attacks ("quantum hacking"), cf. [29, 19]. Even the physical basis of QKD has been questioned [31, 10]. Despite this, QKD may still usefully augment existing technologies in point-to-point applications, such as intra or inter data-centre communications or in metropolitan networks. Given this context, we should consider the possibility of incorporating QKD-based keying material into our hybrid protocol designs, resulting in three sources of keying material to combine: classically-secure (e.g. ECDHE), post-quantum secure (e.g.

NewHope, SIDH, or another NIST candidate), and QKD-based. Having established this context, we can now begin to describe our contributions.

1.1 Our Contributions

The HAKE security framework: We introduce a flexible framework for capturing and analysing Hybrid Authenticated Key Exchange (AKE) protocols that combine a wide variety of symmetric and asymmetric primitives. The HAKE framework is the result of heavily modifying the classic Bellare-Rogaway [7] model for AKE, incorporating security notions such as perfect forward secrecy and post-compromise security (referring to the ability of a key exchange protocol to recover security in the event of a catastrophic compromise of all its secrets) and smoothly caters for different strengths of adversary (quantum or even classical). It features a particularly simple and novel abstraction of QKD protocols to allow them to be modelled in a standard computational setting: pairs of parties are given private access to a shared source of secret random bits.

The Muckle AKE protocol: To exercise the HAKE framework, we also present the Muckle AKE protocol,[3] its security analysis, details of a working software implementation of Muckle, and benchmarking results. Muckle securely combines keying material obtained from a quantum key distribution (QKD) protocol with that from a post-quantum-secure key encapsulation mechanism (KEM) and a classically-secure KEM. Muckle is a one-round (1-RTT) protocol which exploits the presence of a QKD component to simplify the authentication of protocol messages. Specifically, QKD protocols typically assume the presence of an initial or pre-shared key (PSK) between the pair of communicating parties. This is used to bootstrap an authenticated channel for exchanging basis measurement information.[4] Muckle's design assumes the presence of a second PSK (since the cost of establishing two such keys is not any greater in practice than the cost of establishing just one), and uses it as the basis for authenticating its protocol messages via MACs. Muckle evolves this key and associated state, updating them with material obtained from the post-quantum and classical primitives as well as the QKD itself, see Figure 2. Our approach avoids costly post-quantum secure signatures, but Muckle and its security analysis can be extended to rely on them instead of PSKs.

Benchmarking Muckle: We instantiate and implemented Muckle in 'C' (which we denote C-Muckle) and benchmarked it in different network settings, selecting specific schemes in order to fix a concrete design , in particular instantiating the post-quantum component with SIDH and the classical component with ECDH, but other algorithmic choices are of course possible. We profile the cost of the underlying C-Muckle functions in terms of the median execution wall-time and clock cycle counts. We also contrast the wall-time profiling

[3] The name Muckle derives from the traditional English phrase "Many a mickle makes a muckle": many small things can add up to make a big thing.

[4] As a side-note, this is why QKD in this normal form does not solve the key distribution problem, but only the key expansion problem.

of C-Muckle functions when it runs over a LAN with the same profiling when it is run between London and Paris (approximately 500km, somewhat more than the current maximum range of single-hop QKD systems). These experiments are done without a real QKD system, which is simulated via access to a file of keying material.

Security analysis of Muckle: Finally, we demonstrate that Muckle achieves AKE security as defined by our HAKE framework. This allows us to make security statements about Muckle in the presence of quantum adversaries (assuming post-quantum variants of standard cryptographic assumptions), or under the catastrophic failure of all but one of its distinct components. The latter includes scenarios where, for example, all public key cryptography evaporates (and only Muckle's QKD component remains secure). It also includes the situation where the QKD component turns out to be badly engineered and therefore insecure and where the classical component becomes vulnerable to a quantum computer, but where its post-quantum counterpart remains secure.

Organisation: For readability, we describe the Muckle protocol (Section 2) and its performance (Section 3) before presenting the HAKE security framework (Section 4) and then the formal security analysis of Muckle using the framework (Section 5). The paper closes with conclusions and directions for future work (Section 6).

1.2 Related Work

While the analysis of "fully classical" hybrid schemes have appeared in the past (for instance, work on combining multiple public-key encryption schemes [32]), little work has been done on combining post-quantum and classical cryptographic primitives. Bindel et al. [12] examine a variety of hybrid digital signature schemes in quantum and post-quantum settings. They also formalise the notion of *separability*, which captures the ability of an attacker to separate the hybrid scheme into its individual cryptographic components. Bindel et al. [11] is most closely related to our work, considering hybrid key exchange in a similar setting to our own, but is focussed on quantum-secure KEM combiners. Their setting and security model are less general than ours in some regards (our HAKE framework can accommodate KEMs, theirs is limited to KEMs), but considers a heirarchy of attackers depending on quantum-computing capability and quantum access to the protocol participants. In addition, their compromise paradigm is less fine-grained, considering only the compromise of long-term and session keys. Complementing our approach, Mosca et al. [24] analyse the security of the QKD protocol BB84 [8], using an AKE security model in the tradition of Bellare-Rogaway to formalise the protocol in their notation. They prove the security of BB84 in this security model, and their notions of keys output by the QKD protocol match our assumptions. The concept of *breakdown resilience* was introduced by Brendel et al. [15]; this concept considers the effect on overall protocol security of failures of individual cryptographic components. They also extend Bellare-Rogaway security models by providing an interface for an attacker to

break individual cryptographic components, similar to our approach of providing specific key exposure oracles. Earlier work by Bos et al. [13] considered the integration of the NewHope lattice-based scheme in TLS, but not in the sense of a hybrid scheme, and provided a security model and formal analysis specific to that setting. There have been a couple of recent IETF drafts [26, 30] describing hybrid approaches for TLS 1.3, but without any accompanying formal security analysis as far as we are aware.

2 The Muckle Protocol

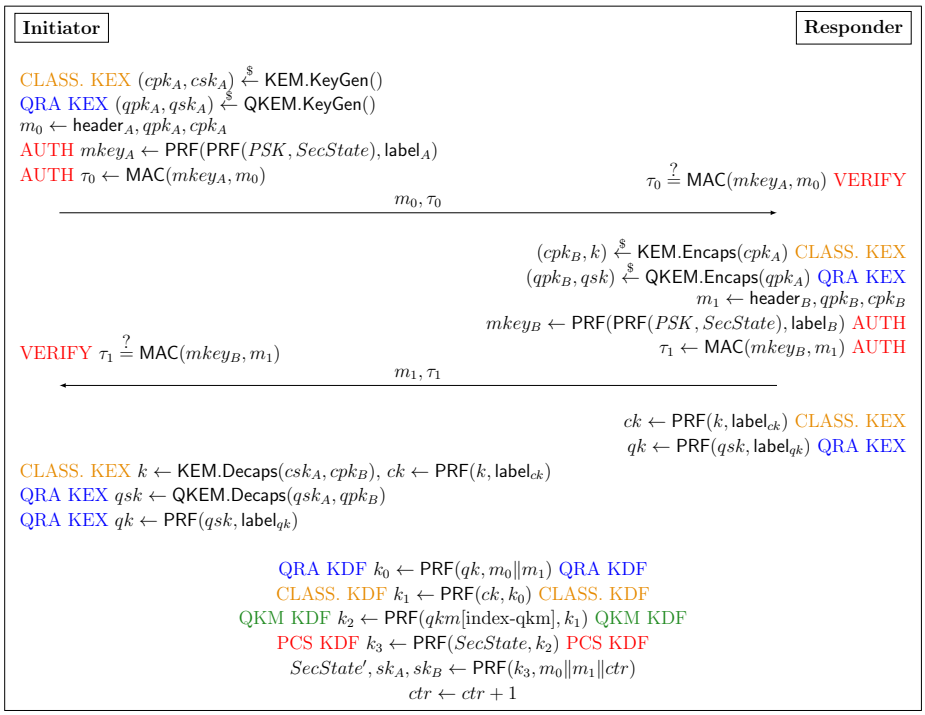

Fig. 1: A single stage of the Muckle protocol. The definitions of the KEM, PRF and MAC algorithms described here can be found in Appendix A.

Here we introduce the Muckle hybrid key exchange protocol; see Figure 1 for an overview. At a high-level, Muckle simultaneously executes post-quantum and classical key encapsulation primitives, and draws key material from a QKD protocol, represented abstractly in the protocol as a shared array of bits into which the two parties can index. The three distinct types of key material are used as inputs to a sequence of key derivation steps that we refer to as the Muckle key schedule, see Figure 2. The design of the Muckle key schedule allows us to prove

that the session keys produced by Muckle are resilient to vulnerabilities in the underlying QKD or key exchange primitives. Muckle is a multi-stage protocol, where the initiator and responder repeatedly run the single stage shown in Figure 1, updating the session keys sk_A, sk_B and the secret, shared state $SecState$ of the protocol at each stage. We highlight the key features of Muckle below:

- One round trip (1-RTT) to establish post-quantum-secure session keys.
- Multi-stage design and the inclusion of an updating secret state ($SecState$) allows Muckle to achieve post-compromise security, i.e. recover security after full compromise attacks.[5]
- Hybrid key exchange approach allows Muckle to be secure against classical adversaries even if the QKD and post-quantum components fail.
- Use of symmetric cryptography (of an appropriate key-length) allows Muckle to achieve post-quantum authentication without the use of computationally-expensive and bandwidth-intensive post-quantum signatures.
- Modular design allows implementers to easily replace underlying key exchange primitives if vulnerabilities are discovered.
- Key confirmation and full message transcript agreement of previous stages are provided in successive stages via the computation of authentication keys.

We expand on these below, and explain the different components of Muckle.

Message Structure: There are four elements to a Muckle message: a header (referred to in Figure 1 as header_A and header_B), containing message identifiers, cryptographic primitive identifiers and party identifiers; a classical ephemeral key encapsulation, (which we instantiate with elliptic-curve-based Diffie-Hellman (ECDH) in C-Muckle); a post-quantum ephemeral key encapsulation, (which we instantiate with Supersingular Isogeny-based Diffie-Hellman (SIDH) in C-Muckle); and a MAC tag computed over the message. Appendix E contains additional details.

QKD: Muckle assumes that a QKD scheme is running between pairs of communicating parties. QKD schemes make use of classically-authenticated communication channels, and such channels are (in practice) built using symmetric keys (though they could use other cryptographic techniques, such as digital signatures). Thus, Muckle assumes the presence of pre-shared symmetric keys (PSKs) between pairs of communicating parties. Likewise, this makes it possible to assume the existence of pre-established party identifiers in the protocol. Theses two value allow us to achieve post-quantum-secure authentication of the Muckle messages without incurring the significant computational or communication overhead that would be associated with a post-quantum signature scheme.

In our description of Muckle, we abstract the QKD protocol by modelling its output as an array of independent, uniformly-random bits (denoted $qsk[\cdot]$ in Figure 1) that is available to both parties in the protocol, otherwise treating the QKD component as a black box. Thus, we assume that the QKD system is implemented perfectly. This significantly simplifies our security analysis task, since it avoids the need for us to integrate existing QKD security models with

[5] Under certain restrictions, see Section 4.

our HAKE security framework. However, this is an idealisation that we plan to relax in future work, see Section 6 for more discussion. The pre-shared key is denoted PSK in Figure 1, and is 256 bits in size. The party identifiers are 32-byte strings (they do not appear explicitly in Figure 1, but instead are implicit in label_A and label_B).

Authentication: MAC tag computations use freshly-generated keys ($mkey_A$, $mkey_B$) for each new stage. Specifically, $mkey_A \leftarrow \mathsf{PRF}(\mathsf{PRF}(PSK, SecState), \mathsf{label}_A)$, and $mkey_B \leftarrow \mathsf{PRF}(\mathsf{PRF}(PSK, SecState), \mathsf{label}_B)$. Note that $SecState$ is updated with each new stage, and thus $mkey_A$ and $mkey_B$ are similarly fresh.

Key Schedule: The key schedule (see Figure 2) is run after the initiator and responder have sent and received their respective messages. The straightforward iterative design simplifies the analysis of the protocol. Each step takes as input some key material and a chaining key, and outputs a new chaining key used in the next iteration. We also include a counter ctr (a 256-bit integer) in the final PRF computation; ctr is incremented after each stage.

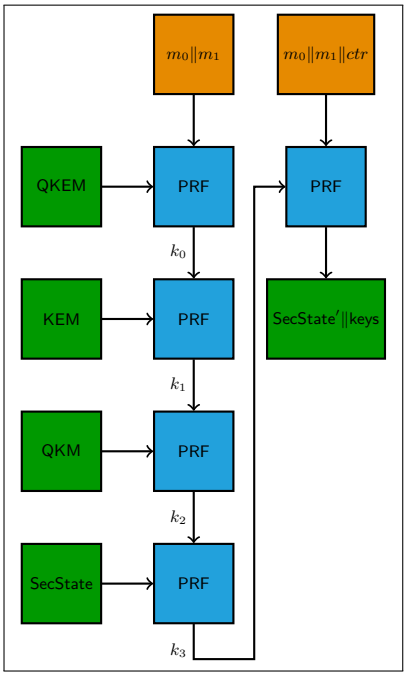

Fig. 2: A diagram of the Muckle key schedule. The PRF function (defined in Section A) takes input a key k, and some arbitrary-length bitstring in, outputs a fixed-length bitstring out. In our diagram, the left input is k, and the top input is in.

State Update: The secret state $SecState$ is updated at the end of a Muckle stage ,[6] when the session keys are computed. Specifically, $SecState, sk_A, sk_B \leftarrow \mathsf{PRF}(k_3, m_0 \| m_1 \| ctr)$, taking as input the final chain key k_3 from the key schedule, and the concatenation of the message transcript and counter: $m_0 \| m_1 \| ctr$. Thus, consecutive Muckle stages provide implicit key confirmation (since in order to derive the same $SecState$, protocol participants must also derive the same session keys sk_A, sk_B) of previous stages, as well as full message transcript agreement of previous stages.

Post-Compromise Security (PCS): At a high-level, PCS is the ability of a key exchange protocol to recover security when an attacker has compromised all secrets of a session. Obviously, this is impossible against an attacker that remains active, as it can continue to act as a Man-in-the-Middle (MITM) and using the long-term secrets, inject its own protocol messages. However, if during some stage of the protocol, the attacker becomes passive, it is possible to recover security (in the form of key indistinguishability) if the protocol participants honestly complete that stage. In addition, at that point, the adversary should no longer be able to act as an MITM. Our Muckle design achieves PCS by virtue of the inclusion of the secret state $SecState$ in the MAC computations and in the derivation of the session keys.

3 Implementation of Muckle

This section describes our reference implementation of Muckle in 'C', which we denote C-Muckle [2]. C-Muckle follows the same governing design principles as Muckle, favouring simplicity and verifiability. As a result, we optimise for readability and reproducibility and sacrifice features such as fully performance-optimised code.

3.1 Instantiation and Implementation

C-Muckle targets 128-bit post-quantum security. To instantiate C-Muckle and achieve this goal, we have made the following choice of parameters and cryptographic algorithms:

KEM: Ephemeral elliptic-curve Diffie-Hellman key exchange using the elliptic curve curve25519 [9].

QKEM: Supersingular isogeny Diffie-Hellman key exchange using field arithmetic over the prime p503, construction and parameters by Costello et. al. [17].

PRF: Pseudo-random function instantiated by the key derivation scheme HDKF [20] using 256-bit keys.

MAC: Message authentication code instantiated by HMAC [6] using 256-bit keys.

We describe how to generically instantiate key encapsulation mechanisms with Diffie-Hellman key exchange in Appendix B.

[6] In the first stage, $SecState$ is set to a constant, public value

Dependencies: To provide support for the chosen cryptographic components C-Muckle relies upon two libraries: mbedtls [1] version 2.13.0 and PQCrypto-SIDH [3] version 3. The former is used to support the ECDHE, PRF, and MAC cryptographic components as well as random number generation, while the latter is used to support SIDH.

Message Format: The C-Muckle message structure[7] is 476 bytes in length, identical in both directions, and consists of four fields. The first is the `header`, consisting of three sub-fields: `type`, `version` and `partyIdentifier`. The sub-fields `type` and `version` are both a single byte in length and indicate the direction of the message, and the underlying cryptographic primitives, respectively. One could think of the `version` sub-field as a ciphersuite indicator, similar to how a TLS handshake message indicates the underlying cryptographic components used in the handshake protocol. Currently, this field is set to a fixed value, and it is included for future extensibility. The last sub-field `partyIdentifier` is a 32-byte Muckle party identifier, a public string indicating the identity of the party sending the message. The next two fields of a C-Muckle message contain the ECDHE and SIDH public keys for either the initiator or responder, depending on the direction of the message. The ECDHE field has a size of 32 bytes while the SIDH field has a size of 378 bytes. The last field in a message contains the 32-byte MAC tag, computed over the first three fields of the message. These messages are transported over TCP/IP.

QKD Bits: Currently, our software implementation of C-Muckle does not engage with real QKD devices. The process of obtaining the bits produced by a QKD protocol is therefore emulated.

We provide two distinct methods for doing this. The first method is to store a static array of bits in the source code. During an execution, bits are read from the array depending on an index. The second method reads from a file, with bits similarly read from the file depending on an index. In both cases the bits should be uniformly random. The method of emulation can be changed during compile-time. Currently, C-Muckle defaults to using the static array method. These methods are solely implemented for experimental use and should not be used in any production system. C-Muckle is designed to allow easy switching to a method that provides true access to QKD key material.

3.2 Performance Study

Here we profile and discuss the performance of C-Muckle. Our experiments aim at conveying the cryptographic costs associated with the different components of Muckle, as well as the total cost of executing a complete run of the protocol. To achieve this, we benchmark different parts of C-Muckle as well as the core cryptographic API calls made to external libraries.

[7] For full details of the C-Muckle message, see Appendix E.

Methodology: We measure the performance of C-Muckle using two metrics: clock cycles and wall-time. For each metric, a single stage execution of C-Muckle is measured and recorded. The cost of lower layer functions responsible for performing cryptographic operations is also measured and separately recorded. Below we list these functions and describe the cryptographic operation they each perform:

muckle_ecdh_gen(): Generates an ECDHE public key pair.
muckle_ecdh_compute(): Computes the ECDHE secret.
muckle_sidh_gen(): Generates the SIDH public key pair.
muckle_sidh_compute(): Computes the SIDH shared secret.
muckle_read_qkd_keys(): Reads the QKD keying material using a method described above.
muckle_derive_keys(): Derives the secret state and session keys according to the key schedule defined in Section 2.

Note that the functions above perform more than just cryptographic operations. Additional operations include initialisation, copying between buffers, and general glue-code. We further discuss the overhead relative to the cryptographic operations for a subset of these functions.

