A Logic and an Interactive Prover for the Computational Post-Quantum Security of Protocols

Cas Cremers*, Caroline Fontaine†, Charlie Jacomme*

*CISPA Helmholtz Center for Information Security, Germany

† Université Paris-Saclay, CNRS, ENS Paris-Saclay, Laboratoire Méthodes Formelles, 91190, Gif-sur-Yvette, France

Abstract—We provide the first mechanized post-quantum sound security protocol proofs. We achieve this by developing PQ-BC, a computational first-order logic that is sound with respect to quantum attackers, and the corresponding mechanization support in the form of the PQ-SQUIRREL prover.

Our work builds on the classical BC logic [7] and its mechanization in the SQUIRREL [5] prover. Our development of PQ-BC requires making the BC logic sound for a single interactive quantum attacker. We implement the PQ-SQUIRREL prover by modifying SQUIRREL, relying on the soundness results of PQ-BC and enforcing a set of syntactic conditions; additionally, we provide new tactics for the logic that extend the tool's scope.

Using PQ-SQUIRREL, we perform several case studies, thereby giving the first mechanical proofs of their computational post-quantum security. These include two generic constructions of KEM based key exchange, two sub-protocols from IKEv1 and IKEv2, and a proposed post-quantum variant of Signal's X3DH protocol. Additionally, we use PQ-SQUIRREL to prove that several classical SQUIRREL case studies are already post-quantum sound.

Index Terms—Security Protocols, Post Quantum, Formal Methods, Observational Equivalence, Computational Security, Interactive Prover.

I. INTRODUCTION

In recent years, multiple highly-successful tools have been developed to analyze and verify cryptographic protocols and primitives [9], [15], [16], [48], [57]. They have proven the usefulness and necessity of computer-aided cryptography, both uncovering critical attacks against widely deployed protocols and helping in the design of new standards [3], [10], [13], [22], [25]–[27], [29], [30], [43], [44], [47].

In anticipation of developments in quantum computing that would break a lot of widely-used cryptographic primitives, the security community has started to develop many new security primitives and protocols, and revisit old protocols. Additionally, an extensive multi-year NIST standardization process is ongoing to develop new primitives and protocols. These efforts aim to establish mechanisms that are provably secure against quantum attackers. At some level of abstraction, this implies (i) designing new primitives and prove (or assume) that they are secure against a quantum attacker, and (ii) proving that a concrete protocol that uses such primitives is indeed secure against a quantum attacker. In this work, we focus on the latter, and in particular how we can mechanize such proofs.

A classical strategy for proving a protocol's security is a so-called *reduction* proof, which yields computational security guarantees against a polynomial-time attacker. This approach is used in most pen-and-paper proofs by cryptographers, and involves constructing a reduction from any attack on the protocol to an attack on the used cryptographic assumptions, and then reasoning by contradiction. This is a well studied approach with respect to classical (non-quantum) attackers: different flavors of such proofs can be mechanized by tools such as CRYPTOVERIF [15] and EASYCRYPT [9]. However, some proof steps commonly used in reductions that are valid for a classical attacker, such as *rewinding*, are in general not valid anymore for quantum attackers. This result is similar to the *no-cloning* theorem [61], which implies that one must be careful when talking about the state of a quantum attacker. As a consequence, a classical reduction proof of a protocol (even based on post-quantum sound primitives) may not be valid for quantum attackers. Unfortunately, there exists no formal framework nor mechanization dedicated to computational proofs of a protocol's security versus a quantum attacker.

In this work, we address this problem by developing PQ-BC, a post-quantum sound variant of a computationally sound protocol logic, and a corresponding tool called the PQ-SQUIRREL prover, by extending the logic's tool support for the post-quantum setting, as well as adding tactics. We use our new tool to provide the first mechanized post-quantum computational security proofs for several protocols.

Concretely, our work builds on the BC logic [7] and its mechanization in the SQUIRREL prover [5]. The BC logic can be used to construct security proofs that provide computational guarantees against a classical (non-quantum) attacker, while only working inside a logical framework in which many intricate details have been abstracted. It has notably been used for manual proofs of real-world protocols, see e.g., proofs of RFID based protocols [24], AKA [45], e-voting protocols [6], key-wrapping API [53], and SSH through a composition framework [23]. Reasoning in BC was recently mechanized and extended in the SQUIRREL prover [5], dedicated to the formal proofs of protocols. Notably, reasoning in BC (and therefore SQUIRREL) is not sound with respect to a quantum attacker, because the framework allows reduction steps that cannot be reproduced with a quantum attacker.

To develop PQ-BC, we have to make the BC logic sound for a single interactive quantum attacker, while it previously relied on a set of deterministic one-shot attackers. This seemingly small change triggers a cascade of technical changes. We design a new term interpretation for the logic and identify three syntactic conditions for proofs that help ensure their post-quantum soundness. We provide mechanization for PQ-BC

in the form of the PQ-SQUIRREL prover. PQ-SQUIRREL's soundness relies on the soundness results of PQ-BC and implements the syntactic conditions; additionally, we design and implement new tactics that extend the tool's scope.

Contributions. We see our main contributions as the following:

- First, we develop PQ-BC, the first computational first-order logic to prove guarantees of security protocols whose results are provably sound with respect to a quantum attacker.
- Second, we develop the PQ-SQUIRREL prover, a mechanized tool support for establishing such guarantees.
- Third, we use our tool to provide the first mechanized proofs of the post-quantum computational security of 11 security protocols as case studies. These include two KEM-based key exchanges [18], [35], a post-quantum variant of Signal's X3DH [38], and two protocols from the IKE standards [20], [41] – confirming claims in [34].

Overview: We provide in Section II the necessary background on the BC logic and the SQUIRREL prover. Then, in Section III, we give a high-level overview of the design of the PQ-BC logic and how it differs from the BC logic. In Section IV we formally define PQ-BC, its syntax and semantics, and its rules; in Section V we describe PQ-SQUIRREL and perform case studies. We discuss current limitations and future work in Section VI, and conclude in Section VII.

Upon first reading, the reader may get a high-level understanding of the paper by skipping Section IV and directly continuing with Section V.

We provide all source code, protocol models, and the long version of this paper with full details, at [1].

Additional related Work: Issues regarding the validity of classical cryptographic reductions in the post-quantum setting have mostly started with [60], which identified the "no-cloning theorem" [61] as a key issue, followed, e.g., by [4].

Key details and difficulties when moving to post-quantum security are generally discussed [36], [56]. They provide some guidelines that are suitable for game-based approaches and gave us many insights, but those guidelines are not suited for the BC logic approach.

There exists many tools for security proofs, we only discuss the most widely used. At one extreme of the spectrum are tools like EASYCRYPT [9] and CRYPTOVERIF [15], which provide strong computational guarantees for detailed models of cryptographic primitives, but for whom scaling to larger constructs is more challenging; at the other end of the spectrum, we have tools like TAMARIN [48] and PROVERIF [16], which can analyze much larger protocol mechanisms by using a more abstract symbolic model, but cannot provide computational guarantees. SQUIRREL, and thus PQ-SQUIRREL, lies in the middle ground between those two ends: on one side, it provides computational guarantees, that are thus stronger than the one given by PROVERIF and TAMARIN; on the other side, it operates at a higher level of abstraction than EASYCRYPT and CRYPTOVERIF. Consequently, SQUIRREL is less expressive and thus less suited to reason about cryptographic primitives, but

tends to scale better to larger construct. However, SQUIRREL does not provide any concrete security bounds, and in security proofs over unbounded protocols, the number of sessions is arbitrary and not attacker chosen. For a detailed comparison between SQUIRREL, EASYCRYPT and CRYPTOVERIF, we refer the reader to [5, Appendix E].

CRYPTOVERIF does not have any support for quantum attackers yet, it might be possible to make it quantum-sound by using ideas from our work, such as forbidding some manipulations over the attacker state, and ensuring that a unique quantum attacker process can continue without having to alter or inspect its internal state.

The previously mentioned EASYCRYPT is a toolset for constructing cryptographic proofs, which currently mainly targets cryptographic primitives. It was first adapted to the quantum setting with qRHL [59], a formal security prover based on a quantum relational Hoare logic, and later (in concurrent work to ours) to the post-quantum setting with EASYPQC [8]. The qRHL approach works on quantum constructions, which substantially complicates proving classical constructions. For example, there is no equivalent to the classical implication operation over quantum predicates (see e.g., [28]). EASYPQC avoids this overhead by only considering classical constructions.

Similar to our approach, EASYPQC adds new side conditions to its core logic, such as forbidding case distinctions on the attacker's internal state. It is difficult to compare their side conditions to ours, since the conditions are deeply linked to the underlying logics, which are of a very different nature. Notably, EASYPQC supports reasoning in the Qantum Random Oracle Model (QROM). The BC logic does not yet support the ROM (nor QROM), and hence neither do we. This is not an inherent restriction of the logic and could be future work. For our current case studies, we prefer the use of the PRF assumption over the QROM.

EASYPQC and our approach inherit their focus from their starting points: the EASYCRYPT approach is more geared towards cryptographic primitives, while BC is designed for protocols. All current EASYPQC case studies are cryptographic primitives, whereas our case studies are protocols. In particular, our case studies for KEM based key exchanges are the first mechanized proofs with computational guarantees of such protocols.

II. BACKGROUND: THE CLASSICAL BC LOGIC AND SQUIRREL

Below we first recall the main elements of the original BC logic [7] that are relevant for understanding our work in the following sections. In Section II-E we describe the SQUIRREL prover [5], which mechanizes reasoning in the BC logic.

In the computational model, the security of a protocol is established by showing that the protocol cannot be distinguished from its idealized version by any polynomial-time attacker w.r.t. a security parameter. Such security proofs of protocols rely on two ingredients: a computational hardness assumption, and a security reduction showing that an attacker that can break the security of the protocol can break the hardness assumption.

But the construction of security reductions is difficult and error-prone. To ease this process, the BC logic proves the security of protocols inside a first-order logic. This approach requires that within BC, everything is modeled using only terms, i.e., purely syntactic constructs. This is very different from the game-based modeling, where protocols are expressed as abstract programs with procedure calls, states, and side effects. With terms, all protocol actions become pure functional calls, which tends to ease the formal reasoning. This leads to some core elements in the design of the logic: one needs to

1) define terms, as well as an interpretation from protocols to terms so that the terms syntactically describe all the behaviours of the protocol,
2) define logical predicates and rules (which include axioms) to reason about our terms, and
3) show that the rules are sound, i.e., that the rule applications correspond to correct reductions.

We provide an overview of the first two elements in the following, and refer the reader to [7] for details of the rules and their soundness. However, all three elements will be discussed when we present the modified post-quantum sound logic in the following sections.

A. Specifying protocol behaviours using syntactic terms

1) From protocols to terms: Let us consider a very simple example protocol process P, using an informal syntax.

Example 1 (Protocol).

$$P := \textbf{new sk}.\textbf{in}(x).\textbf{new r}.\textbf{out}(enc(x, \mathsf{r}, \mathsf{sk})).$$
$$\textbf{in}(y).\textbf{new r}'.\textbf{out}(enc(y, \mathsf{r}', \mathsf{sk})).$$

Process P samples a secret key sk*, and uses it to encrypt some attacker input x using the random seed* r *(explicitly modeling probabilistic encryption). It then encrypts a second input y with random seed* r'*.*

Equivalently, in the game-based notation with a stateful probabilistic attacker $\mathcal{A}$ and security parameter η, the experiment that returns the attacker-observable values is defined as:

$$
\begin{array}{l}
\underline{\textit{Experiment } Exp^P_{\mathsf{enc}, \mathcal{A}}(\eta)} \\[4pt]
\mathsf{sk} \xleftarrow{\$} \{0,1\}^\eta \\
x \xleftarrow{\$} \mathcal{A}(1^\eta) \\
\mathsf{r} \xleftarrow{\$} \{0,1\}^\eta \\
y \xleftarrow{\$} \mathcal{A}(1^\eta, \mathsf{enc}(x, \mathsf{r}, \mathsf{sk})) \\
\mathsf{r}' \xleftarrow{\$} \{0,1\}^\eta \\
\textbf{return } (\mathsf{enc}(x, \mathsf{r}, \mathsf{sk}), \mathsf{enc}(y, \mathsf{r}', \mathsf{sk}))
\end{array}
$$

To syntactically represent such observable sequences of values using terms, one can use the following constructions:

- fresh values n sampled from an infinite set $\mathcal{N}$, representing randomly sampled bitstrings, such as r, r', and sk above;
- public function symbols $f \in \Sigma$, to model e.g., encryption functions such as enc; and
- variables such as x and y from the set of variables $\mathcal{X}$, modeling attacker inputs.

