

Message Transmission with Reverse Firewalls— Secure Communication on Corrupted Machines

Yevgeniy Dodis^{1*}, Ilya Mironov², and Noah Stephens-Davidowitz^{1**}

¹ Dept. of Computer Science, New York University.

² Google.

Abstract. Suppose Alice wishes to send a message to Bob privately over an untrusted channel. Cryptographers have developed a whole suite of tools to accomplish this task, with a wide variety of notions of security, setup assumptions, and running times. However, almost all prior work on this topic made a seemingly innocent assumption: that Alice has access to a trusted computer with a proper implementation of the protocol. The Snowden revelations show us that, in fact, powerful adversaries can and will corrupt users' machines in order to compromise their security. And, (presumably) accidental vulnerabilities are regularly found in popular cryptographic software, showing that users cannot even trust implementations that were created honestly. This leads to the following (seemingly absurd) question: “Can Alice securely send a message to Bob even if she cannot trust her own computer?!”

Bellare, Paterson, and Rogaway recently studied this question. They show a strong lower bound that in particular rules out even semantically secure public-key encryption in their model. However, Mironov and Stephens-Davidowitz recently introduced a new framework for solving such problems: reverse firewalls. A secure reverse firewall is a third party that “sits between Alice and the outside world” and modifies her sent and received messages so that *even if the her machine has been corrupted*, Alice's security is still guaranteed. We show how to use reverse firewalls to sidestep the impossibility result of Bellare et al., and we achieve strong security guarantees in this extreme setting.

Indeed, we find a rich structure of solutions that vary in efficiency, security, and setup assumptions, in close analogy with message transmission in the classical setting. Our strongest and most important result shows a protocol that achieves interactive, concurrent CCA-secure message transmission with a reverse firewall—i.e., CCA-secure message transmission on a possibly compromised machine! Surprisingly, this protocol is quite efficient and simple, requiring only four rounds and a small constant number of public-key operations for each party. It could easily be used in practice. Behind this result is a technical composition theorem that shows how key agreement with a sufficiently secure reverse firewall can be used to construct a message-transmission protocol with its own secure reverse firewall.

1 Introduction

We consider perhaps the simplest, most fundamental problem in cryptography: secure message transmission, in which Alice wishes to send a plaintext message to Bob without leaking the plaintext to an eavesdropper. Of course, this problem has a rich history, and it is extremely well-studied with a variety of different setup assumptions and notions of security (e.g., [BDPR98]). There are many beautiful solutions, based on symmetric-key encryption, public-key encryption, key agreement, etc.

However, in the past few years, it has become increasingly clear that the real world presents many vulnerabilities that are not captured by the security models of classical cryptography. The revelations of Edward Snowden show that the United States National Security Agency successfully gained access to secret information by extraordinary means, including subverting cryptographic standards [PLS13, BBG13] and intercepting and tampering with hardware on its way to users [Gre14]. Meanwhile, many

* Partially supported by gifts from VMware Labs and Google, and NSF grants 1319051, 1314568, 1065288, 1017471.

** Partially supported by National Science Foundation under Grant No. CCF-1320188. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

(apparently accidental) security flaws have been found in widely deployed pieces of cryptographic software, leaving users completely exposed [LHA⁺12, CVE14a, CVE14b, CVE14c, Jun15]. Due to the complexity of modern cryptographic software, such vulnerabilities are extremely hard to detect in practice, and, ironically, cryptographic modules are often the easiest to attack, as attackers can often use cryptographic mechanisms to mask their activities or opportunistically hide their communications within encrypted traffic (as in the case of the Heartbleed vulnerability). This has led to a new direction for cryptographers (often called “post-Snowden cryptography”), which in our context is summarized by the following (seemingly absurd) question: “How can Alice and Bob possibly communicate securely when an eavesdropper may have corrupted their computers?!”