Our experiments were performed between two Amazon Web Service (AWS) dedicated m5.large EC2 instances in two different availability zones (AZs) in the London Region. Each instance runs lLinux 4.14 with an Intel Xeon Platinum 8175M 2.5 Ghz CPU. We chose this relatively short distance between the initiator and responder to remain faithful to the practical restrictions on the deployment of Muckle. The QKD is an inherent part of the protocol, and deployment of a QKD network currently has a maximum distance of approximately 100km between nodes. For both metrics, the median over 100 samples is reported and each process is pinned to a single CPU.

To contrast running C-Muckle over this short distance with a more typical real-world setting, we performed the same experiment between two m5.large EC2 instances in two different regions, London and Paris, but only measuring the wall-time. The profiling results of this experiment can be found in Figure 3.

Wall-time Complete Execution: The complete execution time for a Muckle protocol run between two AZs is approximately 12.9ms. In comparison, the complete execution time between two regions is 26ms. The measurement scope is the execution of one entire stage of Muckle, including networking, initialising contexts, running clean up functions and executing general glue-code. By contrast, the round-trip-times for simple pings between two AZs and two regions were 0.745 ms and 8.224 ms.

Wall-time Function Profiling: Figure 4 provides a more granular view of the cost for specific functions in C-Muckle between two AZs. For the initiator, approximately 7.22ms is spent on various cryptographic function calls, with more than 68% of the 7.22ms spent performing SIDH-related operations. The same behaviour can be observed for the responder. The relative cost of cryptographic

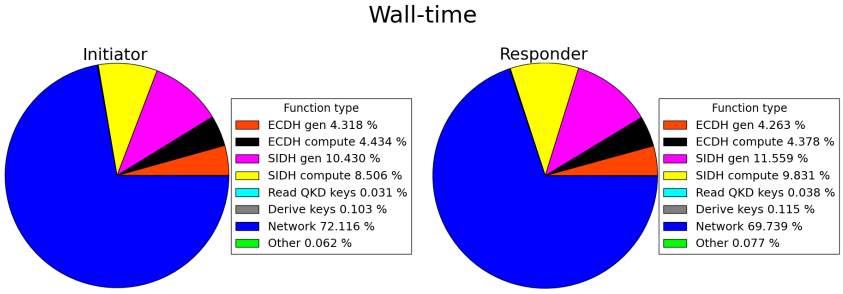

Fig. 3: Results of the wall-time measurement experiment between two AWS EC2 instances in two different regions (London and Paris). The top 6 categories for each chart are functions that correspond to C-Muckle functions described earlier. The network category includes time taken to initialise of the socket, as well as sending and receiving messages. The percentage for the Other category is computed by subtracting the median wall-time for the top 6 functions and the median time for networking from the entire median wall-time of the participant. (**Left**) C-Muckle initiator. (**Right**) C-Muckle responder.

operations in C-Muckle is therefore more than 65% when it is run over the short distances between AZs.

Clock Cycle Function Profiling: Table 1 contains an overview of the measured number of clock cycles for various functions. Each cell contains two functions: the first in each cell is the C-Muckle function described above, while the second is the function from the library dependencies,[8] used to implement the cryptographic operations in the C-Muckle function, i.e. during the execution of e.g. muckle_ecdh_gen() the function mbedtls_ecdh_gen_public() from the mbedtls library will be called. The table therefore highlights the absolute overhead of the cryptographic operations as implemented in C-Muckle compared to the core cryptographic operation supported via one of the two library dependencies. We have excluded the two functions muckle_read_qkd_keys() and muckle_derive_keys() because their cost is negligible relative to the total cost of the execution flow in C-Muckle.

The overhead in cryptographic C-Muckle functions relative to the corresponding external library functions is less than 15,000 clock cycles with one spike at 77,000 clock cycles. The overhead is predominately from copying buffers, initialisation and retrieving parameters. The "compute" functions also involve key derivation steps.

The number of clock cycles for the ECDHE mbedtls functions using the elliptic curve curve25519, are far from state-of-the-art. For example, Bernstein [9] reports a total of 832,457 clock cycles for both key generation and secret key computation. It should therefore be possible to significantly improve the ECDHE

[8] Either mbedtls or PQCrypto-SIDH

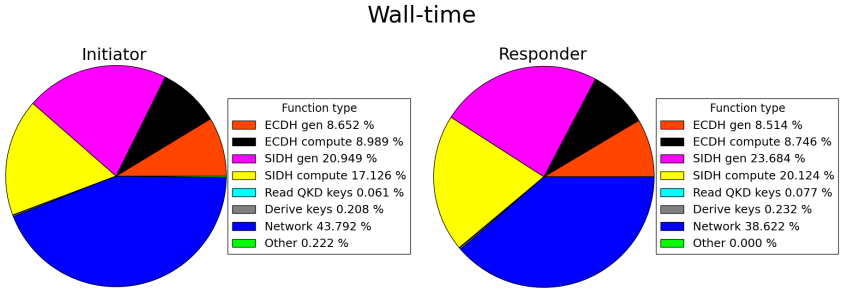

Fig. 4: Results of the wall-time measurement experiment between two AWS EC2 instances in two different availability zones located in the same region (London). Specifically, the chart captures the relative median wall-time spent executing various functions in the C-Muckle execution flow. The top 6 categories for each chart are functions that correspond to C-Muckle functions described in the text. The network category includes time taken to intialise of the socket, as well as sending and receiving messages. The percentage for the Other category is computed by subtracting the median wall-time for the top 6 functions and the median time for networking from the entire median wall-time of the participant. (**Left**) C-Muckle initiator. (**Right**) C-Muckle responder.

Function	**Clock cycles**
muckle_ecdh_gen()	2,769,893
mbedtls_ecdh_gen_public()	2,768,317
muckle_ecdh_compute()	2,875,367
mbedtls_ecdh_calc_secret()	2,846,614
initiator muckle_sidh_gen()	6,852,319
EphemeralKeyGeneration_A_SIDHp503()	6,775,268
initiator muckle_sidh_compute()	5,630,939
EphemeralSecretAgreement_A_SIDHp503()	5,613,257
responder muckle_sidh_gen()	7,531,586
EphemeralKeyGeneration_B_SIDHp503()	7,526,757
responder muckle_sidh_compute()	6,399,884
EphemeralSecretAgreement_B_SIDHp503()	6,391,934

Table 1: (**Left column**) The first function in each cell is a C-Muckle function described in the text. The second in each cell is the function from the library dependency, used to implement the C-Muckle function. The functions prefixed with mbedtls are from the mbedtls library, otherwise they are from the PQCrypto-SIDH library. (**Right column**) The median number of clock cycles over 100 samples.

performance in C-Muckle using a different library to mbedtls. However, we have found the mbedtls library to be easier to work with than other available libraries (like OpenSSL).

4 Hybrid Security Framework

Here we introduce our multi-stage hybrid authenticated key exchange (AKE) security framework HAKE for the analysis of our new protocol. HAKE follows the tradition of standard Bellare-Rogaway-based AKE models, and cleanly captures adversaries of differing strength (quantum and classical) via a fine-grained key compromise interface. Specifically, we model quantum adversaries by allowing them to corrupt non-post-quantum key exchange mechanisms (for instance, discrete logarithm-based key exchange algorithms). We highlight that our HAKE framework is flexible, and extends beyond Muckle, as HAKE captures (for example) the use of long-term asymmetric secrets, which are not used within Muckle. This allows HAKE to capture a variety of hybrid schemes, and is not simply to restricted to the use case of Muckle. We explain the HAKE framework in Section 4.2 (and give an algorithmic description of the security model in Figure 5 of Appendix C), and describe the corruption abilities of the adversary in Section 4.3. We then describe cleanness and partnering definitions in Section 4.4 as well as Section 4.5.

4.1 Secret Key Generation

HAKE addresses secret key generation (the output of a "KeyGen" algorithm) of individual key exchange components explicitly, and categorises them into *long-term* (i.e. generated once and used in every execution of the protocol), and *ephemeral* (i.e. generated on a per-stage basis) secret generation. We further divide these into the following sub-categories:

- *Post-quantum asymmetric secret generation.* The generation of a public-key pair for post-quantum code-based signature schemes is an example of a long-term variant. We denote the algorithm that generates these secrets as LQKeyGen. An algorithm that generates SIDH public-key pairs is an example of an ephemeral variant, which we denote as EQKeyGen.
- *Classical asymmetric secrets.* An algorithm that generates long-term RSA public-key pairs for signatures (that do not offer post-quantum security) would be denoted via LCKeyGen. Similarly, the generation of ECDHE public-key pairs would be done via ECKeyGen.
- *Symmetric secrets.* Long-term preshared secret keys would be generated via LSKeyGen, while (for instance), we consider that the ephemeral keying material generated by a quantum key distribution protocol to be captured as a ephemeral symmetric secret generation algorithm, which we denote ESKeyGen.

With this context, we now formally define the HAKE execution environment, capturing how an adversary can interact with a hybrid AKE protocol. HAKE as currently defined specifies only a single protocol of each category, i.e. one ephemeral classical key exchange primitive, one ephemeral post-quantum key exchange primitive, but we note that one can generalise the number of components of each type, allowing more flexibility in how HAKE protocols are constructed.

4.2 Execution Environment

Consider an experiment $\mathsf{Exp}^{\mathsf{HAKE},\mathsf{clean},\mathcal{A}}_{\Pi,n_P,n_S,n_T}(\lambda)$ played between a challenger $\mathcal{C}$ and an adversary $\mathcal{A}$. $\mathcal{C}$ maintains a set of n_P parties $P_1,\ldots,P_{n_P}$ (representing users interacting with each other in protocol executions), each capable of running up to (potentially parallel) n_S sessions of a probabilistic key-exchange protocol Π. Each session can consist of up to n_T consecutive stages, each an execution of the key-exchange protocol Π, represented as a tuple of algorithms $\Pi = (f, \mathsf{EQKeyGen}, \mathsf{ECKeyGen}, \mathsf{ESKeyGen}, \mathsf{LQKeyGen}, \mathsf{LCKeyGen}, \mathsf{LSKeyGen})$. We use π_i^s to refer to both the identifier of the s-th instance of the Π being run by party P_i and the collection of per-session variables maintained for the s-th instance of Π run by P_i, and f is a algorithm capturing the honest execution of the protocol Π by protocol participants. Due to space restrictions, we give the full list of algorithms in Appendix C, but describe generically the algorithms below:

$\Pi.f(\lambda, \boldsymbol{pk}_i, \boldsymbol{sk}_i, \boldsymbol{pskid}_i, \boldsymbol{psk}_i, \pi, m) \xrightarrow{\$} (m', \pi')$ is a (potentially) probabilistic algorithm that takes a security parameter λ, the set of long-term asymmetric key pairs $\boldsymbol{pk}_i, \boldsymbol{sk}_i$ of the party P_i, a collection of per-session variables π and an arbitrary bit string $m \in \{0,1\}^* \cup \{\emptyset\}$. f outputs a response $m' \in \{0,1\}^* \cup \{\emptyset\}$ and an updated per-session state π', behaving as an honest protocol implementation.

We describe a set of algorithms $\Pi.\mathsf{XYKeyGen}(\lambda) \xrightarrow{\$} (pk, sk)$, where $X \in \{\mathsf{E},\mathsf{L}\}$ and $Y \in \{\mathsf{C},\mathsf{Q}\}$. $\Pi.\mathsf{XYKeyGen}$ is a probabilistic *post-quantum ephemeral* (if $\mathsf{XY} = \mathsf{EQ}$), *post-quantum long-term* (if $\mathsf{XY} = \mathsf{LQ}$), *classic ephemeral* (if $\mathsf{XY} = \mathsf{EC}$), or *classic long-term* (if $\mathsf{XY} = \mathsf{LC}$) asymmetric keygen algorithm, taking a security parameter λ and outputting a public-key/secret-key pair (pk, sk).

We describe a set of algorithms $\Pi.\mathsf{ZSKeyGen}(\lambda) \xrightarrow{\$} (psk, pskid)$, where $Z \in \{\mathsf{E},\mathsf{L}\}$. $\Pi.\mathsf{ZSKeyGen}$ is a probabilistic *ephemeral* (if $\mathsf{Z} = \mathsf{E}$), or *long-term* (if $\mathsf{Z} = \mathsf{L}$) symmetric key generation algorithm taking as input a security parameter λ and outputting some symmetric keying material and (potentially) a keying material identifier $(psk, pskid)$, (or $(qkm, qkmid)$, respectively).

$\mathcal{C}$ runs $\Pi.\mathsf{LQKeyGen}(\lambda)$, $\Pi.\mathsf{LCKeyGen}(\lambda)$ and $\Pi.\mathsf{LSKeyGen}(\lambda)$ n_P times to generate asymmetric post-quantum and classical key pairs (which we denote with $\boldsymbol{pk}_i, \boldsymbol{sk}_i$) for each party $P_i \in \{P_1,\ldots,P_{n_P}\}$ as well as a symmetric keys and identifier $(\boldsymbol{psk}, \boldsymbol{pskid})$ and delivers all public-keys $\boldsymbol{pk}_i$, $\boldsymbol{pskid}$ for $i \in \{1,\ldots,n_P\}$ to $\mathcal{A}$. The challenger $\mathcal{C}$ then randomly samples a bit $b \xleftarrow{\$} \{0,1\}$ and interacts with the adversary via the queries listed in Section 4.3, also maintaing a set of corruption registers, described in Appendix C. Eventually, $\mathcal{A}$ terminates and outputs a guess d of the challenger bit b. The adversary wins the HAKE key-indistinguishability experiment if $d = b$, and additionally if the test session π satisfies a cleanness predicate clean, which we discuss in more detail in Section 4.5. We give an algorithmic description of this experiment in Figure 5 below. Each session maintains a set of per-session variables:

$\rho \in \{\texttt{init}, \texttt{resp}\}$**:** The role of the party in the current session. Note that parties can be directed to act as `init` or `resp` in concurrent or subsequent sessions.

$pid \in \{1, \ldots, n_P, \star\}$**:** The intended communication partner, represented with $\star$ if unspecified. Note that the identity of the partner session may be set during the protocol execution, in which case *pid* can be updated once.

$stid \in [n_T]$**:** The current (or most recently completed) stage of the session.

$\alpha \in \{\texttt{active}, \texttt{accept}, \texttt{reject}, \bot\}$**:** The status of the session, initialised with $\bot$.

$\mathbf{m}_i[stid] \in \{0,1\}^* \cup \{\bot\}$**, where** $\mathbf{i} \in \{\mathbf{s}, \mathbf{r}\}$**:** An array of the concatenation of messages sent (if $\mathbf{i} = \mathbf{s}$) or received (if $\mathbf{i} = \mathbf{r}$) by the session in each stage. Initialised by $\bot$ and indexed by the stage identifier *stid*.

$\mathbf{k}[stid] \in \{0,1\}^* \cup \{\bot\}$**:** An array of the session keys from each stage, or $\bot$ if no session key has yet been computed. Indexed by the stage identifier *stid*

$\mathbf{exk}[stid] \in \{0,1\}^* \cup \{\bot\}$**, where** $\mathbf{x} \in \{\mathbf{q}, \mathbf{c}, \mathbf{s}\}$**:** An array of the *post-quantum ephemeral asymmetric* (if $\mathbf{x} = \mathbf{q}$), *classic ephemeral asymmetric* (if $\mathbf{x} = \mathbf{c}$), or *ephemeral symmetric* (if $\mathbf{x} = \mathbf{s}$) secret keys used by the session in each stage. Initialised by $\bot$ and indexed by the stage identifier *stid*.

$\mathbf{pss}[stid] \in \{0,1\}^* \cup \{\bot\}$**:** Any per-stage secret state that is established during protocol execution for use in the following stage. Sessions use $\mathbf{pss}[stid - 1]$ during the protocol execution of stage *stid*. Indexed by *stid*.

$\mathbf{st}[stid] \in \{0,1\}^*$**:** Any additional state used by the session in each stage.

Finally, the challenger manages sets of corruption registers, which maintain the leakage status of various secrets (i.e. which secrets $\mathcal{A}$ has revealed). We describe the full set of registers in Appendix C.

4.3 Adversarial Interaction

Our HAKE framework considers a traditional AKE adversary, in complete control of the communication network, able to modify, inject, delete or delay messages. They are able to compromise several layers of secrets: (a) long-term private keys, modelling the misuse or corruption of long-term secrets in other sessions, and additionally allowing our model to capture forward-secrecy notions and quantum adversaries. (b) ephemeral private keys, modelling the leakage of secrets due to the use of bad randomness generators, or potentially bad cryptographic primitives or quantum adversaries. (c) preshared symmetric keys, modelling the leakage of shared secrets, potentially due to the misuse of the preshared secret by the partner, or the forced later revelation of these keys due to the compromise of partner devices. (d) ephemeral keying material, modelling attacks on the quantum key distribution. For instance, capturing things such as photon splitting attacks. (e) session keys, modelling the leakage of keys by their use in bad cryptographic algorithms. The adversary interacts with the challenger $\mathcal{C}$ via the queries below:

$\mathsf{Create}(i, j, role) \to \{(s), \bot\}$**:** Allows the adversary $\mathcal{A}$ to initialise a new session owned by party P_i, where the role of the new session is *role*, and intended communication partner party P_j. Note that if $\mathcal{A}$ has already initialised the intended partner session, $\mathcal{A}$ must give the session index r (indicating the intended partner session π_j^r) in order to synchronise ephemeral symmetric

keys. If a session π_i^s has already been created, $\mathcal{C}$ returns $\perp$. Otherwise, $\mathcal{C}$ returns (s) to $\mathcal{A}$.

$\mathsf{Send}(i, s, m) \rightarrow \{m', \perp\}$**:** Allows $\mathcal{A}$ to send messages to sessions for protocol execution and receive the output. If the session $\pi_i^s.\alpha \neq \texttt{active}$, then $\mathcal{C}$ returns $\perp$ to $\mathcal{A}$. Otherwise, $\mathcal{C}$ computes $\Pi.f(\lambda, \boldsymbol{pk}_i, \boldsymbol{sk}_i, \boldsymbol{pskid}_i, \boldsymbol{psk}_i, \pi_i^s, m) \rightarrow (m', \pi_i^{s\prime})$, sets $\pi_i^s \leftarrow \pi_i^{s\prime}$, updates transcripts $\pi_i^s.\mathbf{m}_r$, $\pi_i^s.\mathbf{m}_s$ and returns m' to $\mathcal{A}$.