Any protocol computation can then be modeled as applications of public functions to either fresh values modeling randomly generated values (such as nonces and secret keys), or variables modeling attacker inputs. This is essentially the Dolev-Yao model [31], where the protocol can be described with the following term sequence:

$$\mathsf{enc}(x, \mathsf{r}, \mathsf{sk}), \ \mathsf{enc}(y, \mathsf{r}', \mathsf{sk}) \tag{1}$$

This sequence of terms, that we will refer to as the *frame* of the protocol, represents the possible messages that an attacker can observe during the protocol's execution.

2) Modeling attacker computations: The sequence of terms in Eq. (1) is not yet sufficient to reason syntactically about protocols, because it does not capture that y probabilistically depends on the value of $\mathsf{enc}(x, \mathsf{r}, \mathsf{sk})$. From a high-level point of view the logic must satisfy **locality**: a term must explicitly contain all its probabilistic dependencies. Essentially, it needs to syntactically capture that y is the result of an attacker's unknown computation, which depends on the previous messages, e.g., in our example, y depends on r and sk. The BC logic uses free function symbols $\mathtt{att}_i$ that represent unknown pieces of code, i.e., attacker computations that receive as arguments the previous messages seen by the attacker.

The previous frame can now be expressed as:

$$\mathsf{enc}(\mathtt{att}_0(), \mathsf{r}, \mathsf{sk}), \ \mathsf{enc}(\mathtt{att}_1(\mathsf{enc}(\mathtt{att}_0(), \mathsf{r}, \mathsf{sk})), \mathsf{r}', \mathsf{sk}) \tag{2}$$

$\mathtt{att}_0()$ representing the first message (x) computed by the attacker, when it does not have access to any information from the protocol, and $\mathtt{att}_1(\mathsf{enc}(\mathtt{att}_0(), \mathsf{r}, \mathsf{sk}))$ being its second message (y), which is a function of the protocol's first output.

The term-based notation is more akin to a functional view: one cannot use variables such as x to refer to previously computed values. For this reason, $\mathtt{att}_0()$ occurs twice in Eq. (2). However, attacker computations such as $\mathtt{att}_0$ are probabilistic algorithms, and hence two different invocations might yield different results. This is not the intended interpretation, and it therefore introduces a new requirement on terms: **stability** - two occurrences of the same term must evaluate to the same value. When one evaluates the value of the previous frame, this implies that the two occurrences of $\mathtt{att}_0()$ evaluate to the same value. This requirement captures that within a single protocol execution, identical terms refer to the same value and not to separate (probabilistic) attacker calls.

3) Reasoning about terms: The BC logic contains in its syntax a binary predicate $\sim$. This predicate expects two sequences of terms and intuitively represents their indistinguishability. Note that $\sim$ has low operator precedence, but we often add parentheses for readability.

Example 2 (BC indistinguishability formula).

$$\left(\mathsf{n}', \textbf{if } \mathtt{att}_0(\mathsf{n}') \doteq \mathsf{n} \textbf{ then } \mathsf{ko} \textbf{ else } \mathsf{ok}\right) \sim \left(\mathsf{n}', \mathsf{ok}\right) \tag{3}$$

The protocol modeled on the left-hand side produces a fresh value n', sends it, and then waits to receive another value (represented by $\mathtt{att}_0(\mathsf{n}')$). If the received value is the same as a freshly generated value n, the protocol outputs

ko, otherwise ok. The formula expresses that the attacker cannot distinguish this protocol from the protocol that sends a fresh value n', waits for an input, and then outputs ok. Note that the **if _ then _ else _** notation is syntactic sugar for a ternary function symbol $\mathtt{ite}(\_,\_,\_)$, with which one models conditionals in BC. To give a flavor of how proofs are performed in the BC logic, we provide some rule examples and prove that Eq. (3) holds in the logic.

Example 3 (Rules and proofs)**.**

$$\frac{}{(t \doteq \mathsf{n}) \sim \textit{false}} \;{=}\text{IND} \quad \textit{when } \mathsf{n} \textit{ does not occur in } t$$

$$\frac{\phi \sim \textit{false} \quad (u,v) \sim w}{(u, \textbf{if } \phi \textbf{ then } s \textbf{ else } v) \sim w} \;\text{I\textsc{f-f}} \qquad \frac{}{u \sim u} \;\text{R\textsc{efl}}$$

Logical rules are read bottom-up, where to prove the formula on the bottom, one can prove the formulas on the top. Rule ${=}\text{IND}$ *means that a term t that does not syntactically contain a fresh value n cannot be equal to it, and* I\textsc{f-f} *tells us that if a conditional is always false, we can only consider its negative branch. The* R\textsc{efl} *rule encodes that $\sim$ is reflexive.*

With those rules, we can prove the simple property in Eq. (3):

$$\frac{\dfrac{}{(att_0(\mathsf{n}') \doteq \mathsf{n}) \sim \textit{false}}\;{=}\text{IND} \quad \dfrac{}{(\mathsf{n}', \mathsf{ok}) \sim (\mathsf{n}', \mathsf{ok})}\;\text{R\textsc{efl}}}{(\mathsf{n}', \textbf{if } att_0(\mathsf{n}') \doteq \mathsf{n} \textbf{ then } \mathsf{ko} \textbf{ else } \mathsf{ok}) \sim (\mathsf{n}', \mathsf{ok})}\;\text{I\textsc{f-f}}$$

B. A faithful computational interpretation

We previously described the BC way to syntactically describe the behaviour of a protocol interacting with an attacker. To ultimately get to a logic that provides computational guarantees, those syntactic terms must capture all the behaviours of the protocol. To do so, BC provides a formal way to interpret those terms, so that their possible evaluations match those of the real protocol. Intuitively, given an attacker against the real world protocol, and given the syntactic frame (e.g., Eq. (2)), it should be possible to build a simulator producing the same results as the protocol. If this is possible, then the frame does indeed capture all the possible behaviours of the protocol, and we can thus use it to reason about its security.

1) Interpreting terms: We now describe how BC interprets a term, i.e., computes the probabilistic result of a term, while satisfying both locality and stability. The interpretation has three parameters: the security parameter η and two infinite random bitstrings ρ_s and ρ_r. The interpretation extracts the randomness used in probabilistic protocol functions from ρ_s, and the attacker randomness from ρ_r. By universally quantifying over those parameters and the attacker computations, it captures all the possible executions of a protocol.

The interpretation depends on a set of Polynomial Time Turing Machines (PTTM) that compute the evaluation function $\mathcal{T}$ of a term, i.e.,

- a machine $\mathcal{T}_\mathsf{n}$ for n: outputs the value of a given fresh value, depending on η and ρ_s, which may typically extract a sequence of η bits from ρ_s;


Fig. 1. $\mathtt{enc}(att_0(), r, sk)$

Fig. 2. $att_0() \doteq att_0()$

- a machine $\mathcal{T}_f$ for f: computes the output of a public function depending on its arguments;
- a machine $\mathcal{T}_{att_i}$ for att_i: performs some attacker computation, depending on its arguments, η and ρ_r.

To interpret a term, one can consider the term as a tree of (sub)terms, and recursively call the corresponding Turing machine from the leaves to the root. Going back to the term $\mathtt{enc}(att_0(), r, sk)$ occurring in Eq. (2), its possible values are obtained by running the simulation described in Fig. 1.

By design, all Turing machines in the interpretation are deterministic: the randomness is explicitly passed to capture probabilistic behaviour (using ρ_s and ρ_r). This ensures that the interpretation satisfies stability: two occurrences of $att_0()$ in the term imply two computations of $\mathcal{T}_{att_0}$ that deterministically evaluate to the same value. For example, if we consider the function symbol $\doteq$ that models equality testing, the evaluation of Fig. 2 always returns true, which is the expected behaviour.

Note that $\mathcal{T}_{att_0}$ and $\mathcal{T}_{att_1}$ do not share any implicit state and, more generally, nor would $\mathcal{T}_{att_i}$ and $\mathcal{T}_{att_{i+1}}$. This may seem counter-intuitive, as att_i and att_{i+1} usually represent two unknown computations of the same attacker $\mathcal{A}$, that can in practice maintain a state between successive calls. The BC logic does not explicitly ensure this, but by modeling conventions, we always give all the previous inputs of att_i also to att_{i+1}. As a result, in the simulations one can make $\mathcal{T}_{att_{i+1}}$ re-perform all the computations of $\mathcal{T}_{att_i}$ to recompute its state. Modelling each attacker call by a distinct machine is a crucial design choice of the BC logic that ensures locality.

2) Protocol interactions: We now focus on protocol interactions, as it will play an important role when moving to the quantum setting. We informally describe an attacker $\mathcal{A}$ in the real world as a sequence of unknown operations $\{\mathcal{A}^i\}$, that produces the message m_i at step i, and implicitly maintains a state ϕ_i between each computation. We can consider a protocol P as a sequence of operations split by **in** (receive) operations. The attacker's operations interact with the protocol sequence: a protocol operation P^i typically uses **out** to send messages, providing more knowledge to the attacker, which can then use this knowledge to produce a message for P^{i+1}. E.g., considering Example 1, we define $\mathcal{T}_{att_0}$ s.t. it simulates $\mathcal{A}^0$ and returns m_0, and $\mathcal{T}_{att_1}$ s.t. on input u it first simulates $\mathcal{A}^0$ to get ϕ_0, and then runs $\mathcal{A}^1$ on ϕ_0 and u.

Still considering P from Example 1, with the corresponding syntactic frame Eq. (2), we depict in Fig. 3 the simulation that produces the same result as the real-world interaction.

C. Indistinguishability predicate and logical rules

Once equipped with an interpretation for terms, BC defines the predicate $\sim$ that will return true over two sequence of terms if the advantage of any final distinguisher is negligible in η. As

4

Fig. 3. Final BC simulation of attacker/protocol interactions. Note the recomputation of, e.g., $\mathcal{T}_{\text{att}_0}$.

for any other attacker computation step, this final distinguisher does not inherit any states from the previous attacker calls, but it may recompute those states if it is given the same arguments.

Consider two protocols: one can build as previously described two sequences of terms that correctly capture all the behaviours of their interactions with an attacker. If there exists an attacker that can distinguish the two protocols, then there is a set of attacker Turing machines $\{\mathcal{T}_{\text{att}_i}\}$ that shows that the simulation of the two sequences of terms can be distinguished, and thus that $\sim$ does not hold over them. This constitutes the soundness of the BC logic.

Finally, one needs to prove the soundness of the logical reasoning rules by using reductions. For instance, a simple reduction would prove that a rule such as IF-F from Example 3 is indeed a valid rule to reason about $\sim$.

The rules presented in this example can be proven sound for any Turing machine attacker, without limiting its computational power. Such rules form the core of the logic, and are referred to as structural rules. In addition, one also needs to integrate the classical cryptographic assumptions inside the logic.