Motivated by such concerns, Bellare, Paterson, and Rogaway consider the problem of securely encrypting a message when the encrypting party might be compromised [BPR14a]. They consider the case in which the corrupted party’s behavior is indistinguishable from that of an honest implementation. Even in this setting, their main result is a strong lower bound, showing that even a relatively weak adversary can break any scheme that “non-trivially uses randomness.” (They also provide an elegant deterministic symmetric-key solution, which we use as a subprotocol in the sequel.) In particular, it is easy to see that a semantically secure public-key message transmission is impossible in their framework. (See [BH15] for an analysis of weaker notions of security for public-key encryption in this setting.)

1.1 Reverse Firewalls

Due to the strong lower bound in [BPR14a], we consider a relaxation of their model in which we allow for an additional party, a (*cryptographic*) *reverse firewall* (RF) as recently introduced by Mironov and Stephens-Davidowitz [MS15]. We provide formal definitions in Section 2.1, but since RFs are quite a new concept (and they can be rather confusing at first), we now provide a high-level discussion of some of the salient aspects of the reverse-firewall framework.

A reverse firewall for Alice is an autonomous intermediary that modifies the messages that Alice’s machine sends and receives. The hope is that the protocol equipped with a reverse firewall can provide meaningful security guarantees for Alice *even if her own machine is compromised*. As we explain in detail below, the firewall is *untrusted* in the sense that it shares no secrets with Alice, and in general we expect Alice to place no more trust in the firewall than she does in the communication channel itself.

More concretely, Mironov and Stephens-Davidowitz start by considering an arbitrary cryptographic protocol that satisfies some notions of functionality (i.e., correctness) and security.¹ For example, perhaps the simplest non-trivial case is semantically secure message transmission from Alice to Bob, which has the functionality requirement that Bob should receive the correct plaintext message from Alice and the security requirement that a computationally bounded adversary “should not learn anything about Alice’s plaintext message” from the transcript of a run of the protocol. Formally, we can model this functionality by providing Alice with an input plaintext and requiring Bob’s output to match this, and we can model semantic security by a standard indistinguishability security game.

A reverse firewall for Alice in such a protocol is a party that “sits between” Alice and the outside world and modifies the messages that Alice sends and receives. Such a firewall *maintains functionality* if the resulting protocol achieves the same functionality as the original. E.g., in the

¹ The notion of functionality in [MS15] is quite simple, and it should not be confused with the much more complicated concept of functionality used in the universal composability framework. Formally, Mironov and Stephens-Davidowitz define a functionality requirement as any condition on the output of the parties that may depend on the input, and in practice, these requirements are straightforward.

case of message transmission, Bob should still receive Alice’s message—his output should still match Alice’s input. More interestingly, the firewall *preserves security* if the protocol with the firewall is secure even when we replace Alice’s computer with some arbitrarily corrupted party. For example, a reverse firewall for Alice preserves semantic security of message transmission if a computationally bounded adversary “learns nothing about Alice’s plaintext message” from the transcript of messages sent between the firewall and Bob, *regardless of how Alice behaves*. E.g., the firewall may rerandomize the messages that Alice sends in a way that makes them indistinguishable from random from the adversary’s perspective, regardless of Alice’s original message.

Note that it also makes sense to consider reverse firewalls for the receiver, Bob. For example, consider a protocol in which Bob first sends his public key to Alice, and Alice responds with an encryption of her message under this key. Clearly, if Bob’s computer is corrupted in such a protocol, this can compromise security, even if Alice behaves properly. In such a protocol, a firewall for Bob might rerandomize his public key. Of course, to maintain correctness, this firewall must also intercept Bob’s incoming messages and convert ciphertexts under this rerandomized key to encryptions under Bob’s original key. (We analyze such protocols in a stronger setting in Section 3.)

A key feature of protocols with reverse firewalls, as defined in [MS15], is that they should be functional and secure *both* with the reverse firewall *and* without it. I.e., there should be a well-defined *underlying protocol* between Alice and Bob that satisfies classical functionality and security requirements. This is one important difference between reverse firewalls and some similar models, such as the mediated model [AsV08] and divertible protocols [BD91, OO90, BBS98], and it comes with a number of benefits. ([MS15] contains a thorough comparison of many different related models.) First, it means that these protocols can be implemented and used without worrying about whether reverse firewalls are present—one protocol works regardless; we simply obtain additional security guarantees with an RF.