$\mathsf{Reveal}(i, s, t)$**:** Allows $\mathcal{A}$ access to the session keys computed by a session. $\mathcal{C}$ checks if $\pi_i^s.\alpha[t] = \texttt{accept}$ and if so, returns $\pi_i^s.\mathbf{k}[t]$ to $\mathcal{A}$. In addition, the challenger checks if there exists another session π_j^r that *matches* with π_i^s, and also sets $\mathbf{SK}_j^r[t] \leftarrow \texttt{corrupt}$. Otherwise, $\mathcal{C}$ returns $\perp$ to $\mathcal{A}$.

$\mathsf{Test}(i, s, t) \rightarrow \{k_b, \perp\}$**:** Allows $\mathcal{A}$ access to a real-or-random session key k_b used in determining the success of $\mathcal{A}$ in the key-indistinguishability game. If a session π_i^s exists such that $\pi_i^s.\alpha = \texttt{accept}$, then the challenger $\mathcal{C}$ samples a key $k_0 \xleftarrow{\$} \mathcal{D}$ where $\mathcal{D}$ is the distribution of the session key, and sets $k_1 \leftarrow \pi_i^s.\mathbf{k}[t]$. $\mathcal{C}$ then returns k_b (where b is the random bit sampled during set-up) to $\mathcal{A}$. Otherwise $\mathcal{C}$ returns $\perp$ to $\mathcal{A}$.

$\mathsf{CorruptSK}(i, j) \rightarrow \{psk, \perp\}$**:** Allows $\mathcal{A}$ access to the secret preshared key $psk_i^j = psk_j^i$ jointly shared by parties P_i and P_j prior to protocol execution. If the preshared key has already been corrupted previously, then $\mathcal{C}$ returns $\perp$ to $\mathcal{A}$.

$\mathsf{CorruptQK}(i) \rightarrow \{qpk_i, \perp\}$**:** Allows $\mathcal{A}$ access to the secret post-quantum long-term key qsk_i generated for the party P_i prior to protocol execution. If the secret post-quantum long-term key has already been corrupted previously, then $\mathcal{C}$ returns $\perp$ to $\mathcal{A}$.

$\mathsf{CorruptCK}(i) \rightarrow \{cpk_i, \perp\}$**:** Allows $\mathcal{A}$ access to the secret classical long-term key csk_i generated for the party P_i prior to protocol execution. If the secret classical long-term key has already been corrupted previously, then $\mathcal{C}$ returns $\perp$ to $\mathcal{A}$.

$\mathsf{CompromiseQK}(i, s, t) \rightarrow \{\mathbf{eqk}[t], \perp\}$**:** Allows $\mathcal{A}$ access to the secret ephemeral post-quantum key $\pi_i^s.\mathbf{eqk}[t]$ generated for the session π_i^s prior to protocol execution in stage t. If $\pi_i^s.\mathbf{eqk}[t]$ has already been corrupted previously, then $\mathcal{C}$ returns $\perp$ to $\mathcal{A}$.

$\mathsf{CompromiseCK}(i, s, t) \rightarrow \{\mathbf{eck}[t], \perp\}$**:** Allows $\mathcal{A}$ access to the secret ephemeral classical key $\pi_i^s.\mathbf{eck}[t]$ generated for the session π_i^s prior to protocol execution in stage t. If $\pi_i^s.\mathbf{eck}[t]$ has already been corrupted previously, then $\mathcal{C}$ returns $\perp$ to $\mathcal{A}$.

$\mathsf{CompromiseSK}(i, s, t) \rightarrow \{\mathbf{esk}[t], \perp\}$**:** Allows $\mathcal{A}$ access to the secret ephemeral symmetric key $\pi_i^s.\mathbf{esk}[t]$ generated for the session π_i^s prior to protocol execution in stage t. Note that if there exists another session π_j^r such that $\pi_i^s.\mathbf{esk}[t] = \pi_j^r.\mathbf{esk}[t]$, then that session's ephemeral symmetric key is also considered corrupted. If $\pi_i^s.\mathbf{esk}[t]$ has already been corrupted previously, then $\mathcal{C}$ returns $\perp$ to $\mathcal{A}$.

$\mathsf{CompromiseSS}(i, s, t) \rightarrow \{\mathbf{pss}[t], \perp\}$**:** Allows the adversary access to the secret per-session state $\pi_i^s.\mathbf{pss}[t]$ generated by a session π_i^s during protocol execution. for use in the next stage of the session's protocol execution. Note that if there exists another session π_j^r such that $\pi_i^s.\mathbf{pss}[t] = \pi_j^r.\mathbf{pss}[t']$, then that

session's per-stage secret state is also considered corrupted. If $\pi_i^s.\mathbf{pss}[t]$ has already been corrupted previously, then $\mathcal{C}$ returns $\perp$ to $\mathcal{A}$.

4.4 Partnering Definition

To evaluate the secrets that $\mathcal{A}$ can reveal without trivially breaking the security of the protocol, key-exchange models must first define how sessions are *partnered*. Otherwise, $\mathcal{A}$ would simply run a protocol between two sessions, faithfully delivering all messages, Test the first session to receive the real-or-random key, and Reveal the other session's key. If the keys are equal, then the Test key is real, and otherwise the session key has been sampled randomly.

In our work, we use both the matching definition *matching sessions* defined in the original eCK model [22], and *origin sessions*, introduced by Cremers and Feltz [18]. On a high level, π_i^s is an origin session of π_j^r if π_i^s has received the messages that π_j^r sent without modification, even if the reply that π_i^s sent back has not been received by π_j^r. If all messages sent and received by π_i^s and π_j^r are identical, then the sessions *match*. We give a precise pseudocode description of these functions in Appendix C.

Definition 1 (Matching Sessions). *We consider $\pi_i^s.\mathbf{m}_s[t]$ and $\pi_i^s.\mathbf{m}_r[t]$ to be the concatentation of all messages sent and received (respectively) by a session π_i^s in a stage t. We say that π_i^s* matches *a session π_j^r in stage t if $\pi_i^s.pid = j$, $\pi_j^r.pid = i$, $\pi_i^s.\rho \neq \pi_j^r.\rho$, $\pi_i^s.\mathbf{m}_r[t] = \pi_j^r.\mathbf{m}_s[t]$ and $\pi_i^s.\mathbf{m}_s[t] = \pi_j^r.\mathbf{m}_r[t]$.*

We now turn to defining origin sessions for use in the HAKE security experiment.

Definition 2 (Origin Sessions). *We consider $\pi_i^s.\mathbf{m}_s[t]$ and $\pi_i^s.\mathbf{m}_r[t]$ to be the concatenation of all messages sent and received (respectively) by a session π_i^s in stage t. We say that π_i^s* matches *a session π_j^r in stage t if $\pi_j^r.\mathbf{m}_s[t] = \pi_i^s.\mathbf{m}_r[t]$. We say that π_i^s* prefix-matches *a session π_j^r in stage t if $\pi_j^r.\mathbf{m}_s[t] = \pi_i^s.\mathbf{m}_r[t]'$ where $\pi_i^s.\mathbf{m}_r[t]'$ is $\pi_i^s.\mathbf{m}_s[t]$ truncated to the length of $|\pi_j^r.\mathbf{m}_s[t]|$. Finally, we say that a session π_i^s has an* origin session *with π_j^r if π_i^s prefix-matches π_j^r (and π_i^s has sent the last message) or π_i^s matches π_j^r (and π_j^r has sent the last message).*

4.5 Cleanness Predicates

We now define the exact combinations of secrets that an adversary $\mathcal{A}$ is allowed to compromise without trivially breaking a hybrid key exchange protocol. However, we note that the cleanness predicate defined below is specific to Muckle, and the threat model that Muckle intends to defend against. Other predicates, both stronger and weaker, can be constructed.

We wish to capture security against a quantum-equipped adversary, so a successful adversary is allowed to compromise the long-term and ephemeral classical asymmetric secrets without penalty. Since Muckle itself does not use public-key cryptography to authenticate its messages, we allow $\mathcal{A}$ to compromise the long-term asymmetric secrets (however, the challenger $\mathcal{C}$ will respond to CorruptCK and CorruptQK queries with $\perp$).

Since we wish to capture perfect forward secrecy, we allow a successful adversary to have issued a Test query to a session π_i^s owned by a party P_i (with no origin session) that has had its long-term symmetric key compromised previously, as long as the session was completed before the $\mathsf{CorruptSK}(i, j)$ query was issued (and $\pi_i^s.pid = j$). In addition, our construction should be post-compromise secure (as explored by Cohn-Gordon et al. [16]), so our cleanness predicate allows an adversary to have compromised *all* ephemeral secrets associated with a particular stage as long as there exists some stage previous that has not had all its ephemeral secrets compromised and the adversary has been passive in all stages between the "Test" stage and the previous "clean" stage.

Coming full circle then, a "clean" stage intuitively is one where the adversary has not compromised *all* of: (a) the ephemeral classic secrets of the Test session and its matching partner in the tested stage (b) the ephemeral post-quantum secrets of the Test session and its matching partner in the tested stage (c) the previous per-stage secrets shared by the Test session and its matching session in the tested stage, and (d) the quantum keying material / ephemeral symmetric secrets shared by the Test session and its matching session in the tested stage.

We formalise this intuition below in Definition 3.

Definition 3 ($\mathsf{clean}_{q\mathsf{HAKE}}$)**.** *A session π_i^s in stage t such that $\pi_i^s.\alpha[t] = \texttt{accept}$ and $\pi_i^s.pid = j$ in the security experiment defined in Figure 5 is $\mathsf{clean}_{q\mathsf{HAKE}}$ if all of the following conditions hold:*

1. *The query $\mathsf{Reveal}(i, s, t)$ has not been issued.*
2. *For all $(j, r, t) \in n_P \times n_S \times n_T$ such that π_i^s matches π_j^r in stage t, the query $\mathsf{Reveal}(j, r, t)$ has not been issued.*
3. *If there exists a session π_j^r such that π_j^r matches π_i^s in stage t, then at least one of the following sets of queries has not been issued:*
   - $\mathsf{CompromiseQK}(i, s, t)$, $\mathsf{CompromiseQK}(j, r, t)$ *have not been issued, where π_j^r* matches *π_i^s in stage t.*
   - $\mathsf{CompromiseSK}(i, s, t)$, $\mathsf{CompromiseSK}(j, r, t)$ *have not been issued, where π_j^r* matches *π_i^s in stage t.*
   - $\mathsf{CompromiseQK}(i, s, t')$, $\mathsf{CompromiseQK}(j, r, t')$ *have not been issued, where π_j^r* matches *π_i^s in stages u such that $t' \leq u < t$ and no $\mathsf{CompromiseSS}(i, s, u)$, $\mathsf{CompromiseSS}(j, r, u)$ queries have been issued.*[9]
   - $\mathsf{CompromiseSK}(i, s, t')$, $\mathsf{CompromiseSK}(j, r, t')$ *have not been issued, where π_j^r* matches *π_i^s in stages u such that $t' \leq u < t$ and no $\mathsf{CompromiseSS}(i, s, u)$, $\mathsf{CompromiseSS}(j, r, u)$ queries have been issued.* [10]

[9] This condition was added to capture post-compromise security (PCS). Typically, in key-exchange frameworks that capture PCS the cleanness predicate is recursive, i.e. there is a condition that says that *this* stage is clean if there exists a *previous* stage that is *also* clean, and the adversary is restricted from issuing queries that would prevent this "cleanness" from propagating into future stages. We make explicit in our model *which* queries that the adversary cannot have issued in the previous and intermediate stages for the purpose of clarity.

[10] Refer to footnote 9.

4. *If there exists no* $(j, r, t) \in n_P \times n_S \times n_T$ *such that* π_j^r *is an origin session of* π_i^s *in stage* t*, then* $\mathsf{CorruptSK}(i, j)$ *and* $\mathsf{CorruptSK}(j, i)$ *have not been issued before* $\pi_i^s.\alpha[t] \leftarrow \texttt{accept}$*. If there exists a* $(j, r, t) \in n_P \times n_S \times n_T$ *such that* π_j^r *is an origin session of* π_i^s *in stage* t*, then* $\mathsf{CorruptSK}(i, j)$ *and* $\mathsf{CorruptSK}(j, i)$ *have not been issued before* $\pi_j^r.\alpha[t] \leftarrow \texttt{accept}$*.*

It may also be desirable to determine the security guarantees that Muckle provides in the event of a new vulnerability discovered in the underlying post-quantum asymmetric key-exchange primitive, or a side-channel attack being discovered in the hardware of the QKD system. In order to capture this scenario, we provide a second cleanness predicate that captures non-quantum-equipped adversaries, which we denote $\mathsf{clean}_{c\mathsf{HAKE}}$. It is more-or-less identical to $\mathsf{clean}_{q\mathsf{HAKE}}$, with the following additional restricted sets of queries in condition three:

- $\mathsf{CompromiseCK}(i, s, t)$, $\mathsf{CompromiseCK}(j, r, t)$ have not been issued, where π_j^r *matches* π_i^s in stage t.
- $\mathsf{CompromiseCK}(i, s, t')$, $\mathsf{CompromiseCK}(j, r, t')$ have not been issued, where π_j^r *matches* π_i^s in stages u such that $t' \leq u < t$ and no $\mathsf{CompromiseSS}(i, s, u)$, $\mathsf{CompromiseSS}(j, r, u)$ queries have been issued.

We give the definition for this predicate in Appendix D. Next, we formalise the advantage of a QPT algorithm $\mathcal{A}$ in winning the HAKE key indistinguishability experiment in the following way:

Definition 4 (HAKE Key Indistinguishability). *Let* Π *be a key-exchange protocol, and* n_P, n_S, $n_T \in \mathbb{N}$*. For a particular given predicate* clean*, and a QPT algorithm* $\mathcal{A}$*, we define the advantage of* $\mathcal{A}$ *in the* HAKE *key-indistinguishability game to be :*

$$\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean},\mathcal{A}}_{\Pi,n_P,n_S,n_T}(\lambda) = 2 \cdot \left| \Pr\left[\mathsf{Exp}^{\mathsf{HAKE},\mathsf{clean},\mathcal{A}}_{\Pi,n_P,n_S n_T}(\lambda) = 1\right] - \frac{1}{2} \right|.$$

We say that Π *is post-quantum* HAKE*-secure if, for all* $\mathcal{A}$*,* $\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean},\mathcal{A}}_{\Pi,n_P,n_S,n_T}(\lambda)$ *is negligible in the security parameter* λ*.*

5 Security Analysis

This section is dedicated to proving our main result Theorem 1. As discussed in Section 4.5, it is also desirable to assess security of Muckle with respect to classic probabilistic polynomial-time adversaries. We prove that Muckle is HAKE-secure with cleanness predicate $\mathsf{clean}_{c\mathsf{HAKE}}$, see Theorem 2 in Section 5.1. Recall that $\mathsf{clean}_{c\mathsf{HAKE}}$ used here is a generalisation of $\mathsf{clean}_{q\mathsf{HAKE}}$, allowing us to establish key indistinguishability security in the scenario where the classical cryptographic component of Muckle remains secure and uncompromised, even if the security of the post-quantum and QKD components both fail.

Theorem 1. *The* Muckle *key exchange protocol is* HAKE*-secure with cleanness predicate* $\mathsf{clean}_{q\mathsf{HAKE}}$ *(capturing perfect forward secrecy and post-compromise security) under the* prf, eufqcma, dual-prf, *and* ind-cpa *assumptions of* PRF, MAC, PRF *and* KEM*, respectively. That is, for any QPT algorithm* $\mathcal{A}$ *against the* HAKE *key-indistinguishability game (defined in Figure 5)* $\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{q\mathsf{HAKE}},\mathcal{A}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda)$ *is negligible , with:*

$$\begin{aligned}
&\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{q\mathsf{HAKE}},\mathcal{A}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) \leq \\
&\quad 2 \cdot n_P^2 n_S n_T \cdot \big(\mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \mathsf{Adv}^{\mathsf{eufcma}}_{\mathsf{MAC},\mathcal{A}}(\lambda)\big) \\
&\quad + n_P^2 n_S^2 n_T \cdot \big(\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{KEM},\mathcal{A}}(\lambda) + (10 + 2 \cdot n_T) \cdot \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) \\
&\quad + (11 + 2 \cdot n_T) \cdot \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda)\big).
\end{aligned}$$

Proof. We begin by dividing the proof into three separate cases (and denote with $\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{\mathsf{HAKE}},\mathcal{A},C_l}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda)$ the advantage of the adversary in winning the key-indistinguishability game in Case l) where the query $\mathsf{Test}(i,s,t)$ has been issued:

1. The session π_i^s (where $\pi_i^s.\rho = \texttt{init}$) has no origin session in stage t.
2. The session π_i^s (where $\pi_i^s.\rho = \texttt{resp}$) has no origin session in stage t.
3. The session π_i^s in stage t has a matching session.

It follows then that

$$\begin{aligned}
\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{q\mathsf{HAKE}},\mathcal{A}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) \leq\ & \mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{q\mathsf{HAKE}},\mathcal{A},C_1}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda), \\
&+ \mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{q\mathsf{HAKE}},\mathcal{A},C_2}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda), \\
&+ \mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{q\mathsf{HAKE}},\mathcal{A},C_3}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda).
\end{aligned}$$

We then bound the probability of each case, and show that under certain assumptions, the probability of the adversary winning in the key-indistinguishability game is negligible.

Case 1: Test `init` session without origin session

We begin by showing that $\mathcal{A}$ has negligible change in causing π_i^s to reach an `accept` state without an origin session.