D. Cryptographic assumptions

In BC, cryptographic assumptions are encoded as rules (also referred to as axioms). BC includes provenly-sound rules to encode assumptions such as PRF, IND-CCA, EUF-CMA, ENC-KP, INT-CTXT, OTP, and DDH. Given a cryptographic assumption, one can translate it into a BC rule using terms: proving soundness of such a rule is done by providing a reduction showing that if there exists a model that breaks the BC assumption, then there exists an attacker that breaks the original cryptographic assumption. As the translation is often natural, the soundness proofs of the axioms are relatively straightforward, and are usually fully black-box reductions. We refer the reader to [7], [46] for the soundness proofs of the rules corresponding to the main cryptographic assumptions.

Example 4 (DDH assumption). *In* BC *proofs, the full* DDH *assumption is not usually made: the encoded as-*

sumption can be seen as a generic version that is agnostic about the implementation. Concretely, one can see the DDH *assumption of* BC *as a generic assumption over two function symbols, an extractor* ext *and a combiner* comb, *such that, for any names* n_1, n_2, n_3 *we have*
$$\big(\text{ext}(n_1), \text{ext}(n_2), \text{comb}(\text{ext}(n_1), n_2), \text{comb}(\text{ext}(n_2), n_1)\big)$$
$$\sim \big(\text{ext}(n_1), \text{ext}(n_2), \quad \text{ext}(n_3), \quad \text{ext}(n_3)\big).$$
For a cyclic group with generator g, *the extractor function would be* $\text{ext}(n) := g^n$ *and the combiner* $\text{comb}(t, u) := t^u$, *and the* DDH *over the group assumption implies our generic assumption over the functions.*

Example 5 (OTP assumption). *For the exclusive-or operator* $\oplus$, *we only assume that it is a binary function symbol such that for any term* t *and fresh name* n, *it allows for a one-time pad, i.e.,* $\big((t \oplus n) \sim n\big)$.

For any security proof of a protocol that uses $\oplus$, *one typically does not need to specify that* $\oplus$ *may be associative and commutative. Such a proof then gives us security guarantees, regardless of whether the concrete implementation of* $\oplus$ *is associative or not.*

To summarize, for proving with BC the full (computational) security of a protocol that uses concretely instantiated primitives, we first require a proof that the primitives instantiate cryptographic assumptions. Then, we use a translation of those assumptions in the BC proof to prove properties of the protocol. For example, to give a computational proof of a signed Diffie-Hellman protocol that uses Ed25519 signatures, we would use the proof that Ed25519 signatures satisfy EUF-CMA, and then use the corresponding EUF-CMA rule in the BC proof.

Note that it is also not uncommon to leave this question unanswered: i.e., to propose a protocol that relies on an assumption (e.g., PRF), for which we do not know an instantiation that provably meets it, and the guarantee is then phrased as "assuming that primitive X is a PRF, . . .". Such a proof can still be meaningful since it proves the absence of a class of attacks, or anticipates a future provable instantiation.

E. The SQUIRREL Prover

The SQUIRREL Prover mechanizes the BC logic, and offers additional features not present in BC. Two downsides of the original BC logic, which are solved by SQUIRREL, are:

1) to prove the security of a protocol, one must make a proof for each possible action orderings of the protocol, and
2) the logic only considers finite protocols, and therefore notably only a bounded number of sessions.

As the security proof of a protocol may be very similar for many action orderings, which are also called (symbolic) traces, the first point implies a lot of tedious repetitions inside proofs. The second point can be solved by performing induction over the number of sessions, but this cannot be done inside the BC logic, and instead requires external mathematical reasoning.

1) Protocol Specification: The SQUIRREL prover's protocol specification language is close to the well-known applied pi-calculus. In Fig. 4 we give a simplified version of a key exchange protocol based on asymmetric encryption. An Initiator

Fig. 4. Encryption based key exchange

with secret and public keys (sk_I, pk_I) sends some encrypted ephemeral secret k_I to a responder with keys (sk_R, pk_R), that answers with its own encrypted ephemeral share k_R. Then, a shared key is derived as $h(s, k_I) \oplus h(s, k_R)$, where h is a PRF hash function and s a public seed. The derived key should only be computable by the intended peer. The protocol does not offer any explicit authentication. A possible implicit authentication guarantee could be: If a party X derives a key k with intended peer Y, then whoever can also compute the key must be Y, and agree on the communication partners.

A basic model of this protocol is expressed in SQUIRREL as illustrated in Listing 1. The **hash** and **aenc** instructions declare function symbols that respectively satisfy either the PRF assumption or the IND-CCA assumption, and the **name** command declares fresh values. In the process declarations, **out** and **in** specify outputs and inputs, **new** specifies sampling fresh values, and | is for parallel composition. In the last step of the Initiator, the instruction sIR:= is used to store the derived key in a state. The last line of this example captures a security property, which we discuss in next Section.

Note that the threat model is part of the input file, and SQUIRREL can be used to verify many scenarios. For instance, more complex versions of this running example that will be discussed in Section V-C can include an unbounded number of sessions and participants, and $pk(sk_I)$ may not be hard-coded inside the Responder process but received as an input.

```
hash H
aenc enc,dec,pk
name s,skI,skR,kIR : message

process Initiator =
 new kI; new rI;
 out(cI,  enc(kI, rI,pk(skR)) );
 in(cR,ctR);
 let kR = dec(ctR,skI) in
 sIR :=  H(s,kI) XOR H(s,kR).

process Responder =
   new kR;
   new rR;
   in(cI, ctI);
   out(cR, enc(kR, rR, pk(skI))).

system [KE] out(cI,s); (Responder | Initiator).

equiv [KE] forall t, frame@t,sIR ~ frame@t,kIR.
```

Listing 1. Simplified key exchange in SQUIRREL

2) The SQUIRREL *meta-logic:* The SQUIRREL paper [5] introduces a meta-logic that introduces universal quantification over the possible traces of a protocol, as well as its number of sessions. As one can then reason at an abstract level on the traces, a single proof in the meta-logic covers multiple traces of the protocol, and it is now possible to consider unbounded protocols, thus solving both previous issues.

To do so, the logic is extended with timestamp variable τ to quantify over possible points in a trace, and macros that depend on a protocol P, such as:

- $\mathsf{input}^P@\tau$, to refer to the input received by the protocol P at timestamp τ;
- $\mathsf{output}^P@\tau$ for the output;
- $\mathsf{frame}^P@\tau$ for the full frame for the protocol up to this point, which is the sequence of all previous outputs.

If we denote by $\mathsf{pre}(\tau)$ the function that points to the previous timestamp in the trace, and abstracting away some details, $\mathsf{frame}^P@\tau$ would expand to $\mathsf{frame}^P@\mathsf{pre}(\tau), \mathsf{output}^P@\tau$.

The indistinguishability of two protocols P and Q can be expressed in a single meta-logic formula $\mathsf{frame}^P@\tau \sim \mathsf{frame}^Q@\tau$. This meta-logic formula then holds if for all possible traces of the protocol, and all possible τ over those traces, the BC frames of the protocol are indistinguishable. In the example of the *KE* protocol of Listing 1, the last line declares the security goal expressing the fact that the key derived by the initiator is indistinguishable from some fresh value k_{IR}, even given all the messages sent over the network. The proof in SQUIRREL of this goal will be described at a high-level in Section V-C3.

We remark however that even though this implies a proof of security for each possible number of sessions, it does not imply the classical security for an unbounded number of sessions because no concrete security bounds are obtained. Over this meta-logic, there are two sets of rules. The first allows proving that some probabilistic statement is true with overwhelming probability, and the second is for proving indistinguishability properties. These two sets interact: for instance, one can first prove authentication properties of a key exchange using rules from the first set, and then rely on these properties to prove the secrecy of the key using the second set.

3) SQUIRREL*'s mechanization:* The SQUIRREL prover is available at [2], along with the code of the running example and the case studies. It is an interactive prover with some built-in tactics that simplify low-level reasoning.

From a high-level point of view, most proofs inside SQUIRREL follow the following schema:

1) prove a set of probabilistic statements inside the dedicated prover, such as matching conversations, or that some bad states cannot be reached;
2) prove the desired indistinguishability, using the probabilistic statements as lemmas.

The second step often relies on an induction over the length of the trace. Concretely, the proof is usually done by assuming that $\mathsf{frame}^P@\mathsf{pre}(\tau) \sim \mathsf{frame}^Q@\mathsf{pre}(\tau)$ is true, and proving that

frame^P@$\tau \sim$ frame^Q@τ holds. By definition of the frame, this is equal to proving that

frame^P@pre(τ), output^P@$\tau \sim$ frame^Q@pre(τ), output^Q@τ

and making a case distinction over possible values of output.

As a concrete example: Fig. 4 in SQUIRREL takes 20 lines to model, 20 lines of proof, and is then proved by SQUIRREL in under a second.

III. ADAPTING THE BC LOGIC AND SQUIRREL TO THE POST-QUANTUM WORLD

Proofs in the BC logic as presented in the previous section do not guarantee security against quantum attackers.

Recall that computational proofs rely on (i) computational assumptions, and (ii) reductions. In the context of the BC logic, the computational assumptions are the easiest to deal with when going to the post-quantum world: if we know that no post-quantum secure instantiations exist for a particular assumption, we can no longer use it in proofs. For example, if a previous proof relied on an axiom that states that integer factorization is hard, the proof is unsound as we can no longer use this axiom in the post-quantum setting.

The complexity of adapting the reductions that follow from the logic mostly revolves around revisiting the assumptions around the term interpretation and the proof rules, which implicitly encode attacker manipulations. Where previously the attacker was modeled using polynomial time Turing machines, we now need to change this to a quantum attacker.

It turns out that this modification is less straightforward than we expected. As explained previously, the classical BC logic uses a term representation that essentially corresponds to a tree in which computations, including those of the attacker, are performed in a purely local fashion. In the classical BC setting, this allows for modeling a single attacker using multiple local attackers: instead of explicitly communicating, subsequent local attackers effectively recompute the results of previous ones. This recomputation was possible since we made all operations deterministic and moved the probabilistic aspects (of both protocol and attacker) to two explicit random tapes. This modeling relies on a classical result for modeling probabilistic Turing machines as deterministic machines with a random tape.

However, for a quantum polynomial time attacker, we have no corresponding result. Notably, a quantum attacker may produce internal random state that we cannot influence nor extract. This phenomenon manifests itself in the so-called *no-cloning theorem* [61]: We cannot duplicate the quantum state of a quantum machine, nor can we run a quantum machine twice and expect to obtain the same inner state.

This directly breaks the stability of the previous interpretation, and for instance, the evaluation of Fig. 2 could now return false. Furthermore, in the interpretation of Fig. 3, we would not be able to recompute inside $\mathcal{T}_{\texttt{att}_1}$ the same attacker state as the one computed during $\mathcal{T}_{\texttt{att}_0}$, and therefore this evaluation would not correspond to the run of a single interactive quantum attacker, as was the case for classical attacker. Similarly, the classical BC interpretation also allows for *rewinding* the


Fig. 5. Quantum compatible simulation of attacker/protocol interactions without recomputation in PQ-BC. Compare this to Fig. 3.

attacker, which is impossible for quantum attackers based on the same no-cloning theorem.

We solve these issues by defining an interpretation where the attacker is a single black-box interactive machine, as opposed to a set of one-shot Turing machines. This allows us to verify a set of conditions that when met, ensure that we can provide a sound interpretation for a quantum attacker. Concretely, to interpret $\texttt{att}_0$, we query the interactive machine once, and to interpret $\texttt{att}_1(t)$, we interpret t, and give it as input to the same interactive machine. This effectively changes the term interpretation from a tree to a directed acyclic graph. We will formally define such an interpretation later, but it would intuitively be pictured as in Fig. 5. Changing the term interpretation in this fashion leads to a wide range of subtle changes throughout the logic.

We provide three conditions to ensure that we can give a sound interpretation for a quantum attacker. We provide intuition for them below, and formally define them in Section IV.