Second, and more importantly, this definitional choice provides an elegant solution to a natural concern about reverse firewalls: What happens when the firewall itself is corrupted? Of course, if *both* Alice’s own machine and her firewall are compromised, then we cannot possibly hope for security. But, if Alice’s own implementation is correct and the firewall has been corrupted, then we can view the firewall as “part of” the adversary in the firewall-free protocol between Alice and Bob. Since this underlying protocol must itself be secure, it trivially remains secure in the presence of a corrupted firewall.² This is why we can say that the firewall is trusted no more than the communication channel. (In contrast, the divertible protocols framework assumes that the “warden” is trusted, and the mediated model handles security “against the mediator” separately.)

Of course, the advantage of using a firewall comes when Alice’s machine is corrupted but the firewall is implemented correctly, in which case the firewall provides Alice with a security guarantee that she could not have had otherwise. In short, *the firewall can only help*. In fact, we even require firewalls to be “stackable,” so that arbitrarily many firewalls may be deployed, and security is guaranteed as long as *either* (1) Alice’s own machine is uncorrupted; or (2) at least one of these firewalls is implemented correctly.

Finally, it is convenient to identify a class of *functionality-maintaining corruptions*: compromised implementations that are “technically legal” in the sense that they may deviate arbitrarily from the protocol, as long as they do not break its functionality. Some of our reverse firewalls are only secure against this type of corruption. (This model is introduced by [MS15], and the authors call security against unrestricted compromise *strong* security.) We emphasize that, while this restricted class of compromised implementations is not ideal, it is still quite large. In particular, all of the real-world compromises mentioned above fall into this category [PLS13, BBG13, Gre14, Sup15, LHA⁺12,

² Technically, this statement only holds if the underlying protocol is secure against active adversaries.

CVE14b, CVE14a, CVE14c, Jun15], as do essentially all other forms of compromise considered in prior work, such as backdoored PRNGs [DGG⁺15], Algorithm Substitution Attacks [BPR14a], subliminal channels [Sim84], etc. (We discuss functionality-maintaining corruption in our setting in Section 2.4. See [MS14] for a detailed comparison of the general reverse firewall framework with prior work.)

1.2 Our results

In this section, we walk through the results that we obtain in different settings, starting with simpler cases and working our way up to our stronger results. In what follows, Alice is always the sender and Bob is always the receiver of the message. All of our security notions apply to the concurrent setting, in which the adversary may instantiate many runs of the protocol simultaneously.

The symmetric-key setting. In the first and simplest scenario, Alice and Bob have a shared secret key. (See Appendix A.) Quite naturally, Alice may want to use a symmetric-key encryption scheme to communicate with Bob. Using a standard scheme (e.g., AES-CBC) would, however, expose her to a number of “algorithm-substitution attacks” (ASA, i.e., what we call corruption or compromise) described by Bellare, Paterson, and Rogaway [BPR14a], such as IV-replacement or a biased-ciphertext attacks. To defend against ASA, Bellare et al. propose using an elegant solution: a *deterministic* encryption scheme based on either a counter or a nonce. We briefly consider this case, observing that their solution corresponds to a one-round protocol in our model (in which the firewall simply lets messages pass unaltered).

Unfortunately, we show that strong security (i.e., security against corrupted implementations of Alice that are not necessarily functionality-maintaining) is not achievable without using (less efficient) public-key primitives, even in the reverse-firewalls framework. This provides further motivation to study reverse firewalls in the public-key setting.