We do so via the sequence of game hops:

Game 0 This is the standard HAKE security game. Thus:

$$\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{q\mathsf{HAKE}},\mathcal{A},C_1}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) = \Pr(break_0).$$

Game 1 In this game, we guess the index (i,s,t) and the intended partner j and abort if, during the execution of the experiment, a query $\mathsf{Test}(i',s',t')$ is received to a session $\pi_{i'}^{s'}$ such that $\pi_{i'}^{s'}.pid = j'$ and $(i,s,t,j) \neq (i',s',t','j)$. Thus:

$$\Pr(break_0) \leq n_P^2 n_S n_T \cdot \Pr(break_1).$$

Game 2 In this game we abort if the test session π_i^s sets the status $\pi_i^s.\alpha[t] \leftarrow$ `reject`. Note that by the previous game we abort if the Test query is issued to a session that is not π_i^s in stage t. If the session π_i^s ever reaches the status $\pi_i^s.\alpha[t] \leftarrow$ `reject`, then the challenger will respond to the $\mathsf{Test}(i, s, t)$ query with $\perp$, and thus the difference in $\mathcal{A}$'s advantage between **Game 2** and **Game 3** is 0. Thus:

$$\Pr(break_1) \leq \Pr(break_2).$$

Game 3 In this game we define an abort event $\mathbf{abort}_\alpha$ that triggers if the test session π_i^s sets the status $\pi_i^s.\alpha[t] \leftarrow$ `accept`. We note that the response to the Test query issued by $\mathcal{A}$ is always $\perp$, and thus $\Pr(break_3) = 0$. In what follows we bound the probability of $\mathcal{A}$ in causing $\mathbf{abort}_\alpha$ to trigger. Thus:

$$\Pr(break_2) \leq \Pr(\mathbf{abort}_\alpha).$$

Game 4 In this game we replace the computation of $mkey_B = \mathsf{PRF}(PSK, SecState)$ a with uniformly random and independent value $\widetilde{mkey_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t. We do so by initialising a post-quantum prf challenger and query $SecState$, and use the output $\widetilde{mkey_B}$ from the prf challenger to replace the computation of $mkey_B$. Since $PSK = psk_i^j$ is itself uniformly random and independent, and $\mathcal{A}$ cannot issue $\mathsf{CorruptSK}(i, j)$ or $\mathsf{CorruptSK}(j, i)$, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{mkey_B} = \mathsf{PRF}(PSK, SecState)$ and we are in **Game 3**. If the test bit sampled by the prf challenger is 1, then $\widetilde{mkey_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 4**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the post-quantum prf security of PRF, and we find:

$$\Pr(\mathbf{abort}_\alpha) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_4).$$

Game 5 In this game, we abort when the session π_i^s accepts without an origin session, i.e. π_i^s receives a MAC tag in stage t that verifies correctly, but there exists no honest session π_j^r that output m_1, τ_1. We do this by interacting with the post-quantum eufcma MAC challenger, but computes $\tau_1 = \mathsf{MAC}(\widetilde{mkey_B}, m_1)$ for π_i^s by querying m_1 received by π_i^s to the MAC challenger. Since the key $\widetilde{mkey_B}$ was already uniformly random and independent of the experiment by **Game 4**, and by the definition of this case, $\mathcal{A}$ has not issued either $\mathsf{CorruptSK}(i, j)$ or $\mathsf{CorruptSK}(j, i)$ queries, this change is indistinguishable. By the definition of Case 1, if π_i^s has no origin session then adversary must have produced a valid MAC tag τ_1 such that $\mathsf{MAC}(\widetilde{mkey_B}, m_1) = \tau_1$. We submit τ_1, m_1 as a forgery to the MAC challenger and aborts. Since π_i^s now aborts when verifying the MAC tag in stage t, it cannot trigger $\mathbf{abort}_\alpha$ and thus we have:

$$\begin{aligned}\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{q_{\mathsf{HAKE}}},\mathcal{A},C_1}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) \leq n_P^2 n_S n_T \cdot \big(&\mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) \\ &+ \mathsf{Adv}^{\mathsf{eufcma}}_{\mathsf{MAC},\mathcal{A}}(\lambda)\big).\end{aligned}$$

The proof of Case 2 follows analogously to Case 1, and with the same bounds. Thus:

$$\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{q\mathsf{HAKE}},\mathcal{A},C_2}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) \leq n_P^2 n_S n_T \cdot \big(\mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \mathsf{Adv}^{\mathsf{eufcma}}_{\mathsf{MAC},\mathcal{A}}(\lambda)\big).$$

Case 3: Test session with matching session

In Case 3, we show that if $\mathcal{A}$ that has issued a $\mathsf{Test}(i,s,t)$ query to a clean session π_i^s in stage t, then $\mathcal{A}$ has negligible advantage in guessing the test bit b. In what follows, we split our analysis of Case 3 into the following sub-cases, each corresponding to a condition necessary for the cleanness predicate clean to be upheld by π_i^s in stage t. These are the subcases:

3.1 $\mathsf{CompromiseQK}(i,s,t)$, $\mathsf{CompromiseQK}(j,r,t)$ have not been issued, where π_j^r *matches* π_i^s in stage t.
3.2 $\mathsf{CompromiseSK}(i,s,t)$, $\mathsf{CompromiseSK}(j,r,t)$ have not been issued, where π_j^r *matches* π_i^s in stage t.
3.3 $\mathsf{CompromiseQK}(i,s,t')$, $\mathsf{CompromiseQK}(j,r,t')$ have not been issued, where π_j^r *matches* π_i^s in stages u such that $t' \leq u < t$ and no $\mathsf{CompromiseSS}(i,s,u)$, $\mathsf{CompromiseSS}(j,r,u)$ queries have been issued.
3.4 $\mathsf{CompromiseSK}(i,s,t')$, $\mathsf{CompromiseSK}(j,r,t')$ have not been issued, where π_j^r *matches* π_i^s in stages u such that $t' \leq u < t$ and no $\mathsf{CompromiseSS}(i,s,u)$, $\mathsf{CompromiseSS}(j,r,u)$ queries have been issued.

It is straightforward to see that the advantage of $\mathcal{A}$ in Case 3 is bound by the sum of the advantage of $\mathcal{A}$ in all subcases. Due to space restrictions, we only detail the proof sketches of subcases 3.1 and 3.3. The other two subcases follow analogously with similar proof strategies. The full proof details can be found in Section 5. We begin by treating the first subcase.

3.1: $\mathsf{CompromiseQK}(i,s,t)$, $\mathsf{CompromiseQK}(j,r,t)$ have not been issued, where π_j^r *matches* π_i^s in stage t.

Game 0 This is the standard HAKE security game. Thus:

$$\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{q\mathsf{HAKE}},\mathcal{A},C_{3.1}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) = \Pr(break_0).$$

Game 1 In this game, we guess the index (i,s,t) and the matching session (j,r,t) and abort if, during the execution of the experiment, a query $\mathsf{Test}(i',s',t')$ is received to a session $\pi_{i'}^{s'}$ such that $\pi_{j'}^{r'}$ matches $\pi_{i'}^{s'}$ in stage t' and $(i,s,t),(j,r) \neq (i',s',t'),(j',r')$. Thus:

$$\Pr(break_0) \leq n_P^2 n_S^2 n_T \cdot \Pr(break_1).$$

Game 2 In this game, we replace the key qsk derived in the test session π_i^s with the uniformly random and independent value $\widetilde{qsk}$. We do so by interacting with a $\mathsf{ind\text{-}cpa}$ KEM challenger (as described in Definition 5) and replace the qpk_A value sent in m_0, and the ciphertext qpk_B sent in m_1 with the public-key pk and the ciphertext c received from the $\mathsf{ind\text{-}cpa}$ KEM challenger. Since π_i^s matches π_j^r in stage t, we know that the public-key and ciphertext sent in m_0 and m_1 respectively were received by the sessions without modification. Detecting the replacement of qsk with a uniformly random value $\widetilde{qsk}$ implies an efficient distinguishing algorithm against the post-quantum $\mathsf{ind\text{-}cpa}$ security of KEM. Thus:

$$\Pr(break_1) \leq \mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{KEM},\mathcal{A}}(\lambda) + \Pr(break_2).$$

Game 3 In this game we replace the computation of the quantum key $qk = \mathsf{PRF}(\widetilde{qsk}, \mathsf{label}_q k)$ with a uniformly random and independent value $\widetilde{qk} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t, and its matching session π_j^r. We do so by initialising a post-quantum prf challenger and query label_{qk}, and use the output $\widetilde{qk}$ from the prf challenger to replace the computation of qk. Since $\widetilde{qsk}$ is uniformly random and independent by **Game 2**, and $\mathcal{A}$ cannot issue $\mathsf{CompromiseQK}(i,s,t)$ or $\mathsf{CompromiseQK}(j,r,t)$, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{qk} = \mathsf{PRF}(\widetilde{qsk}, \mathsf{label}_q k)$ and we are in **Game 2**. If the test bit sampled by the prf challenger is 1, then $\widetilde{qk} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 3**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the post-quantum prf security of PRF, and we find:

$$\Pr(break_2) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_3).$$

Game 4 In this game we replace the computation of the first chaining key $k_0 \leftarrow \mathsf{PRF}(\widetilde{qk}, m_0 \| m_1)$ with a uniformly random and independent value $\widetilde{k_0} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t, and its matching session π_j^r. We do so by initialising a post-quantum prf challenger and query $m_0 \| m_1$, and use the output $\widetilde{k_0}$ from the prf challenger to replace the computation of k_0. Since $\widetilde{qk}$ is uniformly random and independent by **Game 3**, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{k_0} = \mathsf{PRF}(\widetilde{qk}, m_0 \| m_1)$ and we are in **Game 3**. If the test bit sampled by the prf challenger is 1, then $\widetilde{k_0} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 4**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the prf security of post-quantum PRF, and we find:

$$\Pr(break_3) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_4).$$

Game 5 In this game we replace the computation of the second chaining key $k_1 \leftarrow \mathsf{PRF}(ck, \widetilde{k_0})$ with a uniformly random and independent value

$\widetilde{k_0} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t, and its matching session π_j^r. We do so by initialising a post-quantum $\mathsf{dual\text{-}prf}$ challenger and query ck, and use the output $\widetilde{k_1}$ from the $\mathsf{dual\text{-}prf}$ challenger to replace the computation of k_1. Since $\widetilde{k_0}$ is uniformly random and independent by **Game 4**, this is a sound replacement. If the test bit sampled by the $\mathsf{dual\text{-}prf}$ challenger is 0, then $\widetilde{k_1} = \mathsf{PRF}(ck, \widetilde{k_0})$ and we are in **Game 4**. If the test bit sampled by the $\mathsf{dual\text{-}prf}$ challenger is 1, then $\widetilde{k_1} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 5**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the post-quantum $\mathsf{dual\text{-}prf}$ security of PRF, and we find:

$$\Pr(break_4) \leq \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_5).$$

Game 6 In this game we replace the computation of the third chaining key $k_2 \leftarrow \mathsf{PRF}(qkm, \widetilde{k_1})$ with a uniformly random and independent value $\widetilde{k_2} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t, and its matching session π_j^r. We do so by initialising a post-quantum $\mathsf{dual\text{-}prf}$ challenger and query qkm, and use the output $\widetilde{k_2}$ from the $\mathsf{dual\text{-}prf}$ challenger to replace the computation of k_2. Since $\widetilde{k_1}$ is uniformly random and independent by **Game 5**, this is a sound replacement. If the test bit sampled by the $\mathsf{dual\text{-}prf}$ challenger is 0, then $\widetilde{k_2} = \mathsf{PRF}(qkm, \widetilde{k_1})$ and we are in **Game 5**. If the test bit sampled by the $\mathsf{dual\text{-}prf}$ challenger is 1, then $\widetilde{k_2} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 6**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the post-quantum $\mathsf{dual\text{-}prf}$ security of PRF, and we find:

$$\Pr(break_5) \leq \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_6).$$

Game 7 In this game we replace the computation of the final chaining key $k_3 \leftarrow \mathsf{PRF}(SecState, \widetilde{k_2})$ with a uniformly random and independent value $\widetilde{k_3} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t, and its matching session π_j^r. We do so by initialising a post-quantum $\mathsf{dual\text{-}prf}$ challenger and query $SecState$, and use the output $\widetilde{k_3}$ from the $\mathsf{dual\text{-}prf}$ challenger to replace the computation of k_3. Since $\widetilde{k_2}$ is uniformly random and independent by **Game 6**, this is a sound replacement. If the test bit sampled by the $\mathsf{dual\text{-}prf}$ challenger is 0, then $\widetilde{k_3} = \mathsf{PRF}(SecState, \widetilde{k_2})$ and we are in **Game 6**. If the test bit sampled by the $\mathsf{dual\text{-}prf}$ challenger is 1, then $\widetilde{k_3} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 7**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the post-quantum $\mathsf{dual\text{-}prf}$ security of PRF, and we find:

$$\Pr(break_6) \leq \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_7).$$

Game 8 In this game we replace the computation of the updated secret state and session keys $SecState', sk_A, sk_B \leftarrow \mathsf{PRF}(\widetilde{k_3}, m_0 \| m_1 \| ctr)$ with uniformly

random and independent values $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t, and its matching session π_j^r. We do so by initialising a post-quantum prf challenger and query $m_0 \| m_1 \| ctr$, and use the output $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B}$ from the prf challenger to replace the computation of $SecState', sk_A, sk_B$. Since $\widetilde{k_3}$ is uniformly random and independent by **Game 7**, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} = \mathsf{PRF}(\widetilde{k_3}, m_0 \| m_1 \| ctr)$ and we are in **Game 7**. If the test bit sampled by the prf challenger is 1, then $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 8**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the post-quantum prf security of PRF, and we find:

$$\Pr(break_7) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_8).$$

Since $\widetilde{sk_A}, \widetilde{sk_B}$ are now uniformly random and independent values independent of the protocol flow regardless of the value of the test bit b, $\mathcal{A}$ has no advantage in guessing the bit and thus

$$\begin{aligned}\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{q\mathsf{HAKE}},\mathcal{A},C_{3.1}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) \leq n_P^2 n_S^2 n_T \cdot \big(&\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{KEM},\mathcal{A}}(\lambda) \\ &+ 3 \cdot \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) \\ &+ 3 \cdot \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda)\big).\end{aligned}$$

3.2: $\mathsf{CompromiseSK}(i,s,t)$, $\mathsf{CompromiseSK}(j,r,t)$ have not been issued, where π_j^r *matches* π_i^s in stage t.

Game 0 This is the standard HAKE security game. Thus:

$$\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{q\mathsf{HAKE}},\mathcal{A},C_{3.2}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) = \Pr(break_0).$$

Game 1 In this game, we guess the index (i,s,t) and the matching session (j,r,t) and abort if, during the execution of the experiment, a query $\mathsf{Test}(i',s',t')$ is received to a session $\pi_{i'}^{s'}$ such that $\pi_{j'}^{r'}$ matches $\pi_{i'}^{s'}$ in stage t' and $(i,s,t),(j,r) \neq (i',s',t'),(j',r')$. Thus:

$$\Pr(break_0) \leq n_P^2 n_S^2 n_T \cdot \Pr(break_1).$$

Game 2 In this game we replace the computation of the third chaining key $k_2 \leftarrow \mathsf{PRF}(qkm, k_1)$ with a uniformly random and independent value $\widetilde{k_2} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t, and its matching session π_j^r. We do so by initialising a post-quantum prf challenger and query k_1, and use the output $\widetilde{k_2}$ from the prf challenger to replace the computation of k_2. Since qkm is uniformly random and independently generated by $\mathsf{ESKeyGen}$

at the start of the experiment, and by the definition of this subcase $\mathcal{A}$ cannot have issued either $\mathsf{CompromiseSK}(i, s, t)$ or $\mathsf{CompromiseSK}(j, r, t)$ this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{k_2} = \mathsf{PRF}(qkm, k_1)$ and we are in **Game 1**. If the test bit sampled by the prf challenger is 1, then $\widetilde{k_2} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 2**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the post-quantum prf security of PRF, and we find:

$$\Pr(break_1) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_2).$$

Game 3 In this game we replace the computation of the final chaining key $k_3 \leftarrow \mathsf{PRF}(SecState, \widetilde{k_2})$ with a uniformly random and independent value $\widetilde{k_3} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t, and its matching session π_j^r. We do so by initialising a post-quantum dual-prf challenger and query $SecState$, and use the output $\widetilde{k_3}$ from the dual-prf challenger to replace the computation of k_3. Since $\widetilde{k_2}$ is uniformly random and independent by **Game 2**, this is a sound replacement. If the test bit sampled by the dual-prf challenger is 0, then $\widetilde{k_3} = \mathsf{PRF}(SecState, \widetilde{k_2})$ and we are in **Game 2**. If the test bit sampled by the dual-prf challenger is 1, then $\widetilde{k_3} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 3**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the post-quantum dual-prf security of PRF, and we find:

$$\Pr(break_2) \leq \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_3).$$

Game 4 In this game we replace the computation of the updated secret state and session keys $SecState', sk_A, sk_B \leftarrow \mathsf{PRF}(\widetilde{k_3}, m_0 \| m_1 \| ctr)$ with uniformly random and independent values $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t, and its matching session π_j^r. We do so by initialising a post-quantum prf challenger and query $m_0 \| m_1 \| ctr$, and use the output $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B}$ from the prf challenger to replace the computation of $SecState', sk_A, sk_B$. Since $\widetilde{k_3}$ is uniformly random and independent by **Game 3**, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} = \mathsf{PRF}(\widetilde{k_3}, m_0 \| m_1 \| ctr)$ and we are in **Game 3**. If the test bit sampled by the prf challenger is 1, then $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 4**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the post-quantum prf security of PRF, and we find:

$$\Pr(break_3) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_4).$$

Since $\widetilde{sk_A}, \widetilde{sk_B}$ are now uniformly random and independent values independent of the protocol flow regardless of the value of the test bit b, $\mathcal{A}$ has no

advantage in guessing the bit and thus:

$$\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{q\mathsf{HAKE}},\mathcal{A},C_{3.2}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) \leq n_P^2 n_S^2 n_T \cdot \big(2 \cdot \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda)\big).$$

3.3: $\mathsf{CompromiseQK}(i, s, t^*)$, $\mathsf{CompromiseQK}(j, r, t^*)$ have not been issued, where π_j^r *matches* π_i^s in stages u such that $t^* \leq u < t$ and no $\mathsf{CompromiseSS}(i, s, u)$, $\mathsf{CompromiseSS}(j, r, u)$ queries have been issued.