Condition 1 - **Consistency**: In the classical BC logic, the attackers need not be consistent. We provide a concrete example, that intuitively corresponds to the simulation of rewinding the attacker. For two distinct terms u, v, consider the sequence of terms $\texttt{att}_1(\texttt{att}_0(), u)$, $\texttt{att}_1(\texttt{att}_0(), v)$. This corresponds to running $\texttt{att}_0$, running its continuation with u, and then making it forget that we ever gave it u (rewind the attacker) and calling it on v.

Against a quantum attacker or black-box attacker, we cannot do this operation, as it morally implies that we are duplicating the attacker's inner state at the end of $\texttt{att}_0$, using each copy to run $\texttt{att}_1$ once on u and once on v.

We therefore add a syntactic condition on terms to forbid such cases, and force that all occurrences of $\texttt{att}_i$ are made with the same arguments.

Condition 2 - **Monotonicity**: With the new term interpretation, $\texttt{att}_0$ and $\texttt{att}_1$ implicitly share states. This means that when interpreting $\texttt{att}_1$, we will always get an answer that depends on the argument previously given to $\texttt{att}_0$.

Consider for example the terms:

$$\mathtt{att_0(n)},\ \mathtt{att_1() \doteq n}$$

In the classical BC logic, the second term would evaluate to false, as $\mathtt{att_1()}$ does not depend on the distribution of n. But if we are forced to interpret the symbols using a black box attacker, the computation of $\mathtt{att_1()}$ does depend on n. Thus, we lose the locality of our logic on some terms, that we will need to restore by further reducing the set of terms. This result in a second syntactic condition, where the arguments of the successive $\mathtt{att_i}$ must form a growing sequence of terms.

Condition 3 - **Balance***:* In the classical BC logic, we can write formulas that do not model interactions between a single attacker with a protocol, and for instance prove indistinguishability formulas that would be trivially false against a single attacker, but true for an independent set of attackers. For instance, consider the formula $\left(\mathtt{att_0() \doteq n}\right) \sim \left(\mathtt{att_1(att_0(), n') \doteq n}\right)$. Both sides evaluate to false, because the attacker is completely independent of n (this is the =IND rule). In the classical BC logic, the final distinguisher is also a disjoint machine, which typically does not know how many other attacker machines were executed: therefore, the formula is true in classical BC. This reflects that classical BC allows modeling a weaker threat model with independent attackers. In the PQ-BC interpretation, the final attacker will be a final call to the interactive machine, which of course knows how many times it was previously called.

Thus, under the new and strictly stronger interpretation with a single attacker, some formulas that held under the classical BC logic might no longer hold. We therefore have to ensure that none of our logical rules allows for deriving such formulas. We implement this by introducing a syntactic condition for PQ-BC that essentially requires that the number of attacker calls is balanced, i.e., the number of calls is equal on both sides of an indistinguishability operator, and which will be a side condition of all our logical rules.

Design choices for the conditions: While the above three conditions solve issues in designing a post-quantum sound BC logic, they were additionally chosen because they also form a small sufficient set of conditions to derive a usable PQ-BC logic, as we prove in the following section. It would have been possible, for instance, to replace the balance condition by a specific and more refined condition for each logical rule, that as a set would have been equivalent to the balance condition. However, the balance condition is both necessary and more generic, and hence we decide to use this one within PQ-BC. Overall, the three conditions allow for a generic implementation in the SQUIRREL prover, and yield a logic usable in practice.

IV. PQ-BC: A POST-QUANTUM BC LOGIC

In this section we first provide the core of the formal definition of the PQ-BC logic: its term interpretation that is suitable for post-quantum attackers, as well as the consistency and monotonicity conditions needed to ensure the stability and locality of the logic. This interpretation allows to consider a single interactive attacker, rather than a set of single-one shot attackers. In the long version [1] we also provide the computational soundness of the logic, which is a direct adaptation of the original BC soundness proof. Intuitively, the original proof needed to justify how an interactive attacker in the real world could be seen as a set of many one-shot attackers that needed to recompute the state of the previous attackers. With the new interpretation, we can directly use the real-world attacker. This implies that the logic can be used to obtained computational guarantees against quantum attackers. Second, we provide the structural rules of the logic, as well as the balance condition needed to prove the rules sound in the post-quantum setting. Finally, we discuss which cryptographic axioms - and thus which corresponding BC axioms - can be used inside the logic to get post-quantum soundness of a protocol analysis.

While we directly refer to a quantum attacker in definitions and proofs, we actually design a logic with an interpretation and rules that are computationally sound for any interactive black-box attacker. The attacker can be instantiated as a Polynomial Time attacker, or a Quantum Polynomial Time attacker, or even an unbounded Turing machine attacker. It is only the cryptographic assumptions used inside a given proof that restrict the attacker's computational power.

A. Syntax and Semantics

We use terms to model random samplings, public function computations by honest parties, and black-box attacker computations. For random samplings, the BC logic inherits some conventions from the Pi calculus: notably, fresh values (such as nonces) are called *names* (and have nothing to do with identities); by convention, variables called n, n', ... are names and hence freshly generated values. In Example 1 r, r', and sk would also be modeled as names. Let $\mathcal{N}$ be a set of names. Names can be seen as fixed identifiers, where each is a pointer to a uniformly sampled bitstring. Let Σ be a set of function symbols, the set used for public functions and primitives. Let $\{\mathtt{att_i} \mid i \in \mathbb{N}\}$ be a set of function symbols such that $\mathtt{att_i}$ is of arity i for each $i \in \mathbb{N}$.

Definition 1. *We consider terms built according to the syntax:*

$$
\begin{array}{llll}
t & ::= & \mathsf{n} \in \mathcal{N} & \textit{name (fresh value)} \\
  & \mid & f(t_1, \ldots, t_k) & \textit{function symbol } f \in \Sigma \\
  & \mid & att_i(t_1, \ldots, t_i) & \textit{i-th attacker call}
\end{array}
$$

We write $\vec{t}$ or $t_1, \ldots, t_n$ for sequences of terms.

1) Functional model: Recall that while we consider a quantum attacker, we model honest protocol participants as classical Polynomial Time Turing Machines (PTTMs). To interpret terms, we introduce the notion of a *functional model* $\mathcal{M}_f$, a library implementing the public function symbols and names that are used in the protocol: for each function symbol f (encryptions, signatures,...), $\mathcal{T}_f$ is a PTTM, which we view as a deterministic machine with an infinite random tape and taking the security parameter as input. The functional model also contains a PTTM $\mathcal{T}_n$ for each $n \in \mathcal{N}$, which will extract

from the random tape a bitstring of length η. We give η in unary to the PTTMs, as they are expected to be polynomial time w.r.t. η in the computational model.

Definition 2. *A functional model $\mathcal{M}_f$ is a set of PTTMs, one for each name and symbol function, such that:*

1) *if $n \in \mathcal{N}$ (i.e., n is a name), n is associated to the machine $\mathcal{T}_n$ that on input $(1^\eta, \rho_s)$ extracts a word of length η from the tape ρ_s. Different names extract disjoint parts of the random tape.*
2) *if $f \in \Sigma$ is of arity n, f is associated to a machine $\mathcal{T}_f$ which, on input 1^η, expects n more bitstrings, and does not use ρ_s.*

We model public functions as deterministic functions: if randomness is required, it should be given explicitly as an argument to the function symbol. This modeling is needed for the stability of the interpretation.

We can now define the basic interpretation of terms, assuming that we have been given the output bitstring corresponding to each attacker call. Based on this first interpretation, we will then define the one where there is an actual attacker.

Definition 3. *Given a functional model $\mathcal{M}_f$, the security parameter η, a mapping σ from terms $att_i(\phi)$ to bitstrings, and an infinite sequence of bitstrings ρ_s, we define the interpretation of terms such that all occurrences of $att_i(\phi_i)$ are in the domain of σ as:*

$$
\begin{aligned}
[\![n]\!]^{\eta,\sigma}_{\mathcal{M}_f,\rho_s} &:= \mathcal{T}_n(1^\eta, \rho_s) && \text{, if } n \in \mathcal{N} \\
[\![f(\vec{u})]\!]^{\eta,\sigma}_{\mathcal{M}_f,\rho_s} &:= \mathcal{T}_f([\![\vec{u}]\!]^{\eta,\sigma}_{\mathcal{M}_f,\rho_s}) && \text{, if } f \in \Sigma \\
[\![att_i(\vec{u})]\!]^{\eta,\sigma}_{\mathcal{M}_f,\rho_s} &:= att_i(\vec{u})\sigma && \text{, for all } i
\end{aligned}
$$

We assume that the functional model contains function symbols that expresses propositional formulas, which are interpreted as expected. We denote those connectives by the $\dot{=}, \dot{\wedge}, \dot{\Rightarrow}, \ldots$ – note these are marked with a dot.

We will ultimately use two different sets of logical connectives: (1) the dot variants, used in terms, and (2) the variants without a dot that are part of the logic we are building. We will illustrate their combination in Example 7.

2) Computational Model: To define the interpretation of terms with attacker function symbols, we view terms as directed acyclic graphs from leaves to their root.

Example 6. *Consider a variant of Example 1, where we denote tuples using $\langle \ldots \rangle$:*

$$
P := \textbf{new } sk. \textbf{ in}(c, x). \textbf{ new } r. \textbf{ out}(c, enc(x, r, sk)).
$$
$$
\textbf{in}(c, y). \textbf{ new } r'. \textbf{ out}(c, enc(\langle y, x, y \rangle, r', sk)).
$$

In the second step, the protocol encrypts the tuple made with twice the second protocol input on and once the first input.

The frame corresponding to this protocol would be t_0, t_1:

- $t_0 := enc(att_0(), r, sk)$
- $t_1 := enc(\langle att_1(t_0), att_0(), att_1(t_0) \rangle, r', sk)$

We give the original terms and the acyclic graph variant for this frame in Fig. 6.

The acyclic representation leads to a natural interpretation that we can execute even when we are only given access to a black-box straight line attacker. We illustrate this interpretation in Fig. 7. When we model the interactions with an interactive attacker, this corresponds to the high-level change between Fig. 3 and Fig. 5.

We assume that the attacker is an oracle Quantum Turing Machine (QTM) to obtain post-quantum soundness. We provide the formal definition of such machines in the long version [1]. For our purpose here, it suffices to know that such a machine behaves as a quantum computer that interactively performs oracle queries and receives the answers. Importantly, the oracle queries and answers are classical bitstrings, and do not contain any quantum state. This models a quantum computer interacting over a network with a classical protocol. One can also abstract such an attacker as a straight-line black-box interactive process, which is what we do in most of our proofs.

Definition 4. *A computational model $\mathcal{M}^{\mathcal{A}}$ is an extension of a functional model $\mathcal{M}_f$, which provides an additional oracle QTM $\mathcal{A}$ that takes as input a security parameter 1^η.*

Given $\mathcal{M}^{\mathcal{A}}$, η, σ, ρ_s and ρ_r, we define the interpretation $[\![t]\!]^\eta_{\mathcal{M}^{\mathcal{A}}, \rho_s}$ of a term t as:

1) *First, evaluate att_0, by running $\mathcal{A}$ on input 1^η until the first oracle query, and store the content of the oracle query tape o inside the substitution $\sigma_0 : \{att_0() \mapsto o\}$.*
2) *Then, we assume that we have a substitution σ_i mapping all occurrences of att_j, $j < i$ to a bitstring. We find the smallest subterm occurrence of a $att_l(t_1, \ldots, t_l)$ in t. Then, for k from i to l, we write on $\mathcal{A}$'s oracle answer tape the value $[\![t_k]\!]^{\eta,\sigma_{k-1}}_{\mathcal{M}_f,\rho_s}$, then continue $\mathcal{A}$ and wait for the next oracle query, and store the content of the oracle query tape o inside the substitution $\sigma_k = \sigma_{k-1} \cup \{att_k(t_1, \ldots, t_k) \mapsto o\}$.*
3) *Finally, given σ_l, where l corresponds to the biggest occurrence of a att_l, return $[\![t]\!]^{\eta,\sigma_l}_{\mathcal{M}_f,\rho_s}$.*

This interpretation does not require that a term contains all intermediate calls to the attacker. Thus, interpreting the term $att_1(t)$ or the sequence $att_0(), att_1(t)$ leads to the same interactions with the attacker.