Rerandomizable encryption. As we mentioned earlier, the simplest non-trivial reverse firewall in the public-key setting uses CPA-secure rerandomizable public-key encryption. In particular, Alice can send her plaintext encrypted under Bob’s public key, and Alice’s reverse firewall can simply rerandomize this ciphertext. We observe that this folklore technique works in our setting. In Section 3, we present a generalization of this idea that does not require any public-key infrastructure, by having Bob send his public key as a first message. Following [MS15], we observe that Bob can have a reverse firewall for such a protocol that rerandomizes his key (and converts Alice’s ciphertext from an encryption under the rerandomized key to an encryption under Bob’s original key). We therefore show a simple two-round protocol with a reverse firewall for each party.

While such protocols are simple and elegant, they have two major drawbacks. First, they are only secure against passive adversaries (an issue to which we will return soon). Second, and arguably more importantly, such protocols require the computation of public-key operations on the entire plaintext. Since plaintexts are often quite long and public-key operations tend to be much slower than symmetric-key operations, it is much faster in practice to use public-key operations to transmit a (relatively short) key for a symmetric-key encryption scheme and then to send the plaintext encrypted under this symmetric key. There are two general methods for transmitting this key in the classical setting: hybrid encryption and key agreement.

Failure of hybrid encryption. Unfortunately, hybrid encryption does not buy us anything in the reverse-firewalls framework. Recall that in a hybrid encryption scheme, Alice selects a uniformly random key rk for a symmetric-key scheme and sends rk encrypted under Bob’s public key together with the encryption of her message under the symmetric-key scheme with key rk . We might naively hope that we can build a reverse firewall for such a scheme by simply applying the “rerandomizing

firewall” to the “public-key part” of Alice’s ciphertext. But, this does not work because of the attack in which a corrupted implementation of Alice chooses a “bad key” rk^* with which to encrypt the message. The “bad key” rk^* might be known to an adversary; might be chosen so that the ciphertext takes a specific form that leaks some information; or might otherwise compromise Alice’s security. So, intuitively, a reverse firewall in such a scheme must be able to rerandomize the key rk , and it therefore must be able to convert an encryption under some key rk into an encryption of the same plaintext under some new key rk' . Unfortunately, we show that any such “key-malleable” symmetric-key encryption scheme implies public-key encryption. Therefore, it cannot be faster than public-key encryption and is useless for our purposes.

Key agreement. Recall that a key-agreement protocol allows Alice and Bob to jointly select a secret key over an insecure channel. Security requires that the resulting key is indistinguishable from random to an eavesdropper. Such a protocol is often used in conjunction with symmetric-key encryption in the classical setting, where it is justified by composition theorems relating the security of the message-transmission protocol to the underlying key-agreement protocol. Indeed, we give an analogous result (Theorem 2) that works in our setting, showing that a carefully designed key-agreement protocol with sufficiently secure reverse firewalls can be combined with symmetric-key encryption to produce an efficient CPA-secure message-transmission protocol with secure reverse firewalls. This motivates the study of key-agreement protocols with secure reverse firewalls.

As a first attempt at constructing such an object, we might try to somehow rerandomize the messages in the well-known Diffie-Hellman key-agreement protocol, in which Alice first sends the message g^a , Bob then sends g^b , and the shared key is g^{ab} . (See Figure 8.) Here, we run into an immediate problem. Since the firewall must maintain correctness, no matter what message A^* the firewall sends to Bob, it must be the case that the final key is A^{*b} , where b is chosen by Bob. But, this allows a corrupt implementation of Bob to influence the key! For example, Bob can repeatedly resample b until, say, the first bit of the key A^{*b} is zero, thus compromising the security. It is easy to see that this problem is not unique to Diffie-Hellman and in fact applies to *any* protocol in which “a party learns what the key is while sending a message that influences the key.”

So, to truly prevent any party from having any control over the final key, we use a three-round protocol in which Bob sends a *commitment of g^b* as his first message. Alice then sends g^a , and Bob then opens his commitment.³ Of course, the commitment scheme that Bob uses must itself be rerandomizable and malleable, so that the firewall can both rerandomize the commitment itself *and* the committed group element. Fortunately, we show that a very simple scheme, a natural variant of the Pedersen commitment, actually suffices.