Game 0 This is the standard HAKE security game. Thus:

$$\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{q\mathsf{HAKE}},\mathcal{A},C_{3.3}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) = \Pr(break_0).$$

Game 1 In this game, we guess the index (i, s, t) and the matching session (j, r, t) and abort if, during the execution of the experiment, a query $\mathsf{Test}(i', s', t')$ is received to a session $\pi_{i'}^{s'}$ such that $\pi_{j'}^{r'}$ matches $\pi_{i'}^{s'}$ in stage t' and $(i, s, t), (j, r) \neq (i', s', t'), (j', r')$. In addition, we also guess the stage t^* such that $\mathcal{A}$ does not issue either a $\mathsf{CompromiseQK}(i, s, t^*)$, $\mathsf{CompromiseQK}(j, r, t^*)$ and π_i^s matches in stages u such that $t^* \leq u < t$ and no $\mathsf{CompromiseSS}(i, s, u)$ or $\mathsf{CompromiseSS}(j, r, u)$ were made. By the definition of this subcase, such a stage *must* exist. Thus:

$$\Pr(break_0) \leq n_P^2 n_S^2 n_T^2 \cdot \Pr(break_1).$$

Game 2 In this game, we replace the key qsk derived in the test session π_i^s with the uniformly random and independent value $\widetilde{qsk}$. We do so by interacting with a ind-cpa KEM challenger (as described in Definition 5) and replace the qpk_A value sent in m_0, and the ciphertext qpk_B sent in m_1 with the public-key pk and the ciphertext c received from the ind-cpa KEM challenger. Since π_i^s matches π_j^r in stage t, we know that the public-key and ciphertext sent in m_0 and m_1 respectively were received by the sessions without modification. Detecting the replacement of qsk with a uniformly random value $\widetilde{qsk}$ implies an efficient distinguishing algorithm against the post-quantum ind-cpa security of KEM. Thus:

$$\Pr(break_1) \leq \mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{KEM},\mathcal{A}}(\lambda) + \Pr(break_2).$$

Game 3 In this game we replace the computation of the quantum key $qk = \mathsf{PRF}(\widetilde{qsk}, \mathsf{label}_q k)$ with a uniformly random and independent value $\widetilde{qk} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t^*, and its matching session π_j^r. We do so by initialising a post-quantum prf challenger and query label_{qk}, and use the output $\widetilde{qk}$ from the prf challenger to replace the computation of qk. Since $\widetilde{qsk}$ is uniformly random and independent by **Game 2**, and $\mathcal{A}$ cannot issue $\mathsf{CompromiseQK}(i, s, t^*)$ or $\mathsf{CompromiseQK}(j, r, t^*)$, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{qk} = \mathsf{PRF}(\widetilde{qsk}, \mathsf{label}_q k)$ and we are in **Game 2**. If the test bit sampled by

the prf challenger is 1, then $\widetilde{qk} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 3**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the post-quantum prf security of PRF, and we find:

$$\Pr(break_2) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_3).$$

Game 4 In this game we replace the computation of the first chaining key $k_0 \leftarrow \mathsf{PRF}(\widetilde{qk}, m_0 \| m_1)$ with a uniformly random and independent value $\widetilde{k_0} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t^*, and its matching session π_j^r. We do so by initialising a post-quantum prf challenger and query $m_0 \| m_1$, and use the output $\widetilde{k_0}$ from the prf challenger to replace the computation of k_0. Since $\widetilde{qk}$ is uniformly random and independent by **Game 3**, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{k_0} = \mathsf{PRF}(\widetilde{qk}, m_0 \| m_1)$ and we are in **Game 3**. If the test bit sampled by the prf challenger is 1, then $\widetilde{k_0} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 4**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the post-quantum prf security of PRF, and we find:

$$\Pr(break_3) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_4).$$

Game 5 In this game we replace the computation of the second chaining key $k_1 \leftarrow \mathsf{PRF}(ck, \widetilde{k_0})$ with a uniformly random and independent value $\widetilde{k_0} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t^*, and its matching session π_j^r. We do so by initialising a post-quantum dual-prf challenger and query ck, and use the output $\widetilde{k_1}$ from the dual-prf challenger to replace the computation of k_1. Since $\widetilde{k_0}$ is uniformly random and independent by **Game 4**, this is a sound replacement. If the test bit sampled by the dual-prf challenger is 0, then $\widetilde{k_1} = \mathsf{PRF}(ck, \widetilde{k_0})$ and we are in **Game 4**. If the test bit sampled by the dual-prf challenger is 1, then $\widetilde{k_1} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 5**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the dual-prf security of post-quantum PRF, and we find:

$$\Pr(break_4) \leq \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_5).$$

Game 6 In this game we replace the computation of the third chaining key $k_2 \leftarrow \mathsf{PRF}(qkm, \widetilde{k_1})$ with a uniformly random and independent value $\widetilde{k_2} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t^*, and its matching session π_j^r. We do so by initialising a post-quantum dual-prf challenger and query qkm, and use the output $\widetilde{k_2}$ from the dual-prf challenger to replace the computation of k_2. Since $\widetilde{k_1}$ is uniformly random and independent by **Game 5**, this is a sound replacement. If the test bit sampled by the dual-prf challenger is

0, then $\widetilde{k_2} = \mathsf{PRF}(qkm, \widetilde{k_1})$ and we are in **Game 5**. If the test bit sampled by the dual-prf challenger is 1, then $\widetilde{k_2} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 6**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the dual-prf security of post-quantum PRF, and we find:

$$\Pr(break_5) \leq \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_6).$$

Game 7 In this game we replace the computation of the final chaining key $k_3 \leftarrow \mathsf{PRF}(SecState, \widetilde{k_2})$ with a uniformly random and independent value $\widetilde{k_3} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t^*, and its matching session π_j^r. We do so by initialising a post-quantum dual-prf challenger and query $SecState$, and use the output $\widetilde{k_3}$ from the dual-prf challenger to replace the computation of k_3. Since $\widetilde{k_2}$ is uniformly random and independent by **Game 6**, this is a sound replacement. If the test bit sampled by the dual-prf challenger is 0, then $\widetilde{k_3} = \mathsf{PRF}(SecState, \widetilde{k_2})$ and we are in **Game 6**. If the test bit sampled by the dual-prf challenger is 1, then $\widetilde{k_3} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 7**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the post-quantum dual-prf security of PRF, and we find:

$$\Pr(break_6) \leq \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_7).$$

Game 8 In this game we replace the computation of the updated secret state and session keys $SecState', sk_A, sk_B \leftarrow \mathsf{PRF}(\widetilde{k_3}, m_0 \| m_1 \| ctr)$ with uniformly random and independent values $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t^*, and its matching session π_j^r. We do so by initialising a post-quantum prf challenger and query $m_0 \| m_1 \| ctr$, and use the output $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B}$ from the prf challenger to replace the computation of $SecState', sk_A, sk_B$. Since $\widetilde{k_3}$ is uniformly random and independent by **Game 7**, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} = \mathsf{PRF}(\widetilde{k_3}, m_0 \| m_1 \| ctr)$ and we are in **Game 7**. If the test bit sampled by the prf challenger is 1, then $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 8**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the post-quantum prf security of PRF, and we find:

$$\Pr(break_7) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_8).$$

At this point we now iteratively do the following two game hops $(t - t^*)$ times in each consecutive stage u where $t^* < u \leq t$. Note that, by the definition of this sub-case, each of the stages following stage t^* must match (i.e. there must exist a matching session π_j^r with π_i^s that agree upon the message transcript):

Game 9' In this game we replace the computation of the final chaining key $k_3 \leftarrow \mathsf{PRF}(\widetilde{SecState}, k_2)$ with a uniformly random and independent value $\widetilde{k_3} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage u, and its matching session π_j^r. We do so by initialising a post-quantum prf challenger and query k_2, and use the output $\widetilde{k_3}$ from the prf challenger to replace the computation of k_3 in stage u. Since $\widetilde{SecState}$ is uniformly random and independent by the previous game, and $\mathcal{A}$ cannot issue a $\mathsf{CompromiseSS}(i,s,u)$ or $\mathsf{CompromiseSS}(j,r,u)$ $\forall t^* \leq u < t$, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{k_3} = \mathsf{PRF}(\widetilde{SecState}, k_2)$ and we are in the previous game. If the test bit sampled by the prf challenger is 1, then $\widetilde{k_3} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 9'**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the post-quantum prf security of PRF, and we find:

$$\Pr(break_8) \leq \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_{9'}),$$
$$\Pr(break_{10'}) \leq \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_{9'})$$

Game 10' In this game we replace the computation of the updated secret state and session keys $SecState', sk_A, sk_B \leftarrow \mathsf{PRF}(\widetilde{k_3}, m_0 \| m_1 \| ctr)$ with uniformly random and independent values $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage u, and its matching session π_j^r. We do so by initialising a post-quantum prf challenger and query $m_0 \| m_1 \| ctr$, and use the output $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B}$ from the prf challenger to replace the computation of $SecState', sk_A, sk_B$. Since $\widetilde{k_3}$ is uniformly random and independent by the previous game, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} = \mathsf{PRF}(\widetilde{k_3}, m_0 \| m_1 \| ctr)$ and we are in the previous game. If the test bit sampled by the prf challenger is 1, then $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 10'**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the post-quantum prf security of PRF, and we find:

$$\Pr(break_{9'}) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_{10'}).$$

Since $\widetilde{sk_A}, \widetilde{sk_B}$ are now uniformly random and independent values in stage t independent of the protocol flow regardless of the value of the test bit b, $\mathcal{A}$ has no advantage in guessing the bit and taking the maximum over all pairs of values t, t^* we find:

$$\begin{aligned}
&\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{q_{\mathsf{HAKE}}},\mathcal{A},C_{3.3}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) \leq \\
&n_P^2 n_S^2 n_T^2 \cdot \big(\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{KEM},\mathcal{A}}(\lambda) + (3+n_T) \cdot \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) \\
&\qquad + (3+n_T) \cdot \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda)\big).
\end{aligned}$$

3.4: $\mathsf{CompromiseSK}(i, s, t^*)$, $\mathsf{CompromiseSK}(j, r, t^*)$ have not been issued, where π_j^r *matches* π_i^s in stages u such that $t' \leq u < t$ and no $\mathsf{CompromiseSS}(i, s, u)$, $\mathsf{CompromiseSS}(j, r, u)$ queries have been issued.

Game 0 This is the standard HAKE security game. Thus:

$$\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{q\mathsf{HAKE}},\mathcal{A},C_{3.4}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) = \Pr(break_0).$$

Game 1 In this game, we guess the index (i, s, t) and the matching session (j, r, t) and abort if, during the execution of the experiment, a query $\mathsf{Test}(i', s', t')$ is received to a session $\pi_{i'}^{s'}$ such that $\pi_{j'}^{r'}$ matches $\pi_{i'}^{s'}$ in stage t' and $(i, s, t), (j, r) \neq (i', s', t'), (j', r')$. In addition, we also guess the stage t^* such that $\mathcal{A}$ does not issue either a $\mathsf{CompromiseSK}(i, s, t*)$, $\mathsf{CompromiseSK}(j, r, t^*)$ and π_i^s matches in stages u such that $t^* \leq u < t$ and no $\mathsf{CompromiseSS}(i, s, u)$ or $\mathsf{CompromiseSS}(j, r, u)$ were made. By the definition of this subcase, such a stage *must* exist. Thus:

$$\Pr(break_0) \leq n_P^2 n_S^2 n_T^2 \cdot \Pr(break_1).$$

Game 2 In this game we replace the computation of the third chaining key $k_2 \leftarrow \mathsf{PRF}(qkm, k_1)$ with a uniformly random and independent value $\widetilde{k_2} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t^*, and its matching session π_j^r. We do so by initialising a post-quantum prf challenger and query k_1, and use the output $\widetilde{k_2}$ from the prf challenger to replace the computation of k_2. Since qkm is uniformly random and independently generated by $\mathsf{ESKeyGen}$ at the start of the experiment, and by the definition of this subcase $\mathcal{A}$ cannot have issued either $\mathsf{CompromiseSK}(i, s, t^*)$ or $\mathsf{CompromiseSK}(j, r, t^*)$ this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{k_2} = \mathsf{PRF}(qkm, k_1)$ and we are in **Game 1**. If the test bit sampled by the prf challenger is 1, then $\widetilde{k_2} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 2**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the post-quantum prf security of PRF, and we find:

$$\Pr(break_1) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_2).$$

Game 3 In this game we replace the computation of the final chaining key $k_3 \leftarrow \mathsf{PRF}(SecState, \widetilde{k_2})$ with a uniformly random and independent value $\widetilde{k_3} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t^*, and its matching session π_j^r. We do so by initialising a post-quantum $\mathsf{dual\text{-}prf}$ challenger and query $SecState$, and use the output $\widetilde{k_3}$ from the $\mathsf{dual\text{-}prf}$ challenger to replace the computation of k_3. Since $\widetilde{k_2}$ is uniformly random and independent by **Game 2**, this is a sound replacement. If the test bit sampled by the $\mathsf{dual\text{-}prf}$ challenger is 0, then $\widetilde{k_3} = \mathsf{PRF}(SecState, \widetilde{k_2})$ and we are in **Game 2**. If the test bit sampled by the $\mathsf{dual\text{-}prf}$ challenger is 1, then $\widetilde{k_3} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 3**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change

can be turned into a successful adversary against the post-quantum dual-prf security of PRF, and we find:

$$\Pr(break_2) \leq \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_3).$$

Game 4 In this game we replace the computation of the updated secret state and session keys $SecState', sk_A, sk_B \leftarrow \mathsf{PRF}(\widetilde{k_3}, m_0 \| m_1 \| ctr)$ with uniformly random and independent values $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t^*, and its matching session π_j^r. We do so by initialising a post-quantum prf challenger and query $m_0 \| m_1 \| ctr$, and use the output $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B}$ from the prf challenger to replace the computation of $SecState', sk_A, sk_B$. Since $\widetilde{k_3}$ is uniformly random and independent by **Game 3**, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} = \mathsf{PRF}(\widetilde{k_3}, m_0 \| m_1 \| ctr)$ and we are in **Game 3**. If the test bit sampled by the prf challenger is 1, then $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 4**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the post-quantum prf security of PRF, and we find:

$$\Pr(break_3) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_4).$$

At this point we now iteratively do the following two game hops $(t - t^*)$ times in each consecutive stage u where $t^* < u \leq t$. Note that, by the definition of this sub-case, each of the stages following stage t^* must match (i.e. there must exist a matching session π_j^r with π_i^s that agree upon the message transcript):

Game 5' In this game we replace the computation of the final chaining key $k_3 \leftarrow \mathsf{PRF}(\widetilde{SecState}, k_2)$ with a uniformly random and independent value $\widetilde{k_3} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage u, and its matching session π_j^r. We do so by initialising a post-quantum prf challenger and query k_2, and use the output $\widetilde{k_3}$ from the prf challenger to replace the computation of k_3 in stage u. Since $\widetilde{SecState}$ is uniformly random and independent by the previous game, and $\mathcal{A}$ cannot issue a $\mathsf{CompromiseSS}(i, s, u)$ or $\mathsf{CompromiseSS}(j, r, u)$ $\forall t^* \leq u < t$, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{k_3} = \mathsf{PRF}(\widetilde{SecState}, k_2)$ and we are in the previous game. If the test bit sampled by the prf challenger is 1, then $\widetilde{k_3} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 5'**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the post-quantum prf security of PRF, and we find:

$$\Pr(break_4) \leq \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_{5'}),$$
$$\Pr(break_{6'}) \leq \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_{5'}).$$

Game 6' In this game we replace the computation of the updated secret state and session keys $SecState', sk_A, sk_B \leftarrow \mathsf{PRF}(\widetilde{k_3}, m_0 \| m_1 \| ctr)$ with uniformly random and independent values $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage u, and its matching session π_j^r. We do so by initialising a post-quantum prf challenger and query $m_0 \| m_1 \| ctr$, and use the output $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B}$ from the prf challenger to replace the computation of $SecState', sk_A, sk_B$. Since $\widetilde{k_3}$ is uniformly random and independent by the previous game, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} = \mathsf{PRF}(\widetilde{k_3}, m_0 \| m_1 \| ctr)$ and we are in the previous game. If the test bit sampled by the prf challenger is 1, then $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 6'**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the post-quantum prf security of PRF, and we find:

$$\Pr(break_{5'}) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_{6'}).$$

Since $\widetilde{sk_A}, \widetilde{sk_B}$ are now uniformly random and independent values in stage t independent of the protocol flow regardless of the value of the test bit b, $\mathcal{A}$ has no advantage in guessing the bit and thus:

$$\begin{aligned}\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{q\mathsf{HAKE}},\mathcal{A},C_{3.4}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) \leq \\ n_P^2 n_S^2 n_T \cdot \big((1+t-t^*) \cdot \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) \\ + (4+t-t^*) \cdot \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda)\big).\end{aligned}$$

Taking the maximum of all pairs of t, t^* we find:

$$\begin{aligned}\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{q\mathsf{HAKE}},\mathcal{A},C_{3.4}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) \leq \\ n_P^2 n_S^2 n_T \cdot \big((1+n_T) \cdot \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) \\ + (4+n_T) \cdot \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda)\big).\end{aligned}$$

5.1 Classical Security of Muckle

Here we prove our second result: the security of Muckle against classical adversaries.

Theorem 2. *The* Muckle *key exchange protocol is* HAKE*-secure with cleanness predicate* $\mathsf{clean}_{\mathsf{HAKE}}$ *(capturing perfect forward secrecy and post-compromise security) under the* prf, eufqcma, dual-prf, ddh *and* ssddh *assumptions. That is, for any PPT algorithm* $\mathcal{A}$ *against the* HAKE *key-indistinguishability game (defined*

in Figure 5) $\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{\mathsf{cHAKE}},\mathcal{A}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda)$ *is negligible, with:*

$$\begin{aligned}
&\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{\mathsf{cHAKE}},\mathcal{A}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) \leq \\
&2 \cdot n_P^2 n_S n_T \cdot \big(\mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \mathsf{Adv}^{\mathsf{eufcma}}_{\mathsf{MAC},\mathcal{A}}(\lambda)\big) \\
&+ n_P^2 n_S^2 n_T \cdot \big(\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{KEM},\mathcal{A}}(\lambda) + (13 + 2 \cdot n_T) \cdot \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) \\
&+ \mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{KEM},\mathcal{A}}(\lambda) + (13 + 2 \cdot n_T) \cdot \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda)\big)
\end{aligned}$$

We give the full proof of our secondary result, Theorem 2, against a classical adversary (i.e. an adversary that is a classical PPT algorithm):

Proof. Similar to the proof of Theorem 1 we begin by dividing the proof into three separate cases (and denote with $\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{\mathsf{cHAKE}},\mathcal{A},C_l}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda)$ the advantage of the adversary in winning the key-indistinguishability game in Case l) where the query $\mathsf{Test}(i,s,t)$ has been issued:

1. The session π_i^s (where $\pi_i^s.\rho = \texttt{init}$) has no origin session in stage t.
2. The session π_i^s (where $\pi_i^s.\rho = \texttt{resp}$) has no origin session in stage t.
3. The session π_i^s in stage t has a matching session.

It follows then that

$$\begin{aligned}
\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{\mathsf{cHAKE}},\mathcal{A}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) \leq\ & \mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{\mathsf{cHAKE}},\mathcal{A},C_1}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) \\
&+ \mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{\mathsf{cHAKE}},\mathcal{A},C_2}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) \\
&+ \mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{\mathsf{cHAKE}},\mathcal{A},C_3}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda)
\end{aligned}$$

We then bound the probability of each case, and show that under certain assumptions, the probability of the adversary winning in the key-indistinguishability game is negligible.