3) Well-defined interpretation: Our previous definition is not defined over all possible terms: there may not be a unique smallest occurrence of a att_l. This is expected, as there exist terms that correspond to experiments that are not realisable with a quantum attacker. Notably, recall that in this context, we cannot interpret the sequence $att_0(), att_1(u), att_1(v)$. Indeed, if the attacker is straight-line and black-box, it means that we can only get one attacker answer corresponding to att_1. Therefore, to ensure that a term corresponds to a valid experiment with respect to a quantum attacker, we require consistency: att_i should always occur with the same arguments in the terms, i.e., corresponds to the same unique call of this attacker's interactive step.

Definition 5 (Consistency). *A sequence of terms $\vec{t}$ is consistent if all function symbol att_i occurs with the same arguments.*

Fig. 6. From the classical BC to the acyclic graph representation in PQ-BC

Fig. 7. Quantum compatible interpretation in PQ-BC

Lemma 1. $[\![\vec{t}]\!]^{\eta}_{\mathcal{M}^{\mathcal{A}},\rho_s}$ *is well-defined if $\vec{t}$ is consistent.*

Proof sketch. In the loops of step 2, all occurrences $\mathtt{att}_i$ will occur with the same sequence of arguments $(t_1, \ldots, t_i)$. Thus, at the end of step 2, σ will indeed be a total mapping over all terms that have as head an attacker function symbol, and the interpretation will terminate. $\square$

4) Restoring locality: In the evaluation of the sequence $\mathtt{att}_0(), \mathtt{att}_1(u), \mathtt{att}_2(v, v')$, the value of $\mathtt{att}_2(v, v')$ depends on u in the new interpretation. Yet, this is not reflected inside the term and introduces an implicit dependency that breaks some BC rules as it breaks locality. In particular, it breaks the =IND rule presented in Section II. To repair it, we introduce the second term restriction.

Definition 6 (Monotonicity)**.** *A sequence of terms $\vec{t}$ satisfies the monotonicity restriction if for all occurrences of $\mathtt{att}_i(t_1, \ldots, t_i)$ and $\mathtt{att}_j(t'_1, \ldots, t'_j)$ with $i < j$, we have for all $1 \le k \le i$ that $t_i = t'_i$.*

Essentially, to restore locality, we require that attacker function symbols must always be called on at least the same set of argument as the previous attacker calls.

Equipped with this new restriction, we can prove a lemma that expresses the locality property of the interpretation. It can be found in the long version [1].

From now on, we only allow in the logic terms that satisfy the consistency and the monotonicity properties.

5) Interpretation of formulas: Atomic formulas of the logic are built using a set of predicate symbols $\sim_n$ of arity $2n$. Given terms $t_1, \ldots, t_n, s_1, \ldots, s_n$, the predicate $\sim_n (t_1, \ldots, t_n, s_1, \ldots, s_n)$ will be interpreted as computational indistinguishability between the two sequences of terms. We use infix notation, and always omit n as it is clear from the context, thus denoting the previous equivalence by $(t_1, \ldots, t_n) \sim (s_1, \ldots, s_n)$. The first order formulas are then built using the usual logical connectives $\vee, \wedge, \top, \bot, \Rightarrow, \exists, \forall, \neg$. Note that these connectives are not marked with a dot, and are part of the logic, not the terms.

Example 7. *Given terms u, v, t, we can write the formula:*

$$\big((u \dot{\Rightarrow} v) \sim t\big) \Leftrightarrow \big((\dot{\neg} \, u \, \dot{\vee} \, v) \sim t\big)$$

This formula holds because we assume that for all functional models, $\dot{\Rightarrow}, \dot{\neg}$ and $\dot{\vee}$ are classically interpreted.

To define a distinguisher between sequences of terms, we must define this distinguisher as a continuation of the interactive attacker $\mathcal{A}$, and thus pass to it the quantum state of $\mathcal{A}$. This is only a technicality, and in practice we simply consider that the continuation is the last stage of the interactive attacker.

Definition 7. *Let $\vec{v}$ be a sequence of term. For any computational model $\mathcal{M}^{\mathcal{A}}$, we denote by $\phi^{\vec{v}}_{\mathcal{M}^{\mathcal{A}},\rho_s,\eta}$ the final (quantum) configuration reached by $\mathcal{A}$ during a computation of $[\![\vec{v}]\!]^{\eta}_{\mathcal{M}^{\mathcal{A}},\rho_s}$.*

Definition 8. *Given a computational model $\mathcal{M}^{\mathcal{A}}$, and two sequences of ground terms $\vec{t}, \vec{u}$, the formula $\vec{t} \sim \vec{u}$ is satisfied by $\mathcal{M}^{\mathcal{A}}$ if, for every polynomial time QTM $\mathcal{B}$,*

$$|\mathbf{Pr}_{\rho_s}\{\mathcal{B}([\![\vec{t}]\!]^{\eta}_{\mathcal{M}^{\mathcal{A}},\rho_s}, 1^{\eta}, \phi^{\vec{t}}_{\mathcal{M}^{\mathcal{A}},\rho_s,\eta}) = 1\}$$
$$-\mathbf{Pr}_{\rho_s}\{\mathcal{B}([\![\vec{u}]\!]^{,\eta}_{\mathcal{M}^{\mathcal{A}},\rho_s}, 1^{\eta}, \phi^{\vec{u}}_{\mathcal{M}^{\mathcal{A}},\rho_s,\eta}) = 1\}|$$

is negligible in η. Here, ρ_s is drawn according to a distribution such that every finite prefix is uniformly sampled, and the probabilities also depend on the inherent probabilistic behavior of the two QTMs. The satisfaction relation is extended to full first-order logic as usual.

In the long version [1] we give the proof of the logic's computational soundness. At a high level, the proof follows naturally from our previous design of a faithful interpretation

$$\frac{\text{=-REFL}}{\mathbb{T}(u \doteq u)} \qquad \frac{\text{=-IND}}{\mathbb{T}(u \neq \mathsf{n})} \text{ when n is not } \atop \text{a subterm of } u \qquad \frac{\text{=-SYM}}{\mathbb{T}(u \doteq v)}{\mathbb{T}(v \doteq u)}$$

Fig. 8. Truth rules of PQ-BC (identical to those of BC)

of terms that is sound for black-box interactive attackers. We directly obtain that by quantifying over all possible $\mathcal{M}^{\mathcal{A}}$, the value of $[\![\vec{t}]\!]^{\eta}_{\mathcal{M}^{\mathcal{A}}, \rho_s}$ describes all possible behaviours of the protocol modelled by t interacting with any black-box interactive attacker. Thus, if there exists an attack on the protocol in the real world, it will correspond to an attack on the (computational) interpretation of the terms. Furthermore, because Definition 8 exactly quantifies universally over all $\mathcal{M}^{\mathcal{A}}$, if there is an attack on the protocol, the predicate is not valid, and if there is none, it is valid. Finally, because an interactive black-box attacker soundly models a quantum attacker, the logic is shown to be post-quantum sound.

6) Overwhelming probabilistic truth: The classical BC logic as well as SQUIRREL has a subset of its rule dedicated to proving the validity of statements of the form $u \sim \mathsf{true}$.

In our case the attacker $\mathcal{A}$ from the computational model and the final distinguisher $\mathcal{B}$ share their (quantum) tape, while in the classical BC definition [7], $\mathcal{A}$ and $\mathcal{B}$ do not share their working tape, but only their source of randomness Thus, compared to the classical BC definition, as soon as u executes an attacker machine, $(u \doteq u) \sim \mathsf{true}$ will not hold, as $\mathcal{B}$ simply checks if an attacker machine was executed.

Note that in the classical BC logic, a proof that $u \sim \mathsf{true}$ is in fact a statement independent of the final distinguisher: we just prove that u is equal to true with overwhelming probability. This is what we call a probabilistic statement.

For PQ-BC, we thus define a predicate dedicated to proving overwhelming probabilistic truth, for which we inherit all the truth rules of BC and SQUIRREL.

Definition 9. *Given a computational model $\mathcal{M}^{\mathcal{A}}$ and a ground term t, $\mathbb{T}(t)$ is satisfied by $\mathcal{M}^{\mathcal{A}}$ if,*

$$\mathbf{Pr}_{\rho_s}\{[\![\vec{t}]\!]^{\eta}_{\mathcal{M}^{\mathcal{A}}, \rho_s} = 1\}$$

is overwhelming in η. The satisfaction relation is extended to full first-order logic as usual.

B. Logical rules

1) Probabilistic statements: We introduce some of the logical rules to reason about the $\mathbb{T}()$ predicate, which are all direct transpositions of the rules of the BC logic that are statements about formulas of the form $u \sim \mathsf{true}$. Their soundness proofs are completely similar, as the corresponding BC rules are in fact sound for any Turing Machine, even with unbounded computational power. This means that if the premices are valid, so are the conclusions. We present in Fig. 8 a subset of such rules that allows to reason about the equality between terms, with the other rules being transposed similarly.

Lemma 2. *The truth rules, shown in Fig. 8, are sound in the PQ-BC interpretation.*

Omitted proofs can be found in [1].

2) Indistinguishability rules: As mentioned earlier, some statements of classical BC are false under the single-attacker interpretation. For example, $u \sim \mathsf{true}$ is false as soon as u contains an attacker symbol, because the final distinguisher also executes the corresponding attacker call, and could therefore differentiate the sides. Yet, such statements are provable inside BC, because it does not assume the final distinguisher sees all attacker calls. Thus, the existing BC rules are not correct in our new interpretation, and we must add a side condition to make PQ-BC sound.

Similar to the rules for the truth predicate $\mathbb{T}()$, the indistinguishability rules in Fig. 9 are also sound for any attacker, without any assumption on their computational power. Thus, the soundness issues only come from the fact that the distinguisher $\mathcal{B}$ now inherits the state of the attacker $\mathcal{A}$. In the end, the main case where an existing BC rule becomes unsound is when it yields in the conclusion a $\sim$ statement where the attacker $\mathcal{A}$ is not called the same number of times on both sides, which corresponds to the *balance* condition.

Definition 10 (Balance). *Given a sequence of terms $\vec{u}$, we denote by $\mathsf{Max}_{att}(\vec{u})$ the biggest index i such that the function symbol $\mathtt{att}_i$ appears in $\vec{u}$.*
We say that $\vec{u} \sim \vec{v}$ satisfies the balance conditions, or is balanced, if $\mathsf{Max}_{att}(u) = \mathsf{Max}_{att}(v)$.

Most BC rules can then be transformed to PQ-BC rules by additionally requiring the *balance* condition over their indistinguishability. In practice, we will thus require that the balance condition holds over all indistinguishabilities appearing inside the proof tree. In the following, however, we only add the side condition where it is needed, in order to pinpoint which rules may break the balance condition. We present in Fig. 9 a subset of the new rules, using the color blue to indicate the new side-conditions. Note that the consistency and monotonicity conditions are enforced globally, and hence also hold for all terms appearing in rules. The CS rule requires a stronger condition than the balance condition. Though we provide it for completeness, this rule is not used in SQUIRREL nor in PQ-SQUIRREL, and the balance condition is thus sufficient for their rules to get post-quantum soundness.

Lemma 3. *The indistinguishability rules, shown in Fig. 9, are sound in the PQ-BC interpretation.*

Intuitively, this means that whenever we construct a proof in the logic, then if there is an attack on the proven formula, there is an attack on the axioms. If we combine this with the soundness of the logic, we get that a proof in the logic implies the existence of a post-quantum sound reduction from an attack on the protocol to an attack on post-quantum sound cryptographic assumptions.