Since this simple protocol is unauthenticated, it cannot be secure against active adversaries (an important problem that we address next). But, we note that it still has a number of benefits. It is very efficient (roughly as efficient as the protocols currently used in practice!), and it does not require any public-key infrastructure. And, passive security might be sufficient in some settings. For example, powerful adversaries are known to passively gather large amounts of web traffic [Gre14]. Such adversaries can then later exploit vulnerabilities (whether accidental or planted) in widespread cryptographic software to read the private communications of individuals of interest to them [PLS13]. The key-agreement-based protocol described above (together with its RF) can defend against such attacks. But, it would be much better to achieve security against active adversaries.

CCA-security from key agreement. Now, we attempt to construct a reverse firewall that preserves CCA-security (i.e., security against active adversaries that may “feed” Alice and Bob arbitrary adversarial messages and read Bob’s output). In this setting, we again prove a generic

³ We note that the problem that we face here is very similar to the problem of key control, and our solution is similar to solutions used in the key-control literature. See, e.g., [DPSW06].

composition theorem, which shows that it suffices to find a key-agreement protocol with reverse firewalls that satisfy certain security properties. In analogy with the passive setting, after agreeing to a key with Bob, Alice can use symmetric-key encryption to send the actual plaintext message. The resulting protocol is CCA-secure, and Alice’s reverse firewall preserves this security. (See Theorem 4.)

To instantiate this scheme, we must construct a key-agreement protocol that is secure against active adversaries and has a reverse firewall that preserves this security. Unfortunately, many of the elegant techniques used in classical key-agreement protocols (or even protocols that are secure against key control) are useless here. In particular, most key-agreement protocols that achieve security against active adversaries do so by essentially having both parties sign the transcript at the end of the protocol. Intuitively, this allows the parties to know if the adversary has tampered with any messages, so that they will never agree to a key if a man-in-the-middle has modified their messages. But in our setting, we actually *want* the firewall to be able to modify the parties’ messages! We therefore need to somehow find some unique information that the parties can use to confirm that they have agreed to the same key without preventing the firewall from modifying the key. Furthermore, we need the firewall to be able to check these signatures, so that it can block invalid messages. Therefore, our primary technical challenge in this context is to find a protocol with some string that (1) uniquely identifies the key; (2) respects the firewall’s changes to the parties’ messages; and (3) is efficiently computable from the transcript. And, of course, the protocol must be secure against active adversaries, even though it is in some sense “designed to help a man-in-the-middle.”

In spite of these challenges, we construct an efficient protocol with a reverse firewall for each party that preserves security against active adversaries. Remarkably, our protocol achieves this extremely strong notion of security with only four rounds and relatively short messages, and the parties themselves (including the firewall) only need to perform a small constant number of operations. This compares quite favorably with protocols that are currently implemented in practice (which of course are completely insecure in our setting), and we therefore believe that this protocol can and should be implemented and used in the real world.

This surprising solution, which we describe in detail in Section 5.1, uses hashed Diffie-Hellman (similar in spirit to [Kil07]) and bilinear maps. We also use unique signatures to prevent the signatures from becoming a channel themselves.

Rerandomizable encryption and active adversaries. Finally, we return to the question of protocols based on rerandomizable encryption (despite the fact that we showed that such protocols cannot be efficient), but now in the setting of active adversaries. We show how to achieve CCA-security in a single round using rerandomizable RCCA-secure encryption [BDPR98]. (See Section 6.) Indeed, we show that such a primitive is actually *equivalent to* a one-round protocol with a firewall that preserves CCA-security. Such schemes are fairly well-studied, and very elegant solutions exist [Gro04, PR07]. But, our work leads to an interesting open question. Currently known schemes are rerandomizable in the sense that the rerandomization of any valid ciphertext is indistinguishable from a fresh ciphertext, even with access to a decryption oracle. We ask whether these schemes can be made “strongly rerandomizable,” in the sense that the same is true even for *invalid* ciphertexts. (See Section 6 for the formal definition.) We show the weaker notion of rerandomizability is equivalent to protocols with firewalls that are secure against functionality-maintaining corruption, while strong rerandomizability is equivalent to security against arbitrary corruption.