The proofs of Case 1 and Case 2 are identical to the proof of Case 1 and Case 2 in Theorem 1, and so we simply state the bounds below:

$$\begin{aligned}
\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{\mathsf{cHAKE}},\mathcal{A},C_1}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) &\leq n_P^2 n_S n_T \big(\mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) \\
&\quad + \mathsf{Adv}^{\mathsf{eufcma}}_{\mathsf{MAC},\mathcal{A}}(\lambda)\big), \\
\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{\mathsf{cHAKE}},\mathcal{A},C_2}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) &\leq n_P^2 n_S n_T \cdot \big(\mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) \\
&\quad + \mathsf{Adv}^{\mathsf{eufcma}}_{\mathsf{MAC},\mathcal{A}}(\lambda)\big).
\end{aligned}$$

Case 3: Test session with matching session

In Case 3, we show that if $\mathcal{A}$ that has issued a $\mathsf{Test}(i,s,t)$ query to a clean session π_i^s in stage t, then $\mathcal{A}$ has negligible advantage in guessing the test bit b. In what follows, we split our analysis of Case 3 into the following sub-cases, each corresponding to a condition necessary for the cleanness predicate clean to be upheld by π_i^s in stage t. These are the subcases:

3.1 $\mathsf{CompromiseQK}(i,s,t)$, $\mathsf{CompromiseQK}(j,r,t)$ have not been issued, where π_j^r *matches* π_i^s in stage t.
3.2 $\mathsf{CompromiseSK}(i,s,t)$, $\mathsf{CompromiseSK}(j,r,t)$ have not been issued, where π_j^r *matches* π_i^s in stage t.
3.3 $\mathsf{CompromiseCK}(i,s,t)$, $\mathsf{CompromiseCK}(j,r,t)$ have not been issued, where π_j^r *matches* π_i^s in stage t.
3.4 $\mathsf{CompromiseQK}(i,s,t')$, $\mathsf{CompromiseQK}(j,r,t')$ have not been issued, where π_j^r *matches* π_i^s in stages u such that $t' \leq u < t$ and no $\mathsf{CompromiseSS}(i,s,u)$, $\mathsf{CompromiseSS}(j,r,u)$ queries have been issued.
3.5 $\mathsf{CompromiseSK}(i,s,t')$, $\mathsf{CompromiseSK}(j,r,t')$ have not been issued, where π_j^r *matches* π_i^s in stages u such that $t' \leq u < t$ and no $\mathsf{CompromiseSS}(i,s,u)$, $\mathsf{CompromiseSS}(j,r,u)$ queries have been issued.
3.6 $\mathsf{CompromiseCK}(i,s,t')$, $\mathsf{CompromiseCK}(j,r,t')$ have not been issued, where π_j^r *matches* π_i^s in stages u such that $t' \leq u < t$ and no $\mathsf{CompromiseSS}(i,s,u)$, $\mathsf{CompromiseSS}(j,r,u)$ queries have been issued.

Note that subcases 3.1, 3.2, 3.4 and 3.5 are exactly identical to the subcases of Case 3 in the proof of Theorem 1, and the proofs follow similarly. To save space we do not reiterate the proofs, but merely give the advantage statements below.

$$\begin{aligned}
\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{\mathsf{cHAKE}},\mathcal{A},C_{3.1}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) &\leq n_P^2 n_S^2 n_T \cdot \big(\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{KEM},\mathcal{A}}(\lambda) \\
&\quad + 3 \cdot \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) \\
&\quad + 3 \cdot \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda)\big), \\
\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{\mathsf{cHAKE}},\mathcal{A},C_{3.2}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) &\leq n_P^2 n_S^2 n_T \cdot \big(2 \cdot \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) \\
&\quad + \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda)\big), \\
\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{\mathsf{cHAKE}},\mathcal{A},C_{3.4}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) &\leq n_P^2 n_S^2 n_T^2 \cdot \big(\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{KEM},\mathcal{A}}(\lambda) \\
&\quad + (3+n_T) \cdot \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) \\
&\quad + (3+n_T) \cdot \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda)\big), \\
\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{\mathsf{cHAKE}},\mathcal{A},C_{3.5}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) &\leq \\
n_P^2 n_S^2 n_T \cdot \big((1+n_T) &\cdot \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) \\
&\quad + (4+n_T) \cdot \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda)\big).
\end{aligned}$$

We now turn to showing that the adversary's advantage in winning in subcase 3.3 is negligible under the ind-cpa security of KEM, and the prf and dual-prf assumptions.

3.3: $\mathsf{CompromiseCK}(i,s,t)$, $\mathsf{CompromiseCK}(j,r,t)$ have not been issued, where π_j^r *matches* π_i^s in stage t.

Game 0 This is the standard HAKE security game with cleanness predicate $\mathsf{clean}_{c\mathsf{HAKE}}$. Thus:

$$\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{c\mathsf{HAKE}},\mathcal{A},C_{3.3}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) = \Pr(break_0).$$

Game 1 In this game, we guess the index (i,s,t) and the matching session (j,r,t) and abort if, during the execution of the experiment, a query $\mathsf{Test}(i',s',t')$ is received to a session $\pi^{s'}_{i'}$ such that $\pi^{r'}_{j'}$ matches $\pi^{s'}_{i'}$ in stage t' and $(i,s,t),(j,r) \neq (i',s',t'),(j',r')$. Thus:

$$\Pr(break_0) \leq n_P^2 n_S^2 n_T \cdot \Pr(break_1).$$

Game 2 In this game we replace the computation of the classical Diffie-Hellman key $dhk = g^{xy}$ with the value g^z, where z is a uniformly-randomly sampled from $\mathbb{Z}_q$ used in the protocol execution of the test session π^s_i in stage t, and its matching session π^r_j. We do so by initialising a ddh challenger, and replace the g^x, g^y values sent in m_0 and m_1 (by the test session and its matching partner in stage t, which by the previous game, we know at the start of the experiment). By the definition of this subcase, $\mathcal{A}$ cannot have issued either a $\mathsf{CompromiseCK}(i,s,t)$ or $\mathsf{CompromiseCK}(j,r,t)$ query, nor inject its own Diffie-Hellman keyshares, this replacement is sound. If the distribution sampled by the ddh challenger is (g^x, g^y, g^{xy}), we are in **Game 1**. If the distribution sampled by the ddh challenger is (g^x, g^y, g^z), then we are in **Game 2**. Thus any adversary capable of distinguishing this change can be turned into a distinguishing algorithm against the ddh assumption, and we find:

$$\Pr(break_1) \leq \mathsf{Adv}^{\mathsf{ddh}}_{q,p,g,\mathcal{A}}(\lambda) + \Pr(break_2).$$

Game 3 In this game we replace the computation of the classical key output $ck \leftarrow \mathsf{PRF}(dhk, \mathsf{label}_{ck})$ with a uniformly random and independent value $\widetilde{ck} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π^s_i in stage t, and its matching session π^r_j. We do so by initialising a prf challenger and query label_{ck}, and use the output $\widetilde{ck}$ from the prf challenger to replace the computation of ck. Since $\mathbb{Z}_q$ is uniformly random and independently sampled from $\mathbb{Z}_q$, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{ck} = \mathsf{PRF}(dhk, \mathsf{label}_{ck})$ and we are in **Game 2**. If the test bit sampled by the prf challenger is 1, then $\widetilde{ck} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 3**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the prf security of PRF, and we find:

$$\Pr(break_2) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_3).$$

Game 4 In this game we replace the computation of the second chaining key $k_1 \leftarrow \mathsf{PRF}(\widetilde{ck}, k_0)$ with a uniformly random and independent value $\widetilde{k_1} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in

the protocol execution of the test session π_i^s in stage t, and its matching session π_j^r. We do so by initialising a prf challenger and query k_0, and use the output $\widetilde{k_1}$ from the prf challenger to replace the computation of k_1. Since $\widetilde{ck}$ is uniformly random and independent by **Game 3**, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{k_1} = \mathsf{PRF}(ck, k_0)$ and we are in **Game 3**. If the test bit sampled by the prf challenger is 1, then $\widetilde{k_1} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 4**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the prf security of PRF, and we find:

$$\Pr(break_3) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_4).$$

Game 5 In this game we replace the computation of the third chaining key $k_2 \leftarrow \mathsf{PRF}(qkm, \widetilde{k_1})$ with a uniformly random and independent value $\widetilde{k_2} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t, and its matching session π_j^r. We do so by initialising a $\mathsf{dual\text{-}prf}$ challenger and query qkm, and use the output $\widetilde{k_2}$ from the $\mathsf{dual\text{-}prf}$ challenger to replace the computation of k_2. Since $\widetilde{k_1}$ is uniformly random and independent by **Game 4**, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{k_2} = \mathsf{PRF}(qkm, \widetilde{k_1})$ and we are in **Game 4**. If the test bit sampled by the $\mathsf{dual\text{-}prf}$ challenger is 1, then $\widetilde{k_2} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 5**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the prf security of PRF, and we find:

$$\Pr(break_4) \leq \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_5).$$

Game 6 In this game we replace the computation of the final chaining key $k_3 \leftarrow \mathsf{PRF}(SecState, \widetilde{k_2})$ with a uniformly random and independent value $\widetilde{k_3} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t, and its matching session π_j^r. We do so by initialising a $\mathsf{dual\text{-}prf}$ challenger and query $SecState$, and use the output $\widetilde{k_3}$ from the $\mathsf{dual\text{-}prf}$ challenger to replace the computation of k_3. Since $\widetilde{k_2}$ is uniformly random and independent by **Game 5**, this is a sound replacement. If the test bit sampled by the $\mathsf{dual\text{-}prf}$ challenger is 0, then $\widetilde{k_3} = \mathsf{PRF}(SecState, \widetilde{k_2})$ and we are in **Game 2**. If the test bit sampled by the $\mathsf{dual\text{-}prf}$ challenger is 1, then $\widetilde{k_3} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 3**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the $\mathsf{dual\text{-}prf}$ security of PRF, and we find:

$$\Pr(break_5) \leq \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_6).$$

Game 7 In this game we replace the computation of the updated secret state and session keys $SecState', sk_A, sk_B \leftarrow \mathsf{PRF}(\widetilde{k_3}, m_0 \| m_1 \| ctr)$ with uniformly random and independent values $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where

$\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t, and its matching session π_j^r. We do so by initialising a prf challenger and query $m_0 \| m_1 \| ctr$, and use the output $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B}$ from the prf challenger to replace the computation of $SecState', sk_A, sk_B$. Since $\widetilde{k_3}$ is uniformly random and independent by **Game 6**, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} = \mathsf{PRF}(\widetilde{k_3}, m_0 \| m_1 \| ctr)$ and we are in **Game 6**. If the test bit sampled by the prf challenger is 1, then $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 7**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the prf security of PRF, and we find:

$$\Pr(break_6) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_7).$$

Since $\widetilde{sk_A}, \widetilde{sk_B}$ are now uniformly random and independent values independent of the protocol flow regardless of the value of the test bit b, $\mathcal{A}$ has no advantage in guessing the bit and thus:

$$\begin{aligned}\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{c\mathsf{HAKE}},\mathcal{A},C_{3.3}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) \leq n_P^2 n_S^2 n_T \cdot \big(&\mathsf{Adv}^{\mathsf{ddh}}_{q,p,g,\mathcal{A}}(\lambda) \\ &+ 3 \cdot \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) \\ &+ 2 \cdot \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda)\big).\end{aligned}$$

3.6: $\mathsf{CompromiseCK}(i,s,t^*)$, $\mathsf{CompromiseCK}(j,r,t^*)$ have not been issued, where π_j^r *matches* π_i^s in stages u such that $t' \leq u < t$ and no $\mathsf{CompromiseSS}(i,s,u)$, $\mathsf{CompromiseSS}(j,r,u)$ queries have been issued.

Game 0 This is the standard HAKE security game with cleanness predicate $\mathsf{clean}_{c\mathsf{HAKE}}$. Thus:

$$\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{c\mathsf{HAKE}},\mathcal{A},C_{3.6}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) = \Pr(break_0).$$

Game 1 In this game, we guess the index (i,s,t) and the matching session (j,r,t) and abort if, during the execution of the experiment, a query $\mathsf{Test}(i',s',t')$ is received to a session $\pi_{i'}^{s'}$ such that $\pi_{j'}^{r'}$ matches $\pi_{i'}^{s'}$ in stage t' and $(i,s,t),(j,r) \neq (i',s',t'),(j',r')$. In addition, we also guess the stage t^* such that $\mathcal{A}$ does not issue either a $\mathsf{CompromiseSK}(i,s,t^*)$, $\mathsf{CompromiseSK}(j,r,t^*)$ and π_i^s matches in stages u such that $t^* \leq u < t$ and no $\mathsf{CompromiseSS}(i,s,u)$ or $\mathsf{CompromiseSS}(j,r,u)$ were made. By the definition of this subcase, such a stage *must* exist. Thus:

$$\Pr(break_0) \leq n_P^2 n_S^2 n_T^2 \cdot \Pr(break_1).$$

Game 2 In this game, we replace the classical decapsulated key k derived in the test session π_i^s with the uniformly random and independent value $\widetilde{k}$. We do so by interacting with a $\mathsf{ind\text{-}cpa}$ KEM challenger (as described in Definition 5) and replace the pk_A value sent in m_0, and the ciphertext pk_B sent in m_1

with the public-key pk and the ciphertext c received from the $\mathsf{ind\text{-}cpa}$ KEM challenger. Since π_i^s matches π_j^r in stage t^*, we know that the public-key and ciphertext sent in m_0 and m_1 respectively were received by the sessions without modification. By the definition of this subcase, $\mathcal{A}$ cannot have issued either a $\mathsf{CompromiseCK}(i, s, t^*)$ or $\mathsf{CompromiseCK}(j, r, t^*)$ queries, not inject its own public-key and ciphertext, and thus this replacement is sound. Detecting the replacement of k with a uniformly random value $\widetilde{k}$ implies an efficient distinguishing algorithm against the classical $\mathsf{ind\text{-}cpa}$ security of KEM. Thus:

$$\Pr(break_1) \leq \mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{KEM},\mathcal{A}}(\lambda) + \Pr(break_2).$$

Game 3 In this game we replace the computation of the classical key output $ck \leftarrow \mathsf{PRF}(\widetilde{k}, \mathsf{label}_{ck})$ with a uniformly random and independent value $\widetilde{ck} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t^*, and its matching session π_j^r. We do so by initialising a prf challenger and query label_{ck}, and use the output $\widetilde{ck}$ from the prf challenger to replace the computation of ck. Since $\widetilde{k}$ is uniformly random and independent by the previous game, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{ck} = \mathsf{PRF}(\widetilde{k}, \mathsf{label}_{ck})$ and we are in **Game 2**. If the test bit sampled by the prf challenger is 1, then $\widetilde{ck} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 3**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the prf security of PRF, and we find:

$$\Pr(break_2) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_3).$$

Game 4 In this game we replace the computation of the second chaining key $k_1 \leftarrow \mathsf{PRF}(\widetilde{ck}, k_0)$ with a uniformly random and independent value $\widetilde{k_1} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t^*, and its matching session π_j^r. We do so by initialising a prf challenger and query k_0, and use the output $\widetilde{k_1}$ from the prf challenger to replace the computation of k_1. Since $\widetilde{ck}$ is uniformly random and independent by **Game 3**, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{k_1} = \mathsf{PRF}(\widetilde{ck}, k_0)$ and we are in **Game 3**. If the test bit sampled by the prf challenger is 1, then $\widetilde{k_1} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 4**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the prf security of PRF, and we find:

$$\Pr(break_3) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_4).$$

Game 5 In this game we replace the computation of the third chaining key $k_2 \leftarrow \mathsf{PRF}(qkm, \widetilde{k_1})$ with a uniformly random and independent value $\widetilde{k_2} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t^*, and its matching session

π_j^r. We do so by initialising a dual-prf challenger and query qkm, and use the output $\widetilde{k_2}$ from the dual-prf challenger to replace the computation of k_2. Since $\widetilde{k_1}$ is uniformly random and independent by **Game 4**, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{k_2} = \mathsf{PRF}(qkm, \widetilde{k_1})$ and we are in **Game 4**. If the test bit sampled by the dual-prf challenger is 1, then $\widetilde{k_2} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 5**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the prf security of PRF, and we find:

$$\Pr(break_4) \leq \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_5).$$

Game 6 In this game we replace the computation of the final chaining key $k_3 \leftarrow \mathsf{PRF}(SecState, \widetilde{k_2})$ with a uniformly random and independent value $\widetilde{k_3} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t^*, and its matching session π_j^r. We do so by initialising a dual-prf challenger and query $SecState$, and use the output $\widetilde{k_3}$ from the dual-prf challenger to replace the computation of k_3. Since $\widetilde{k_2}$ is uniformly random and independent by **Game 5**, this is a sound replacement. If the test bit sampled by the dual-prf challenger is 0, then $\widetilde{k_3} = \mathsf{PRF}(SecState, \widetilde{k_2})$ and we are in **Game 2**. If the test bit sampled by the dual-prf challenger is 1, then $\widetilde{k_3} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 3**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the dual-prf security of PRF, and we find:

$$\Pr(break_5) \leq \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_6).$$

Game 7 In this game we replace the computation of the updated secret state and session keys $SecState', sk_A, sk_B \leftarrow \mathsf{PRF}(\widetilde{k_3}, m_0 \| m_1 \| ctr)$ with uniformly random and independent values $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage t^*, and its matching session π_j^r. We do so by initialising a prf challenger and query $m_0 \| m_1 \| ctr$, and use the output $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B}$ from the prf challenger to replace the computation of $SecState', sk_A, sk_B$. Since $\widetilde{k_3}$ is uniformly random and independent by **Game 6**, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} = \mathsf{PRF}(\widetilde{k_3}, m_0 \| m_1 \| ctr)$ and we are in **Game 6**. If the test bit sampled by the prf challenger is 1, then $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 7**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the prf security of PRF, and we find:

$$\Pr(break_6) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_7).$$

At this point we now iteratively do the following two game hops $(t - t^*)$ times in each consecutive stage u where $t^* < u \leq t$. Note that, by the definition of this

sub-case, each of the stages following stage t^* must match (i.e. there must exist a matching session π_j^r with π_i^s that agrees upon the message transcript):

Game 8' In this game we replace the computation of the final chaining key $k_3 \leftarrow \mathsf{PRF}(\widetilde{SecState}, k_2)$ with a uniformly random and independent value $\widetilde{k_3} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage u, and its matching session π_j^r. We do so by initialising a prf challenger and query k_2, and use the output $\widetilde{k_3}$ from the prf challenger to replace the computation of k_3 in stage u. Since $\widetilde{SecState}$ is uniformly random and independent by the previous game, and $\mathcal{A}$ cannot issue a $\mathsf{CompromiseSS}(i,s,u)$ or $\mathsf{CompromiseSS}(j,r,u)$ $\forall t^* \leq u < t$, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{k_3} = \mathsf{PRF}(\widetilde{SecState}, k_2)$ and we are in the previous game. If the test bit sampled by the prf challenger is 1, then $\widetilde{k_3} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 8′**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the prf security of PRF, and we find:

$$\Pr(break_7) \leq \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_{8'}),$$
$$\Pr(break_{9'}) \leq \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_{8'}).$$