11

$$\text{REFL} \quad \frac{}{\vec{u} \sim \vec{u}}$$

$$\text{TRANS} \quad \frac{\vec{u} \sim \vec{v} \quad \vec{v} \sim \vec{w}}{\vec{u} \sim \vec{w}}$$

$$\text{SYM} \quad \frac{\vec{v} \sim \vec{u}}{\vec{u} \sim \vec{v}}$$

$$\text{FRESH} \quad \frac{\vec{u} \sim_g \vec{v}}{(\vec{u}, \mathsf{n}) \sim_g (\vec{v}, \mathsf{n}')} \text{ when the new names do not occur in } \vec{u}, \vec{v}$$

$$\text{FA} \quad \frac{\vec{u} \sim \vec{v}}{h(\vec{u}) \sim h(\vec{v})}$$

$$\text{RESTR} \quad \frac{(\vec{u}, s) \sim (\vec{v}, t)}{\vec{u} \sim \vec{v}} \text{when } \mathsf{Max}_{\mathrm{att}}(\vec{u}) = \mathsf{Max}_{\mathrm{att}}(\vec{v})$$

$$\text{DUP} \quad \frac{(\vec{u}, t) \sim (\vec{v}, t)}{(\vec{u}, t, t) \sim (\vec{v}, t, t)}$$

$$\text{EQU} \quad \frac{\mathbb{T}(s \doteq t)}{(\vec{u}, s) \sim_g (\vec{u}, t)} \text{when } \mathsf{Max}_{\mathrm{att}}(\vec{u}, s) = \mathsf{Max}_{\mathrm{att}}(\vec{u}, t)$$

$$\text{IFT} \quad \frac{\big(\vec{u}, C[\textbf{if } s \doteq t \textbf{ then } C_0[s] \textbf{ else } w]\big) \sim \vec{v}}{\big(\vec{u}, C[\textbf{if } s \doteq t \textbf{ then } C_0[t] \textbf{ else } w]\big) \sim \vec{v}}$$

$$\text{CS} \quad \frac{(\vec{w}, b, u) \sim (\vec{w}', b', u') \quad (\vec{w}, b, v) \sim (\vec{w}', b', v')}{(\vec{w}, \textbf{if } b \textbf{ then } u \textbf{ else } v) \sim (\vec{w}', \textbf{if } b' \textbf{ then } u' \textbf{ else } v')} \quad \begin{array}{l} \text{when } \mathsf{Max}_{\mathrm{att}}(\vec{w}, b, u) = \mathsf{Max}_{\mathrm{att}}(\vec{w}, b, v) \\ = \mathsf{Max}_{\mathrm{att}}(\vec{w}', b', u') = \mathsf{Max}_{\mathrm{att}}(\vec{w}', b', v') \end{array}$$

Fig. 9. Indistinguishability rules of PQ-BC. We mark new side conditions in blue.

C. Cryptographic assumptions in PQ-BC

In Section II-D we discussed how BC encodes cryptographic assumptions, such as PRF, IND-CCA, EUF-CMA, ENC-KP, INT-CTXT, OTP, and DDH. The original proofs of soundness of these encodings in [7], [46] were aimed at a classical attacker. We revisited all these proofs, and due to their direct black-box nature, it turns out these proofs also directly apply against a post-quantum attacker: if there exist an instantiating that satisfies the assumption against a quantum attacker, then the corresponding BC axiom is post-quantum sound.

However, knowing that a BC rule is post-quantum sound w.r.t. the cryptographic assumption does not mean that we know an instantiation (i.e., a concrete scheme) that is secure with respect to a quantum attacker. While we know instantiations for most of the above assumptions with respect to a classical attacker, at this moment we do not know of a post-quantum secure instantiation of the DDH assumption. In the future, a candidate for post-quantum DDH could be the CSI-DDH [42] assumption, based on the CSIDH assumption [21]. Their concrete security is however the subject of discussions [12], [14], [17], [51]. We therefore omit this for now from list of allowed cryptographic assumptions for PQ-BC.

This yields the following list of currently usable PQ-BC axioms for post-quantum proofs: PRF, IND-CCA, EUF-CMA, ENC-KP, INT-CTXT, and OTP. Concretely, this means that a proof in PQ-BC yields guarantees for post-quantum attackers under the assumption that the previous axioms are instantiated in a post-quantum secure way, e.g., by a protocol that uses CRYSTALS-Dilithium [32] to instantiate EUF-CMA, and see e.g. [62] for a post-quantum sound instantiation of a PRF. Out of these assumptions, the most debatable w.r.t. instantiability is likely ENC-KP, as discussed recently in [37].

V. MECHANIZATION IN PQ-SQUIRREL AND CASE STUDIES

In this section, we describe PQ-SQUIRREL, our extension of SQUIRREL that produces post-quantum sound proofs in the PQ-BC logic. We give an overview of the post-quantum protocol analysis results that we obtained using PQ-SQUIRREL in Table I. As we will show later, it turns out that despite the new term interpretation and corresponding new side conditions, several existing SQUIRREL proofs could be re-interpreted by PQ-SQUIRREL as post-quantum sound proofs in PQ-BC.

A. PQ-SQUIRREL

1) Ensuring post-quantum soundness: To make SQUIRREL post-quantum sound, we must specify which cryptographic axioms are post-quantum sound, and enforce the three syntactic side conditions from Definitions 5, 6 and 10. Furthermore, if there are cryptographic assumptions for which we know instantiations that are secure against classical attackers, but do not know any post-quantum secure instantiations, we no longer rely on them.

Recall that a substantial amount of work in adapting the BC logic to PQ-BC was due to the flexibility in specifying multiple attackers. In contrast, SQUIRREL specifications do not include attacker terms: SQUIRREL automatically produces the attacker terms from the input and output commands in the process specification, under the assumption that there is only one attacker. Concretely, the attacker terms are produced by the interpretation of $\mathsf{input}^P@\tau$, which is equal to $\mathtt{att}_\tau(\mathsf{frame}^P@\mathsf{pre}(\tau))$. While this means that SQUIRREL only supports a subset of the BC logic, it strongly simplifies protocol specification for the user, and prevents users from accidentally modeling a weaker threat model with multiple disjoint attackers. For our post-quantum purposes, this historical choice is very convenient: all terms produced by the meta-logic of SQUIRREL already satisfy the consistency and monotonicity properties.

We still need to ensure that PQ-SQUIRREL verifies the new side conditions w.r.t. the $\mathsf{Max}_{\mathrm{att}}()$ on both sides of the indistinguishability formulas from Fig. 9. Given a meta-logic formula $\vec{u} \sim \vec{v}$, we need to check that for the maximal timestamp element τ of the trace, $\mathsf{input}@\tau$ appears on both sides. If this is the case, we say that a formula satisfies the *synchronization* property.

We cannot check the synchronization property directly since SQUIRREL internally omits the inputs from the frames when they are redundant, and would lead to falsely discarding proof steps. We instead check a generalized property, which is true if either input@τ or frame@τ or frame@pre(τ) occurs in the frame. In the long version [1] we prove a lemma that shows that a proof in which all steps satisfy this generalized property can be mapped to a post-quantum proof.

2) Implementation: PQ-SQUIRREL can operate in classic SQUIRREL mode. Additionally, it offers a post-quantum-mode switch that can be enabled inside proof files. When enabled, PQ-SQUIRREL operates in post-quantum mode: it only allows tactics and axioms that have been proven post-quantum sound, and checks synchronization for every indistinguishability appearing at any step of a proof.

The source-code of PQ-SQUIRREL is available at [2]. Thanks to our identification of a minimal set of simple syntactic conditions, the PQ-SQUIRREL extension could be concisely implemented, and only comprises a few hundred line of codes in addition to SQUIRREL's code base.

B. Case studies

We summarize the case studies we performed using PQ-SQUIRREL in Table I. They fall into two categories: new case studies for the Internet Key Exchange (IKE) standards and of key exchange protocols based on Key-Encapsulation Mechanisms (KEMs); and previous SQUIRREL case studies that we could prove post-quantum sound in PQ-SQUIRREL.

All model files and the prover are at [2].

Required effort: In total, modeling the protocols took in the order of hours, and constructing their proofs in interaction with PQ-SQUIRREL took in the order of weeks. PQ-SQUIRREL verifies each resulting proof file in under 10 seconds on a laptop with a quad-core CPU at 1,8GHz.

As a side contribution, we also developed new generic tactics for the prover that enabled the case studies. We first present case studies in Sections V-C and V-D, and then introduce the new tactics in Section V-E.

C. Key exchange case studies: IKE and KEM-based

1) Threat model and security properties: We modeled five key exchange protocols, for which we proved, e.g.:

- Authentication - if a party accepts, another accepts (with the same parameters).
- Strong Secrecy - the keys derived by the parties are indistinguishable from fresh random values.

In these initial case studies, we consider the same threat model for all key exchange protocols: an arbitrary number of initiators and responders willing to answer to anybody, including to dishonest/compromised identities with attacker-controlled keys. We did not yet prove properties with respect to dynamic corruptions nor more complex security properties like perfect forward secrecy, and we consider them out of scope for this work. We stress however that in a similar fashion to EASYCRYPT and CRYPTOVERIF, PQ-SQUIRREL does not have any hard-coded threat model, and these case studies can be extended in future work.

2) IKE case studies: The IKE standards, version 1 [41] and 2 [41], specify suites of key exchanges. They are Diffie-Hellman key exchanges that support multiple authentication modes. RFC 8784 [34] addresses the issue of quantum computers breaking the DDH assumption, and its authors claim that the authentication mode based on a pre-shared key in version 1 (IKEV1$_{\text{PSK}}$) is post-quantum sound. For the same purpose, they also define a way to extend version 2 so that the final key computation depends on a pre-shared key.

For IKEV1$_{\text{PSK}}$, we use PQ-SQUIRREL to prove that a pre-shared key between two entities allows to derive an authenticated secret key indistinguishable from a random.

We also analyze the version 2 protocol with signatures for the authentication and extended with the pre-shared key, which we call IKEV2$_{\text{PSK}}^{\text{SIGN}}$. The proof of IKEV2$_{\text{PSK}}^{\text{SIGN}}$ is simpler than IKEV1$_{\text{PSK}}$, because the signatures simplify the derivation of the authentication property.

3) KEM based key exchanges: KEMs are currently considered as a possible replacement for DH-like key exchanges. KEMs abstract mechanisms that generate and send fresh key material encrypted to another party, from which both parties derive a fresh shared key. Some generic constructions of KEM based key exchanges have been proposed in [18], [35], and have for instance been expanded into a full alternative to TLS in [54] or as post-quantum sound variants of the Signal X3DH handshake [38]. These key exchanges were specifically designed to not rely on any DH-like operations, and their security instead relies on assumptions on the corresponding KEM constructions, i.e., IND-CCA.

In PQ-SQUIRREL, we generally model KEM-based key exchanges by modeling some common internals of KEMs: generating fresh key material, sending this encrypted to the other party, and then deriving a key from this material at both parties using a key derivation function.

The *basic KEM-based key exchange pattern* is to perform the KEM operation at both parties with respect to their peer's long-term public keys, and then to xor the two resulting fresh keys (one for each direction). This generic pattern was illustrated in our example from Fig. 4. As the knowledge of both fresh keys is needed to derive the final key, the attacker cannot obtain it unless it knows both long-term private keys. Note that such schemes provide implicit authentication, but not (explicit) authentication: only a trusted party can derive the final key, but there is no guarantee that such a party exists.

In our specification, we use `enc` to talk about an abstract KEM construction, while in practice, it is referred to as the encapsulation mechanism, and the decryption is the decapsulation. Presenting it using an encryption symbol directly allows to model it inside PQ-SQUIRREL, but this does not affect the validity of the proofs.