Game 9' In this game we replace the computation of the updated secret state and session keys $SecState', sk_A, sk_B \leftarrow \mathsf{PRF}(\widetilde{k_3}, m_0 \| m_1 \| ctr)$ with uniformly random and independent values $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ (where $\{0,1\}^{\mathsf{PRF}}$ is the output space of the PRF) used in the protocol execution of the test session π_i^s in stage u, and its matching session π_j^r. We do so by initialising a prf challenger and query $m_0 \| m_1 \| ctr$, and use the output $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B}$ from the prf challenger to replace the computation of $SecState', sk_A, sk_B$. Since $\widetilde{k_3}$ is uniformly random and independent by the previous game, this is a sound replacement. If the test bit sampled by the prf challenger is 0, then $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} = \mathsf{PRF}(\widetilde{k_3}, m_0 \| m_1 \| ctr)$ and we are in the previous game. If the test bit sampled by the prf challenger is 1, then $\widetilde{SecState'}, \widetilde{sk_A}, \widetilde{sk_B} \xleftarrow{\$} \{0,1\}^{\mathsf{PRF}}$ and we are in **Game 9′**. Thus any adversary $\mathcal{A}$ capable of distinguishing this change can be turned into a successful adversary against the prf security of PRF, and we find:

$$\Pr(break_{8'}) \leq \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) + \Pr(break_{9'}).$$

After $(t - t^*)$ iterations of the previous two games, $\widetilde{sk_A}, \widetilde{sk_B}$ are now uniformly random and independent values in stage t independent of the protocol flow regardless of the value of the test bit b, $\mathcal{A}$ has no advantage in guessing the bit and thus:

$$\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{\mathsf{cHAKE}},\mathcal{A},C_{3.6}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) \leq$$
$$n_P^2 n_S^2 n_T^2 \cdot \big(\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{KEM},\mathcal{A}}(\lambda) + (2 + t - t^*) \cdot \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda)$$
$$+ (3 + t - t^*) \cdot \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda)\big).$$

Taking the maximum of all pairs of t, t^*, we find:

$$\mathsf{Adv}^{\mathsf{HAKE},\mathsf{clean}_{\mathsf{cHAKE}},\mathcal{A},C_{3.6}}_{\mathsf{Muckle},n_P,n_S,n_T}(\lambda) \leq$$
$$n_P^2 n_S^2 n_T^2 \cdot \big(\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{KEM},\mathcal{A}}(\lambda) + (2 + n_T) \cdot \mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda)$$
$$+ (3 + n_T) \cdot \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda)\big).$$

6 Conclusion and Future Work

In this paper, we have given a framework for the analysis of hybrid key exchange protocols. We have illustrated its usage with our analysis of the Muckle protocol, which combines QKD, classical and and post-quantum components to produce a protocol that is secure against a broad class of adversaries and robust in the face of unanticipated developments in cryptanalysis, advances in quantum computing, and vulnerabilities in immature QKD systems. We benchmarked Muckle in realistic network environments, showing that its computational cost compared to a fully classical approach is quite moderate. Muckle is a basic protocol design that can serve to guide the design of hybrid variants of more complex protocols like TLS (which must address cryptographic agility for key exchange and peer authentication methods, fast key establishment methods such as 0RTT, etc). Our paper opens up many avenues for future work. First, we have strongly abstracted the QKD component in our framework, treating is as an inexhaustible supply of shared, random bits. Yet there is a fine tradition of developing security proofs for QKD systems based purely on physical models. It is a significant challenge to integrate such approaches in our framework. The work of [25] provides one route forward using Universal Composability. In future QKD systems, the bit-rate of key agreement will exceed that which can be achieved by classical communication, at least over short ranges. This suggests adapting Muckle to allow rapid key refreshing from QKD keying material and slower refreshing from other sources. Our analysis framework could be extended to support such "differential refreshing". But this approach also raises implementation challenges, particularly around key synchronisation, which would need to be carefully addressed in order to avoid DoS and other attacks.

Acknowledgments

The research of Dowling was supported by Innovate UK and EPSRC grant EP/L018543/1 (the EQUIP project). The research of Hansen was supported by

the EPSRC and the UK government as part of the Centre for Doctoral Training in Cyber Security at Royal Holloway, University of London (EP/K035584/1). The research of Paterson was supported in part by Innovate UK and EPSRC grants EP/L018543/1, EP/K035584/1 and EP/M013472/1.

Bibliography

[1] ARM mbed TLS. `https://tls.mbed.org/`. Accessed: 12-11-2018.

[2] C-Muckle source code. `https://github.com/himsen/muckle`. Accessed: 29-01-2020.

[3] Microsoft PQCrypto-SIDH. `https://github.com/Microsoft/PQCrypto-SIDH`. Accessed: 12-11-2018.

[4] M. R. Albrecht, B. R. Curtis, A. Deo, A. Davidson, R. Player, E. W. Postlethwaite, F. Virdia, and T. Wunderer. Estimate all the {LWE, NTRU} schemes! In D. Catalano and R. D. Prisco, editors, *Security and Cryptography for Networks - 11th International Conference, SCN 2018, Amalfi, Italy, September 5-7, 2018, Proceedings*, volume 11035 of *Lecture Notes in Computer Science*, pages 351–367. Springer, 2018.

[5] E. Alkim, L. Ducas, T. Pöppelmann, and P. Schwabe. Post-quantum key exchange - A new hope. In T. Holz and S. Savage, editors, *25th USENIX Security Symposium, USENIX Security 16, Austin, TX, USA, August 10-12, 2016.*, pages 327–343. USENIX Association, 2016.

[6] M. Bellare, R. Canetti, and H. Krawczyk. Keying hash functions for message authentication. In N. Koblitz, editor, *CRYPTO'96*, volume 1109 of *LNCS*, pages 1–15. Springer, Heidelberg, Aug. 1996.

[7] M. Bellare and P. Rogaway. Entity authentication and key distribution. In D. R. Stinson, editor, *CRYPTO'93*, volume 773 of *LNCS*, pages 232–249. Springer, Heidelberg, Aug. 1994.

[8] C. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing*, 175(P1), 1984.

[9] D. J. Bernstein. Curve25519: New Diffie-Hellman speed records. In M. Yung, Y. Dodis, A. Kiayias, and T. Malkin, editors, *PKC 2006*, volume 3958 of *LNCS*, pages 207–228. Springer, Heidelberg, Apr. 2006.

[10] D. J. Bernstein. Is the security of quantum cryptography guaranteed by the laws of physics? *CoRR*, abs/1803.04520, 2018.

[11] N. Bindel, J. Brendel, M. Fischlin, B. Goncalves, and D. Stebila. Hybrid key encapsulation mechanisms and authenticated key exchange. In *International Conference on Post-Quantum Cryptography*, pages 206–226. Springer, 2019.

[12] N. Bindel, U. Herath, M. McKague, and D. Stebila. Transitioning to a quantum-resistant public key infrastructure. In T. Lange and T. Takagi, editors, *Post-Quantum Cryptography - 8th International Workshop, PQCrypto 2017, Utrecht, The Netherlands, June 26-28, 2017, Proceedings*, volume 10346 of *Lecture Notes in Computer Science*, pages 384–405. Springer, 2017.

[13] J. W. Bos, C. Costello, M. Naehrig, and D. Stebila. Post-quantum key exchange for the TLS protocol from the ring learning with errors problem. In *2015 IEEE Symposium on Security and Privacy*, pages 553–570. IEEE Computer Society Press, May 2015.

[14] M. Braithwaite. Experimenting with post-quantum cryptography. Retrieved from URL, July 2016.

[15] J. Brendel, M. Fischlin, and F. Günther. Breakdown Resilience of Key Exchange Protocols: NewHope, TLS 1.3, and Hybrids. In K. Sako, S. Schneider, and P. Y. A. Ryan, editors, *Computer Security – ESORICS 2019*, pages 521–541, Cham, 2019. Springer International Publishing.

[16] K. Cohn-Gordon, C. J. F. Cremers, and L. Garratt. On Post-compromise Security. In *IEEE 29th Computer Security Foundations Symposium, CSF 2016, Lisbon, Portugal, June 27 - July 1, 2016*, pages 164–178. IEEE Computer Society, 2016.

[17] C. Costello, P. Longa, and M. Naehrig. Efficient algorithms for supersingular isogeny diffie-hellman. In M. Robshaw and J. Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 572–601. Springer, Heidelberg, Aug. 2016.

[18] C. J. F. Cremers and M. Feltz. Beyond eCK: Perfect forward secrecy under actor compromise and ephemeral-key reveal. In S. Foresti, M. Yung, and F. Martinelli, editors, *ESORICS 2012*, volume 7459 of *LNCS*, pages 734–751. Springer, Heidelberg, Sept. 2012.

[19] A. Huang, S.-H. Sun, Z. Liu, and V. Makarov. Quantum key distribution with distinguishable decoy states. *Phys. Rev. A*, 98:012330, 2018.

[20] H. Krawczyk. Cryptographic extraction and key derivation: The HKDF scheme. In T. Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 631–648. Springer, Heidelberg, Aug. 2010.

[21] K. Kwiatkowski and L. Valenta. The tls post-quantum experiment. Retrieved from URL, October 2010.

[22] J. Li, K. Kim, F. Zhang, and X. Chen. Aggregate proxy signature and verifiably encrypted proxy signature. In W. Susilo, J. K. Liu, and Y. Mu, editors, *ProvSec 2007*, volume 4784 of *LNCS*, pages 208–217. Springer, Heidelberg, Nov. 2007.

[23] D. Moody. What was NIST thinking? Round 2 of the NIST PQC "Competition". Talk at Oxford University, 2019.

[24] M. Mosca, D. Stebila, and B. Ustaoglu. Quantum key distribution in the classical authenticated key exchange framework. In P. Gaborit, editor, *Post-Quantum Cryptography - 5th International Workshop, PQCrypto 2013, Limoges, France, June 4-7, 2013. Proceedings*, volume 7932 of *Lecture Notes in Computer Science*, pages 136–154. Springer, 2013.

[25] J. Müller-Quade and R. Renner. Composability in quantum cryptography. *CoRR*, abs/1006.2215, 2010.

[26] J. Schank and D. Stebila. A Transport Layer Security (TLS) Extension For Establishing An Additional Shared Secret. IETF Draft, 2017.

[27] P. Sibson, C. Erven, M. Godfrey, S. Miki, T. Yamashita, M. Fujiwara, M. Sasaki, H. Terai, M. G. Tanner, C. M. Natarajan, R. H. Hadfield, J. L. O'Brien, and M. G. Thompson. Chip-based quantum key distribution. *Nature Communications*, 8:13984, Feb. 2017.

[28] D. Stebila, S. Fluhrer, and S. Gueron. Design issues for hybrid key exchange in TLS 1.3. IETF Draft, 2019.

[29] A. Vakhitov, V. Makarov, and D. R. Hjelme. Large pulse attack as a method of conventional optical eavesdropping in quantum cryptography. *J. Mod. Opt.*, 48:2023, 2001.
[30] W. Whyte, S. Fluhrer, Z. Zhang, and O. Garcia-Morchon. Quantum-safe hybrid (QSH) key exchange for transport layer security (TLS) version 1.3. IETF Draft, 2017.
[31] H. P. Yuen. Security of quantum key distribution. *IEEE Access*, 4:724–749, 2016.
[32] R. Zhang, G. Hanaoka, J. Shikata, and H. Imai. On the security of multiple encryption or CCA-security+CCA-security=CCA-security? In F. Bao, R. Deng, and J. Zhou, editors, *PKC 2004*, volume 2947 of *LNCS*, pages 360–374. Springer, Heidelberg, Mar. 2004.

A Security Assumptions

In this section we introduce many security assumptions that we rely upon in proving our construction secure. In particular, for the derivation of the keys we rely on post-quantum variants of pseudo-random functions and dual pseudo-random functions security. For our key-exchange, we rely on the classical security of a key encapsulation mechanism (KEMs), and additionally the post-quantum security of a key encapsulation mechanism (QKEM). Finally, for authentication, we rely upon the post-quantum variant of the existential unforgeable security under chosen message attack (eufcma) assumption for a Message Authentication Code (MAC). We begin by introducing "post-quantum" variant, explicitly requiring that the symmetric key algorithm is secure against an algorithm that has access to a quantum computer.

Definition 5 (Key Encapsulation Mechanism). *A key encapsulation mechanism (KEM) is a triple of algorithms* $\mathsf{KEM} = \{\mathsf{KeyGen}, \mathsf{Encaps}, \mathsf{Decaps}\}$ *with an associated keyspace* $\mathcal{K}$*. We describe the algorithms below:*

- $\mathsf{KeyGen}(\lambda) \xrightarrow{\$} (pk, sk)$: KeyGen *is a probabilistic algorithm that takes as input the security parameter* λ *and returns a public/secret key pair* (pk, sk)*.*
- $\mathsf{Encaps}(pk) \xrightarrow{\$} (c, k)$: Encaps *is a probabilistic algorithm that takes as input a public key* pk *and outputs a ciphertext* c *as well as a key* $k \in \mathcal{K}$*.*
- $\mathsf{Decaps}(sk, c) \to (k)$: Decaps *is a deterministic algorithm that takes as input a secret key* sk *and a ciphertext* c *and outputs a key* $k \in \mathcal{K}$*, or a failure symbol* $\perp$*.*

KEM *is correct if* $\forall (pk, sk)$ *such that* $\mathsf{KeyGen}(\lambda) \xrightarrow{\$} (pk, sk)$*, and* (c, k) *such that* $\mathsf{Encaps}(pk) \xrightarrow{\$} (c, k)$*, it holds that* $\mathsf{Decaps}(sk, c) = k$*. We define the* $\mathsf{ind\text{-}cpa}$ *security of a key encapsulation mechanism in the following game played between a challenger* $\mathcal{C}$ *and an adversary* $\mathcal{A}$*.*

1. $\mathcal{C}$ *generates a public-key pair* $\mathsf{KeyGen}(\lambda) \xrightarrow{\$} (pk, sk)$

2. $\mathcal{C}$ *generates a ciphertext and key* $\mathsf{Encaps}(pk) \xrightarrow{\$} (c, k_0)$
3. $\mathcal{C}$ *samples a key* $k_1 \xleftarrow{\$} \mathcal{K}$ *and a bit* b *uniformly at random.*
4. $\mathcal{A}$ *is given* (pk, c, k_b) *and outputs a guess bit* b'

We say that $\mathcal{A}$ *wins the* ind-cpa *security game if* $b' = b$ *and define the advantage of a QPT algorithm* $\mathcal{A}$ *in breaking the* ind-cpa *security of a key encapsulation mechanism* KEM *as* $\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{KEM},\mathcal{A}}(\lambda) = |2 \cdot \Pr(b' = b) - 1|$. *We say that* KEM *is post-quantum secure if for all QPT algorithms* $\mathcal{A}$, $\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{KEM},\mathcal{A}}(\lambda)$ *is negligible in the security parameter* λ.

Definition 6 (Post-Quantum PRF). *A pseudo-random function family is a collection of deterministic functions* $\mathsf{PRF} = \{\mathsf{PRF}_\lambda : \mathcal{K} \times \mathcal{I} \to \mathcal{O} : \lambda \in \mathbb{N}\}$, *one function for each value of* λ. *Here,* $\mathcal{K}$, $\mathcal{I}$, $\mathcal{O}$ *all depend on* λ, *but we suppress this for ease of notation. Given a key* k *in the keyspace* $\mathcal{K}$ *and a bit string* $m \in \mathcal{M}$, PRF_λ *outputs a value* y *in the output space* $\mathcal{O} = \{0,1\}^\lambda$. *We define the security of a pseudo-random function family in the following game between a challenger* $\mathcal{C}$ *and a quantum polynomial-time (QPT) algorithm* $\mathcal{A}$[11], *with* λ *as an implicit input to both algorithms:*

1. $\mathcal{C}$ *samples a key* $k \xleftarrow{\$} \mathcal{K}$ *and a bit* b *uniformly at random.*
2. $\mathcal{A}$ *can now query* $\mathcal{C}$ *with polynomially-many distinct* m_i *values, and receives either the output* $y_i \leftarrow \mathsf{PRF}_\lambda(k, m_i)$ *(when* $b = 0$*) or* $y_i \xleftarrow{\$} \{0,1\}^\lambda$ *(when* $b = 1$*).*
3. $\mathcal{A}$ *terminates and outputs a bit* b'.

We say that $\mathcal{A}$ *wins the* PRF *security game if* $b' = b$ *and define the advantage of a QPT algorithm* $\mathcal{A}$ *in breaking the* pseudo-random function security *of a PRF family* PRF *as* $\mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) = |2 \cdot \Pr(b' = b) - 1|$. *We say that* PRF *is post-quantum secure if for all QPT algorithms* $\mathcal{A}$, $\mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda)$ *is negligible in the security parameter* λ.

We similarly upgrade the Dual-PRF assumption and the eufcma assumption for MACs to post-quantum variants.

Definition 7 (Post-Quantum Dual-PRF). *Let* PRF_λ *be a PRF family. Given a key* k *in the keyspace* $\mathcal{K}$ *and a bit string* $m \in \mathcal{M}$, PRF_λ *outputs a value* y *in the output space* $\mathcal{O} = \{0,1\}^\lambda$. *We define a second PRF family* $\mathsf{PRF}^{\mathsf{dual}} = \{\mathsf{PRF}^{\mathsf{dual}}_\lambda : \mathcal{I} \times \mathcal{K} \to \mathcal{O} : \lambda \in \mathbb{N}\}$ *by setting* $\mathsf{PRF}^{\mathsf{dual}}_\lambda(m, k) = \mathsf{PRF}_\lambda(k, m)$. *We define the security of a pseudo-random function family in the following game between a challenger* $\mathcal{C}$ *and a QPT adversary* $\mathcal{A}$, *with* λ *as an implicit input to both algorithms:*

1. $\mathcal{C}$ *samples a key* $k \xleftarrow{\$} \mathcal{K}$ *and a bit* b *uniformly at random.*

[11] Note that a quantum polynomial-time algorithm is a uniform family of quantum circuits of size polynomial in the security parameter.