For all our KEM based case studies, our models include an unbounded number of initiators with distinct secret decapsulation keys sk_I, each willing to initiate an unbounded number of sessions with any honest responder with encapsulation key pk_R, as well as an unbounded number of honest responders

Protocol	LoC	Primitives and Assumptions				Security properties
		h	sign	enc	⊕	
New case studies of key exchange protocols						
IKEV1$_{PSK}$ [20]	850	PRF				Strong Secrecy & Authentication
IKEV2$_{PSK}^{SIGN}$ [34], [41]	300	PRF	EUF-CMA			Strong Secrecy & Authentication
KE$_{BCGNP}$ [18]	355	PRF		IND-CCA	OTP	Strong Secrecy
KE$_{FSXY}$ [35]	620	PRF		IND-CCA	OTP	Strong Secrecy
SC-AKE [38]	745	PRF	SUF-CMA	IND-CCA	OTP	Strong Secrecy & Authentication
Proving post-quantum soundness of SQUIRREL *case studies* [5]						
Basic Hash [19]	100	PRF				Authentication & Unlinkability
Hash Lock [40]	130	PRF				Authentication & Unlinkability
LAK (with pairs) [39]	250	PRF				Authentication & Unlinkability
MW [49]	300	PRF			OTP	Authentication & Unlinkability
Feldhofer [33]	270			ENC-KP INT-CTXT		Authentication & Unlinkability
Private Authentication [7]	100			ENC-KP IND-CCA		Anonymity

TABLE I

PQ-SQUIRREL CASE STUDIES: WE CONSTRUCTED NEW MODELS OF KEY EXCHANGE PROTOCOLS WITH STATIC KEY COMPROMISE, AND REVISITED PREVIOUS SQUIRREL PROTOCOL MODELS. PQ-SQUIRREL PROVES THAT THESE PROTOCOLS ARE COMPUTATIONALLY POST-QUANTUM SECURE WHEN THEY ARE IMPLEMENTED WITH POST-QUANTUM SECURE PRIMITIVES FOR EACH OF THEIR ASSUMPTIONS.

willing to engage with an unbounded number of sessions with any arbitrary public key that may be attacker-controlled.

Recall that a generic KEM based construction does not provide explicit authentication properties; thus, for KEM based key exchange we only prove the strong secrecy of the derived keys. In our basic example of Fig. 4, this would be achieved in three steps, first by using IND-CCA to hide the secret ephemeral keys k_I and k_R from the attacker; second by using the PRF assumption to derive valid keys, i.e., showing that $kdf(e_x)$ is indistinguishable from a fresh random n_x; third by using the OTP assumptions that enforces the one-time pad property, and thus that the final key is indistinguishable from random, as it is always equal to $n_x \oplus t$ for some t and fresh random bitstring n_x. For illustration purposes, the actual PQ-SQUIRREL proof corresponding to this example can be found along with the other case studies at [2]. Interestingly, the proofs carried out in PQ-SQUIRREL follow this high-level structure for the multiple KEM based case studies, and were thus straightforward to establish.

4) Concrete case studies: The KE$_{BCGNP}$ protocol closely follows the generic pattern. KE$_{FSXY}$ uses an additional ephemeral key used for each session, as well as an additional round of PRF application to the key materials before xor-ing them. We prove the strong secrecy of the derived keys for both protocols.

The SC-AKE protocol, intended as a possible post-quantum replacement of Signal's X3DH, can be seen as a variant of KE$_{FSXY}$ extended with a third message send from the Responder to the Initiator, containing a signature to provide a form of deniable authentication. Instead of deriving a single key $k := kdf(sid, e_i) \oplus kdf(sid, e_r)$, it derives two keys k_s and k_f using $kdf1$ and $kdf2$. The first one is used to xor, and thus hide, the signature of the sid sent for authentication, and the second one is the derived key. Because of these constructions

and their properties, SC-AKE is our most complex case study. The proof first requires proving the authentication of the responder to the initiator, by relying on the EUF-CMA assumption on the signature. After having shown that the material used by the initiator to derive the secret key is from an honest source, we can show that the secret key is strongly secret by following the previous pattern. Such proofs illustrate a strength of the PQ-SQUIRREL prover: it allows interactions between a part of the logic dedicated to proving reachability properties (e.g. authentication), and then use those properties inside indistinguishability proofs (e.g. secrecy).

D. Proving post-quantum soundness of SQUIRREL *case studies*

We used PQ-SQUIRREL to verify the proofs of the nine previous SQUIRREL case studies. Out of those, PQ-SQUIRREL was able to prove that six were post-quantum sound, and three were not, as they relied on the DDH assumption.

Thus, it seems that most existing proofs in SQUIRREL are already post-quantum sound, even though we know it is possible to prove statements in SQUIRREL that are not post-quantum sound. This appears to be because the proofs of realistic protocols rely on an induction on the length of the trace. We then reason on frames of protocols and prove that each of their possible last messages does not break the indistinguishability. This pattern seems to avoid violating the balance condition.

E. Additional tactics

1) A Non-Malleability tactic: SQUIRREL already had a tactic for the IND-CCA axiom, which is the one usually used for KEM. However, the IND-CCA axiom is not only used to provide secrecy in the context of KEM, but also a form of weak authentication. If a party receives the ciphertext of something that corresponds to an honest ephemeral share, then there exists a session of an honest initiator that sent it. To prove such an authentication property, the IND-CCA axiom is ill-suited,

because we rely on the non-malleability property of the scheme, which is implied by IND-CCA [11], [50]. We developed a new tactic that allows to say that encrypted honest secret share cannot have been tampered with by the attacker, and must have been sent by some honest party. In [1], we provide the formal definition of the meta-logic rule, as well as its soundness proof from the original IND-CCA BC axiom. Using this tactic, we could directly prove the weak authentication of the schemes.

2) Global tactics: Inside SQUIRREL proofs, we often consider statements of the form

$$\mathsf{frame@pre}(\tau), t \quad \mathsf{frame@pre}(\tau), t'$$

We then show that t and t' are indistinguishable while leaving the frame abstract. Most SQUIRREL tactics only allow manipulating the terms appearing inside t and t'. However, we sometimes need to perform actions globally: not only on t and t' but also on the terms that may appear inside $\mathsf{frame@pre}(\tau)$, i.e., all the terms inside the protocol. We implemented four new tactics that enable such global manipulations:

- a tactic to globally substitute a name by a fresh name;
- a tactic that allows to prove the indistinguishability of two protocols by proving their functional equivalence, i.e., that they in fact produce exactly the same distributions with overwhelming probability;
- an IND-CCA2 tactic to replace all occurrences of some cipher by a version with perfectly hidden plaintext; and
- a PRF tactic to replace all occurrences of the hash of a given message by the same random.

In comparison, the original SQUIRREL tactic for PRF only allows to replace one instance of a hash inside t by a random as long as one can prove this hash was never computed before by the protocol.

We provide the formal definitions as well as proof of soundness for the tactics in [1].

VI. CURRENT LIMITATIONS AND FUTURE WORK

A. Minimality of the syntactic conditions

Each of the three syntactic conditions is needed to forbid unsound operations over quantum attackers, as illustrated by the three examples in Section III. However, they may not be the minimal possible conditions, and our conditions do reduce expressivity. Notably, our current restrictions rule out zero knowledge proofs, whose security analysis often requires rewinding. However, our current assessment is that any weakening of the conditions would inherently be very domain specific, and thus only useful for a small set of protocols. For instance, while we could have loosened the consistency condition to allow for some particular form of post-quantum sound rewinding [58], [60], all such techniques are dedicated to particular cases. We currently believe it would be very difficult to derive a general post-quantum rewinding technique (see e.g. [4]).

B. Scope of the tool

In terms of security properties, we have already used PQ-SQUIRREL to verify a range of properties like unlinkability, anonymity, strong secrecy, and authentication. In general, it allows expressing properties using arbitrary first-order logic formula, which can mix indistinguishability properties and reachability properties, and can thus be used to express all classical security properties.

For protocols, PQ-SQUIRREL cannot currently carry out proofs that require rewinding or that are in the ROM, but such restriction do not hinder proving a wide range of protocols from the literature. As discussed previously, it is unclear which kind of rewinding should be integrated currently in Squirrel. Integrating the ROM (and QROM) into PQ-BC and PQ-SQUIRREL could be interesting future work. However, while (Q)ROM is often needed for the analysis of primitives, it is less often needed for protocols. A class of protocol we cannot typically verify are e-voting protocols, that often rely on the ROM. However, there is a long-standing debate on the ROM model, because it is not realizable in practice. Typically, it is often preferable to use the PRF assumption when a secret seed is derived in the protocol, as we did in our case studies. In other cases, integrating results such as [52] into PQ-BC might offer a solution in the future.

Besides protocols whose proofs require rewinding or the ROM, we are not aware of any inherent limitations of the logic or the prover that would hinder the proofs of other post-quantum protocols.

C. Refining the case studies

The goal of our case studies is to show the usability and scalability of the tool, not to provide an exhaustive analysis of each of them. For example, our current key exchange analysis only consider a model with static compromise. However, this is not an inherent limitation of the logic, and the threat model is only part of the protocol modeling. As such, a natural future work is to refine our case studies and try to prove more advanced properties such as PFS or PCS.

D. Improving automation

Based on our first set of case studies, PQ-SQUIRREL shows the potential to tackle complex case studies such as the proposal of a recent post-quantum TLS [55]. Performing such complex case studies would benefit from improving the automation in PQ-SQUIRREL. We see two main possible routes to achieve this. First, some low level reasoning about message equalities and inequalities is already automated in many cases, but could be improved by leveraging SMT solvers. Second, the unique abstraction level of our logic enables us to reason both at the high-level of the executions traces as well as at the low level of the indistinguishability of two messages. This abstraction opens the door for the application of more advanced constraint solving techniques, similar to the one used in symbolic tools, which can further improve automation.

VII. Conclusion

We defined the PQ-BC logic for proving protocol security against quantum attackers, and a corresponding PQ-Squirrel prover to mechanize the reasoning. In the process of extending the BC logic for this purpose, we identified three simple syntactic side conditions that are both necessary and sufficient; these conditions, and the new tactics we developed for PQ-Squirrel, can be useful for the classical setting as well. Our initial case studies show that PQ-Squirrel can be effectively used to prove post-quantum protocol security.

References

[1] Long version of this paper. https://hal.inria.fr/hal-03620358.

[2] The Squirrel Prover repository. https://github.com/squirrel-prover/squirrel-prover/.

[3] José Bacelar Almeida, Manuel Barbosa, Gilles Barthe, Matthew Campagna, Ernie Cohen, Benjamin Gregoire, Vitor Pereira, Bernardo Portela, Pierre-Yves Strub, and Serdar Tasiran. A Machine-Checked Proof of Security for AWS Key Management Service. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 63–78, London United Kingdom, November 2019. ACM.

[4] Andris Ambainis, Ansis Rosmanis, and Dominique Unruh. Quantum attacks on classical proof systems: the hardness of quantum rewinding. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 474–483. IEEE, 2014.

[5] David Baelde, Stéphanie Delaune, Charlie Jacomme, Adrien Koutsos, and Solène Moreau. An interactive prover for protocol verification in the computational model. In *42nd IEEE Symposium on Security and Privacy (S&P 2021)*. IEEE, 2021.

[6] Gergei Bana, Rohit Chadha, and Ajay Kumar Eeralla. Formal Analysis of Vote Privacy Using Computationally Complete Symbolic Attacker. In *Computer Security - 23rd European Symposium on Research in Computer Security, ESORICS 2018, Barcelona, Spain, September 3-7, 2018, Proceedings, Part II*, pages 350–372, 2018.

[7] Gergei Bana and Hubert Comon-Lundh. A Computationally Complete Symbolic Attacker for Equivalence Properties. In *Proceedings of the 21st ACM Conference on Computer and Communications Security (CCS'14)*, pages 609–620, Scottsdale, Arizona, USA, November 2014. ACM Press.

[8] Manuel Barbosa, Gilles Barthe, Xiong Fan, Benjamin Grégoire, Shih-Han Hung, Jonathan Katz, Pierre-Yves Strub, Xiaodi Wu, and Li Zhou. EasyPQC: Verifying post-quantum cryptography. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, pages 2564–2586, 2021.

[9] Gilles Barthe, François Dupressoir, Benjamin Grégoire, César Kunz, Benedikt Schmidt, and Pierre-Yves Strub. EasyCrypt: A Tutorial. In *Foundations of Security Analysis and Design VII: FOSAD 2012/2013 Tutorial Lectures*, Lecture Notes in Computer Science, pages 146–166. Springer International Publishing, Cham, 2014.

[10] David Basin, Ralf Sasse, and Jorge Toro-Pozo. Card brand mixup attack: Bypassing the PIN in non-Visa cards by using them for Visa transactions. In *30th USENIX Security Symposium USENIX Security 21)*, 2021.

[11] Mihir Bellare, Anand Desai, David Pointcheval, and Phillip Rogaway. Relations among notions of security for public-key encryption schemes. In *Annual International Cryptology Conference*, pages 26–45. Springer, 1998.

[12] Daniel J Bernstein, Tanja Lange, Chloe Martindale, and Lorenz Panny. Quantum circuits for the CSIDH: optimizing quantum evaluation of isogenies. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 409–441. Springer, 2019.

[13] Karthikeyan Bhargavan, Bruno Blanchet, and Nadim Kobeissi. Verified models and reference implementations for the TLS 1.3 standard candidate. In *IEEE Symposium on Security and Privacy (SP 2017)*, pages 483–502. IEEE, 2017.

[14] Jean-François Biasse, Annamaria Iezzi, and Michael J. Jacobson. A note on the security of CSIDH. In *Progress in Cryptology – INDOCRYPT 2018*, pages 153–168, Cham, 2018. Springer International Publishing.

[15] B. Blanchet. A Computationally Sound Mechanized Prover for Security Protocols. *IEEE Transactions on Dependable and Secure Computing*, 5(4):193–207, October 2008. Conference Name: IEEE Transactions on Dependable and Secure Computing.

[16] Bruno Blanchet. Modeling and Verifying Security Protocols with the Applied Pi Calculus and ProVerif. *Foundations and Trends® in Privacy and Security*, 1(1-2):1–135, October 2016. Publisher: Now Publishers, Inc.

[17] Xavier Bonnetain and André Schrottenloher. Quantum security analysis of CSIDH. *Advances in Cryptology–EUROCRYPT 2020*, 12106:493, 2020.

[18] Colin Boyd, Yvonne Cliff, Juan M. Gonzalez Nieto, and Kenneth G. Paterson. One-round key exchange in the standard model. *International Journal of Applied Cryptography*, 1(3):181, 2009.

[19] Mayla Brusò, Konstantinos Chatzikokolakis, and Jerry den Hartog. Formal verification of privacy for RFID systems. In *CSF*, pages 75–88. IEEE Computer Society, 2010.

[20] David Carrel and Dan Harkins. The Internet Key Exchange (IKE). RFC 2409, November 1998.

[21] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. CSIDH: An efficient post-quantum commutative group action. In *Advances in Cryptology – ASIACRYPT 2018*, pages 395–427, Cham, 2018. Springer International Publishing.

[22] Katriel Cohn-Gordon, Cas Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila. A formal security analysis of the Signal messaging protocol. *Journal of Cryptology*, 33(4):1914–1983, 2020.

[23] Hubert Comon, Charlie Jacomme, and Guillaume Scerri. Oracle simulation: a technique for protocol composition with long term shared secrets. In *Proceedings of the 27st ACM Conference on Computer and Communications Security (CCS'20)*, Orlando, USA, November 2020. ACM Press.

[24] Hubert Comon and Adrien Koutsos. Formal Computational Unlinkability Proofs of RFID Protocols. In *Proceedings of the 30th IEEE Computer Security Foundations Symposium (CSF'17)*, pages 100–114, Santa Barbara, California, USA, 2017. IEEE Computer Society Press.

[25] V. Cortier, C. C. Drăgan, F. Dupressoir, B. Schmidt, P. Strub, and B. Warinschi. Machine-Checked Proofs of Privacy for Electronic Voting Protocols. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 993–1008, May 2017. ISSN: 2375-1207.

[26] Cas Cremers, Marko Horvat, Jonathan Hoyland, Sam Scott, and Thyla van der Merwe. A comprehensive symbolic analysis of TLS 1.3. In *ACM SIGSAC Conference on Computer and Communications Security (CCS 2017)*, pages 1773–1788. ACM, 2017.

[27] Cas Cremers, Benjamin Kiesl, and Niklas Medinger. A formal analysis of IEEE 802.11's WPA2: Countering the kracks caused by cracking the counters. In *29th USENIX Security Symposium (USENIX Security 2020)*, pages 1–17. USENIX Association, 2020.

[28] Maria Luisa Dalla Chiara, Roberto Giuntini, and Richard Greechie. *Reasoning in quantum theory: sharp and unsharp quantum logics*, volume 22. Springer Science & Business Media, 2013.

[29] A. Delignat-Lavaud, C. Fournet, M. Kohlweiss, J. Protzenko, A. Rastogi, N. Swamy, S. Zanella-Beguelin, K. Bhargavan, J. Pan, and J. K. Zinzindohoue. Implementing and Proving the TLS 1.3 Record Layer. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 463–482, May 2017. ISSN: 2375-1207.

[30] Antoine Delignat-Lavaud, Cédric Fournet, Bryan Parno, Jonathan Protzenko, Tahina Ramananandro, Jay Bosamiya, Joseph Lallemand, Itsaka Rakotonirina, and Yi Zhou. A Security Model and Fully Verified Implementation for the IETF QUIC Record Layer. page 17.

[31] D. Dolev and A. Yao. On the security of public key protocols. *IEEE Transactions on Information Theory*, 29(2):198–208, 1983.

[32] Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. Crystals-dilithium: A lattice-based digital signature scheme. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pages 238–268, 2018.

[33] Martin Feldhofer, Sandra Dominikus, and Johannes Wolkerstorfer. Strong authentication for RFID systems using the AES algorithm. In *CHES*, volume 3156 of *Lecture Notes in Computer Science*, pages 357–370. Springer, 2004.

[34] Scott Fluhrer, Panos Kampanakis, David McGrew, and Valery Smyslov. Mixing Preshared Keys in the Internet Key Exchange Protocol Version 2 (IKEv2) for Post-quantum Security. RFC 8784, June 2020.

[35] Atsushi Fujioka, Koutarou Suzuki, Keita Xagawa, and Kazuki Yoneyama. Strongly Secure Authenticated Key Exchange from Factoring, Codes, and Lattices. Technical Report 211, 2012.

[36] Tommaso Gagliardoni. Quantum Security of Cryptographic Primitives. *arXiv:1705.02417 [quant-ph]*, May 2017. arXiv: 1705.02417.

[37] Paul Grubbs, Varun Maram, and Kenneth G. Paterson. Anonymous, Robust Post-Quantum Public Key Encryption. 2021. https://ia.cr/2021/708.

[38] Keitaro Hashimoto, Shuichi Katsumata, Kris Kwiatkowski, and Thomas Prest. An efficient and generic construction for Signal's handshake (X3DH): Post-quantum, state leakage secure, and deniable. In *Public-Key Cryptography – PKC 2021*, pages 410–440, Cham, 2021. Springer International Publishing.

[39] Lucca Hirschi, David Baelde, and Stéphanie Delaune. A method for unbounded verification of privacy-type properties. *J. Comput. Secur.*, 27(3):277–342, 2019.

[40] Ari Juels and Stephen A. Weis. Defining strong privacy for RFID. *ACM Trans. Inf. Syst. Secur.*, 13(1):7:1–7:23, 2009.

[41] Charlie Kaufman, Paul E. Hoffman, Yoav Nir, Pasi Eronen, and Tero Kivinen. Internet Key Exchange Protocol Version 2 (IKEv2). RFC 7296, October 2014.

[42] Tomoki Kawashima, Katsuyuki Takashima, Yusuke Aikawa, and Tsuyoshi Takagi. An efficient authenticated key exchange from random self-reducibility on CSIDH. In *International Conference on Information Security and Cryptology*, pages 58–84. Springer, 2020.

[43] Nadim Kobeissi, Karthikeyan Bhargavan, and Bruno Blanchet. Automated verification for secure messaging protocols and their implementations: A symbolic and computational approach. In *IEEE European symposium on security and privacy (EuroS&P 2017)*, pages 435–450. IEEE, 2017.

[44] Nadim Kobeissi, Georgio Nicolas, and Karthikeyan Bhargavan. Noise Explorer: Fully automated modeling and verification for arbitrary Noise protocols. In *IEEE European Symposium on Security and Privacy (EuroS&P 2019)*, pages 356–370. IEEE, 2019.

[45] A. Koutsos. The 5G-AKA Authentication Protocol Privacy. In *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 464–479, June 2019.

[46] Adrien Koutsos. *Preuves symboliques de propriétés d'indistinguabilité calculatoire*. Theses, Université Paris-Saclay, September 2019.

[47] B. Lipp, B. Blanchet, and K. Bhargavan. A Mechanised Cryptographic Proof of the WireGuard Virtual Private Network Protocol. In *2019 IEEE European Symposium on Security and Privacy (EuroS P)*, pages 231–246, June 2019.

[48] Simon Meier, Benedikt Schmidt, Cas Cremers, and David Basin. The TAMARIN Prover for the Symbolic Analysis of Security Protocols. In *Computer Aided Verification*, Lecture Notes in Computer Science, pages 696–701, Berlin, Heidelberg, 2013. Springer.

[49] David Molnar and David A. Wagner. Privacy and security in library RFID: issues, practices, and architectures. In *CCS*, pages 210–219. ACM, 2004.

[50] Waka Nagao, Yoshifumi Manabe, and Tatsuaki Okamoto. On the equivalence of several security notions of key encapsulation mechanism. *IACR Cryptol. ePrint Arch.*, 2006:268, 2006.

[51] Chris Peikert. He gives C-sieves on the CSIDH. *Advances in Cryptology–EUROCRYPT 2020*, 12106:463, 2020.

[52] Phillip Rogaway. Formalizing human ignorance. In *International Conference on Cryptology in Vietnam*, pages 211–228. Springer, 2006.

[53] Guillaume Scerri and Stanley-Oakes Ryan. Analysis of Key Wrapping APIs: Generic Policies, Computational Security. pages 281–295. IEEE Computer Society, June 2016.

[54] Peter Schwabe, Douglas Stebila, and Thom Wiggers. Post-Quantum TLS Without Handshake Signatures. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, CCS '20, pages 1461–1480, New York, NY, USA, October 2020. Association for Computing Machinery.

[55] Peter Schwabe, Douglas Stebila, and Thom Wiggers. Post-quantum TLS without handshake signatures. In *ACM SIGSAC Conference on Computer and Communications Security (CCS 2020)*, pages 1461–1480. ACM, 2020.

[56] Fang Song. A note on quantum security for post-quantum cryptography. In *International Workshop on Post-Quantum Cryptography*, pages 246–265. Springer, 2014.

[57] Nikhil Swamy, Cătălin Hrițcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Simon Forest, Karthikeyan Bhargavan, Cédric Fournet, Pierre-Yves Strub, Markulf Kohlweiss, Jean-Karim Zinzindohoue, and Santiago Zanella-Béguelin. Dependent types and multi-monadic effects in F*. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '16, pages 256–270, New York, NY, USA, January 2016. Association for Computing Machinery.

[58] Dominique Unruh. Quantum proofs of knowledge. In *Annual international conference on the theory and applications of cryptographic techniques*, pages 135–152. Springer, 2012.

[59] Dominique Unruh. Quantum Relational Hoare Logic. *Proceedings of the ACM on Programming Languages*, 3(POPL):1–31, 2019.

[60] John Watrous. Zero-knowledge against quantum attacks. *SIAM Journal on Computing*, 39(1):25–58, 2009.

[61] W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, October 1982.

[62] Mark Zhandry. How to construct quantum random functions. *Journal of the ACM (JACM)*, 68(5):1–43, 2021.