2. *$\mathcal{A}$ can now query $\mathcal{C}$ with polynomially-many distinct m_i values, and receives either the output $y_i \leftarrow \mathsf{PRF}_\lambda(k, m_i)$ (when $b = 0$) or $y_i \xleftarrow{\$} \{0,1\}^\lambda$ (when $b = 1$).*
3. *$\mathcal{A}$ terminates and outputs a bit b'.*

We say that $\mathcal{A}$ wins the PRF *security game if $b' = b$ and define the advantage of a QPT algorithm $\mathcal{A}$ in breaking the* pseudo-random function security *of a PRF family* PRF *as:*

$$\mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) = \left|2 \cdot \Pr\left(b' = b\right) - 1\right|.$$

We define the advantage of $\mathcal{A}$ in breaking the dual-prf *security of* PRF *as:*

$$\mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda) = \max\left\{\mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda), \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{PRF}^{\mathsf{dual}},\mathcal{A}}(\lambda)\right\},$$

and say that PRF *is a post-quantum secure dual PRF family if, for all QPT algorithms $\mathcal{A}$, $\mathsf{Adv}^{\mathsf{dual\text{-}prf}}_{\mathsf{PRF},\mathcal{A}}(\lambda)$ is negligible in the security parameter λ.*

Definition 8 (Post-Quantum eufcma MAC assumption). *A message authentication code (MAC) scheme is a pair of algorithms* $\mathsf{MAC} = \{\mathsf{KeyGen}, \mathsf{Tag}\}$ *where:*

- KeyGen *is a probabilistic key generation algorithm taking input a security parameter λ and returning a random key k in the keyspace $\mathcal{K}$ of* MAC.
- Tag *is a deterministic algorithm that takes as input a secret key k and an arbitrary message m from the message space $\mathcal{M}$ and returns a MAC tag τ.*

Security is formulated via the following game that is played between a challenger $\mathcal{C}$ and a QPT adversary $\mathcal{A}$:

1. *The challenger samples $k \xleftarrow{\$} \mathcal{K}$*
2. *The adversary may adaptively query the challenger; for each query value m_i the challenger responds with $\tau_i = \mathsf{Tag}(k, m_i)$*
3. *The adversary outputs a pair of values (m^*, τ^*) such that $m^* \notin \{m_0, \ldots m_i\}$*

The adversary $\mathcal{A}$ wins the game if $\mathsf{Tag}(k, m^*) = \tau^*$, producing a MAC forgery. We define the advantage of $\mathcal{A}$ in breaking the existential unforgeability property of a MAC scheme MAC *under chosen-message attack to be:*

$$\mathsf{Adv}^{\mathsf{eufcma}}_{\mathsf{MAC},\mathcal{A}}(\lambda) = \Pr\left(\mathsf{Tag}\left(k, m^*\right) = \tau^*\right)$$

Finally, we say that MAC *is a* post-quantum eufcma *secure if, for all QPT algorithms $\mathcal{A}$, $\mathsf{Adv}^{\mathsf{eufcma}}_{\mathsf{MAC},\mathcal{A}}(\lambda)$ is negligible in the security parameter λ.*

Next we turn to defining the Decisional Diffie-Hellman problem.

Definition 9 (Decisional Diffie-Hellman (DDH) Problem). *Consider a cyclic group G of order q and with generator g. Given a tuple sampled with probability 1/2 from one of the following two distributions:*

– $\left(g^a, g^b, g^{ab}\right)$
– $\left(g^a, g^b, g^c\right)$

where a, b, c are uniformly randomly and independently sampled from $\mathbb{Z}_q$, determine which distribution the tuple is sampled from. We say that the DDH problem is hard *if an adversary has an advantage in solving the DDH problem that is negligibly greater than 1/2. We define the advantage of a PPT algorithm $\mathcal{A}$ in breaking the* ddh *problem as:*

$$\mathsf{Adv}^{\mathsf{ddh}}_{q,p,g,\mathcal{A}}(\lambda) = \left|\Pr\left(\mathcal{A}\left(\left(g^a, g^b, g^{ab}\right)\right) \to 1\right) - \Pr\left(\mathcal{A}\left(\left(g^a, g^b, g^c\right)\right) \to 1\right)\right|.$$

B Instantiating KEMs with Diffie-Hellman key exchange:

Consider a cyclic group G of order q and with generator g. We build a key encapsulation mechanism $\mathsf{KEM} = \{\mathsf{KeyGen}, \mathsf{Encaps}, \mathsf{Decaps}\}$ from a generic Diffie-Hellman key exchange in the following way:

– $\mathsf{KeyGen}(\lambda) \xrightarrow{\$} (pk, sk)$: KeyGen is a probabilistic algorithm that takes as input the security parameters $\lambda := (G, q, g)$, generates a private key $sk := x \xleftarrow{\$} \mathbb{Z}_q$, computes a public key $pk := g^x$ and returns a public/secret key pair (pk, sk).
– $\mathsf{Encaps}(pk) \xrightarrow{\$} (c, k)$: Encaps is a probabilistic algorithm that takes as input a public key $pk := g^x$, and generates a Diffie-Hellman secret key $y \xleftarrow{\$} \mathbb{Z}_q$, computes a ciphertext $c := g^y$, a key $k := (g^x)^y$ and outputs the ciphertext c as well as the key k.
– $\mathsf{Decaps}(sk, c) \to (k)$: Decaps is a deterministic algorithm that takes as input a secret key $sk := x$ and a ciphertext $c := g^y$ and outputs a key $k := (g^y)^x$, or a failure symbol $\perp$.

It is clear that if the Decisional Diffie-Hellman (DDH) problem is hard, then this KEM construction is ind-cpa secure. [12] Since we use MAC tags to authenticate our KEM encapsulations, ind-cpa security suffices for our construction.

C Security Experiment

We begin by showing the full tuple of key exchange algorithms comprising a key-exchange protocol Π, described in Section 4.2.

$\Pi.\mathsf{EQKeyGen}(\lambda) \xrightarrow{\$} (pk, sk)$ is a probabilistic post-quantum ephemeral asymmetric key generation algorithm taking as input a security parameter λ and outputting a public-key/secret-key pair (pk, sk).

[12] For the full details of ddh and ind-cpa security, refer to Appendix A

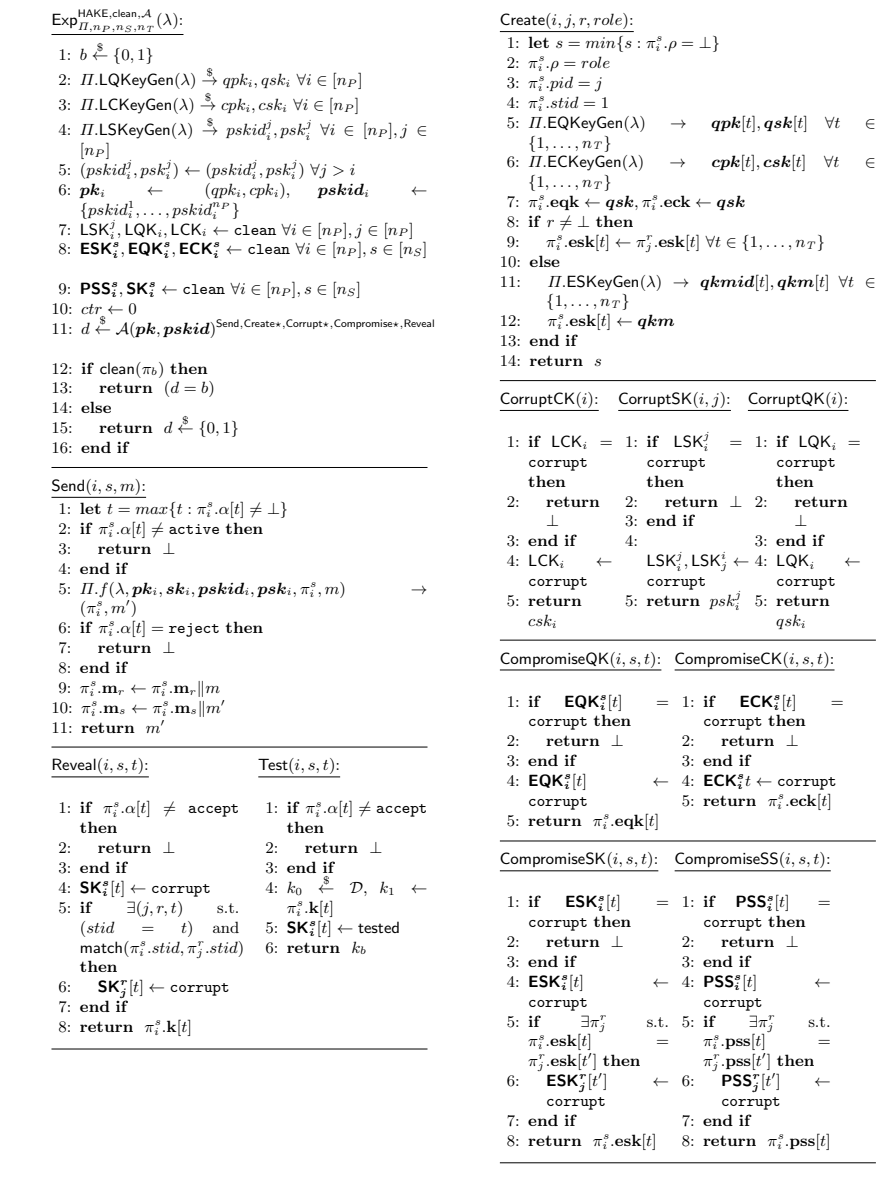

Fig. 5: HAKE experiment for adversary $\mathcal{A}$ against the key-indistinguishability security of protocol Π. Note that the values $\boldsymbol{pk}, \boldsymbol{pskid}$ given as input to $\mathcal{A}$ represent the vectors $\boldsymbol{pk}_i$, $\boldsymbol{pskid}_i$ for all n_P parties. The function match takes as input two sessions π_i^s and π_j^r and determines if they are *matching* according to some matching definition. For the definition of matching sessions used in our HAKE experiment, see Section 4.4.

$\Pi.\mathsf{ECKeyGen}(\lambda) \xrightarrow{\$} (pk, sk)$ is a probabilistic classical (i.e. not quantum-resistant) ephemeral asymmetric key generation algorithm taking as input a security parameter λ and outputting a public-key/secret-key pair (pk, sk).

$\Pi.\mathsf{ESKeyGen}(\lambda) \xrightarrow{\$} (qkm, qkmid)$ is a probabilistic ephemeral symmetric key generation algorithm taking as input a security parameter λ and outputting some symmetric keying material and (potentially) a keying material identifier $(qkm, qkmid)$.

$\Pi.\mathsf{LQKeyGen}(\lambda) \xrightarrow{\$} (pk, sk)$ is a probabilistic post-quantum long-term asymmetric key generation algorithm taking as input a security parameter λ and outputting a public-key/secret-key pair (pk, sk).

$\Pi.\mathsf{LCKeyGen}(\lambda) \xrightarrow{\$} (pk, sk)$ is a probabilistic classical (i.e. not quantum-resistant) long-term asymmetric key generation algorithm taking as input a security parameter λ and outputting a public-key/secret-key pair (pk, sk).

$\Pi.\mathsf{LSKeyGen}(\lambda) \xrightarrow{\$} (psk, pskid)$ is a probabilistic long-term symmetric key generation algorithm taking as input a security parameter λ and outputting some preshared key and (potentially) a preshared key identifier $(psk, pskid)$.

C.1 Corruption Registers

Here we detail the full set of corruption registers, which the challenger uses to determine which secrets the adversary has compromised.

Ephemeral quantum keys: $\{\mathbf{EQK_1^1}, \ldots, \mathbf{EQK_{n_S}^{n_P}}\}$, where $\mathbf{EQK_s^i}[stid] \in \{\mathtt{corrupt}, \mathtt{clean}, \perp\}$, $\forall\ i \in [n_P]$, $s \in [n_S]$ and $stid \in [n_T]$.
Ephemeral classical keys: $\{\mathbf{ECK_1^1}, \ldots, \mathbf{ECK_{n_S}^{n_P}}\}$, where $\mathbf{ECK_s^i}[stid] \in \{\mathtt{corrupt}, \mathtt{clean}, \perp\}$, $\forall\ i \in [n_P]$, $s \in [n_S]$ and $stid \in [n_T]$.
Ephemeral symmetric keys: $\{\mathbf{ESK_1^1}, \ldots, \mathbf{ESK_{n_S}^{n_P}}\}$, where $\mathbf{ESK_s^i}[stid] \in \{\mathtt{corrupt}, \mathtt{clean}, \perp\}$, $\forall\ i \in [n_P]$, $s \in [n_S]$ and $stid \in [n_T]$.
Long-term quantum keys: $\{\mathsf{LQK}_1, \ldots, \mathsf{LQK}_{n_P}\}$, where $\mathsf{LQK}_i \in \{\mathtt{corrupt}, \mathtt{clean}, \perp\},\ \forall\ i \in [n_P]$
Long-term classical keys: $\{\mathsf{LCK}_1, \ldots, \mathsf{LCK}_{n_P}\}$, where $\mathsf{LCK}_i \in \{\mathtt{corrupt}, \mathtt{clean}, \perp\},\ \forall\ i \in [n_P]$
Long-term preshared keys: $\{\mathsf{LSK}_1^1, \ldots, \mathsf{LSK}_{n_P}^{n_P}\}$, where $\mathsf{LSK}_i^j \in \{\mathtt{corrupt}, \mathtt{clean}, \perp\},\ \forall\ i, j \in [n_P]$.
Per-stage secret state: $\{\mathbf{PSS_1^1}, \ldots, \mathbf{PSS_{n_S}^{n_P}}\}$, where $\mathbf{PSS_s^i}[stid] \in \{\mathtt{corrupt}, \mathtt{clean}, \perp\}$, $\forall\ i \in [n_P]$, $s \in [n_S]$ and $stid \in [n_T]$.

Next, we give a complete pseudocode description of our HAKE security model, in Figure 5, as well as a pseudocode description of the matching sessions and origin sessions defined in Definitions 1 and 2.

D Cleanness Predicate $\mathsf{clean}_{c\mathsf{HAKE}}$

Definition 10 ($\mathsf{clean}_{c\mathsf{HAKE}}$). *A session π_i^s in stage t such that $\pi_i^s.\alpha[t] = \mathtt{accept}$ and $\pi_i^s.pid = j$ in the security experiment defined in Figure 5 is $\mathsf{clean}_{c\mathsf{HAKE}}$ if all of the following conditions hold:*

1. *The query $\mathsf{Reveal}(i, s, t)$ has not been issued.*

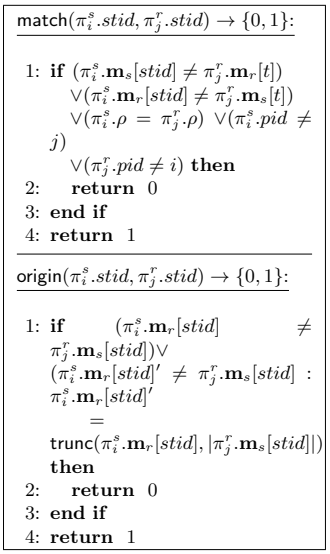

Fig. 6: A pseudocode description of the *matching session* and *origin session* functions defined in Definitions 1 and 2.

2. *For all* $(j,r,t) \in n_P \times n_S \times n_T$ *such that* π_i^s matches π_j^r *in stage* t*, the query* $\mathsf{Reveal}(j,r,t)$ *has not been issued.*
3. *If there exists a session* π_j^r *such that* π_j^r *matches* π_i^s *in stage* t*, then at least one of the following sets of queries has not been issued:*
   - $\mathsf{CompromiseQK}(i,s,t)$, $\mathsf{CompromiseQK}(j,r,t)$ *have not been issued, where* π_j^r matches π_i^s *in stage* t*.*
   - $\mathsf{CompromiseSK}(i,s,t)$, $\mathsf{CompromiseSK}(j,r,t)$ *have not been issued, where* π_j^r matches π_i^s *in stage* t*.*
   - $\mathsf{CompromiseCK}(i,s,t)$, $\mathsf{CompromiseCK}(j,r,t)$ *have not been issued, where* π_j^r matches π_i^s *in stage* t*.*
   - $\mathsf{CompromiseQK}(i,s,t')$, $\mathsf{CompromiseQK}(j,r,t')$ *have not been issued, where* π_j^r matches π_i^s *in stages* u *such that* $t' \leq u < t$ *and no* $\mathsf{CompromiseSS}(i,s,u)$, $\mathsf{CompromiseSS}(j,r,u)$ *queries have been issued.*[13]
   - $\mathsf{CompromiseSK}(i,s,t')$, $\mathsf{CompromiseSK}(j,r,t')$ *have not been issued, where* π_j^r matches π_i^s *in stages* u *such that* $t' \leq u < t$ *and no* $\mathsf{CompromiseSS}(i,s,u)$, $\mathsf{CompromiseSS}(j,r,u)$ *queries have been issued.* [14]
   - $\mathsf{CompromiseCK}(i,s,t')$, $\mathsf{CompromiseCK}(j,r,t')$ *have not been issued, where* π_j^r matches π_i^s *in stages* u *such that* $t' \leq u < t$ *and no* $\mathsf{CompromiseSS}(i,s,u)$, $\mathsf{CompromiseSS}(j,r,u)$ *queries have been issued.* [15]
4. *If there exists no* $(j,r,t) \in n_P \times n_S \times n_T$ *such that* π_j^r *is an origin session of* π_i^s *in stage* t*, then* $\mathsf{CorruptSK}(i,j)$ *and* $\mathsf{CorruptSK}(j,i)$ *have not been issued*

[13] Refer to footnote 9.
[14] Refer to footnote 9.
[15] Refer to footnote 9.

before $\pi_i^s.\alpha[t] \leftarrow$ `accept`. *If there exists a* $(j, r, t) \in n_P \times n_S \times n_T$ *such that* π_j^r *is an origin session of* π_i^s *in stage* t*, then* $\mathsf{CorruptSK}(i, j)$ *and* $\mathsf{CorruptSK}(j, i)$ *have not been issued before* $\pi_j^r.\alpha[t] \leftarrow$ `accept`.

E Muckle Message Structure

Here we give the structure of a Muckle key exchange message. Note that the structure of an initiator message is identical to the structure of a message from a responder.

```
typedef struct muckle_msg_header {
u_int8_t type;
u_int8_t version;
unsigned char partyIdentifier[32];
} MUCKLE_MSG_HEADER;

typedef struct muckle_msg {
MUCKLE_MSG_HEADER header;
unsigned char qraSidhPub[378];
unsigned char tag[32];
} MUCKLE_MSG;
```

We briefly describe each field in the message.

header: Consists of 3 sub-fields:
- **type:** is a flag indicating whether the message came from initiator or responder.
- **version:** is a flag indicating what key exchange primitives are used in Muckle. These can be seen as analogous to ciphersuite indicators in a TLS handshake.
- **partyIdentifier:** is a public string indicating the identity of the party sending the message. This could be seen equally as a preshared key identifier, and assume that some unique mapping exists in the implementation between party identities and preshared keys.

classEcdhPub: is the public keyshare of the elliptic-curve-based key exchange primitive used in Muckle.

qraSidhPub: is the public keyshare of the supersingular isogeny-based key exchange primitive used in Muckle.

tag: is the output message authentication code, computed over the rest of the Muckle message.