1.3 Related work

Message transmission in the classical setting (i.e., without reverse firewalls) is of course extremely well-studied, and a summary of such work is beyond the scope of this paper. We note, however, that our security definitions for message transmission protocols follow closely Dodis and Fiore [DF14].

There have been many different approaches to cryptography in the presence of compromise. [MS15] contains a thorough discussion of many of these (though they naturally do not mention the many relevant papers that appeared simultaneously with or after their publication, such as [AMV15, BJK15, RTYZ15, DFP15, BH15, DGG⁺15, SFKR15]). In particular, [MS15] contains a detailed comparison of the reverse firewalls-framework with many prior models, showing that RFs generalize much of the prior work on insider attacks and various related notions. Here, we focus on works whose setting or techniques are most similar to our own.

In particular, our work can be viewed as a generalization of Bellare, Paterson, and Rogaway’s work [BPR14a] in a number of directions. We consider multi-round protocols in which the parties might not share secret keys, and we consider arbitrarily compromised adversaries. In order to get around the very strong lower bound in [BPR14a], we use the RF framework of [MS15]. While Bellare et al.’s work stresses the danger of randomness in secure message transmission, our work highlights the benefits of randomness. In particular, our schemes rely heavily on rerandomization by the RF. However, we do use the elegant deterministic encryption scheme of Bellare et al. as part of two of our protocols. (See Appendix F.)

Our work is closely related to Mironov and Stephens-Davidowitz [MS15], which introduces the reverse-firewalls framework. [MS15] demonstrate feasibility of this framework by constructing reverse firewalls for parties participating in Oblivious Transfer and Secure Function Evaluation protocols—very strong cryptographic primitives. The fact that such strong primitives can be made secure in this model is quite surprising and bodes well for the reverse-firewalls framework. However, these protocols are very inefficient and therefore mostly of theoretical interest. And, while the primitives considered in [MS15] have very strong functionality, the security notions that they achieve are rather weak (e.g., security in the semi-honest model). To fulfill the promise of reverse firewalls, we need to consider protocols of more practical importance. We construct much more efficient protocols for widely deployed primitives with strong security guarantees. Naturally, we inherit some of the techniques of [MS15], but we also develop many new ideas.

Bellare and Hoang [BH15] build on [BPR14a] in a different direction, showing how to build deterministic and hedged public-key encryption schemes that are secure against randomness subversion and algorithm substitution attacks. Essentially, they show public-key encryption schemes that are secure even when the sender is compromised, provided that (1) the type of compromise is restricted; and (2) the plaintext itself comes from a high-entropy distribution. These are the first constructions of fully secure hedged public-key encryption in the standard model. These notions of security are much weaker than those that we achieve, but they achieve them without the use of an RF.

Recently, Ateniese, Magri, and Venturi studied reverse firewalls for signature schemes and showed a number of elegant solutions [AMV15]. Their work can be considered as complementary to ours, as we are concerned with privacy, while they consider authentication. We also note that our more advanced key-agreement scheme uses unique signatures, and we implicitly rely on the fact that unique signatures have a reverse firewall, as [AMV15] prove. Indeed, the more general primitive of rerandomizable signatures that they consider would also suffice for our purposes and might be more efficient in practice.

Our frequent use of rerandomization to “sanitize” messages is very similar to much of the prior work on subliminal channels [Sim84, BD91, Des90, Des94, BBS98], divertible protocols [BD91, OO90, BBS98], collusion-free protocols [LMs05, AsV08], etc—particularly the elegant work of Blaze, Bleumer, and Strauss [BBS98] and Alwen, shelat, and Visconti [AsV08]. Again, we refer the reader to [MS15] for a thorough discussion of these models and their relationship to the reverse-firewall framework.

Finally, we note that some of our study of key agreement is similar to work on key-agreement protocols secure against active insiders, and the study of key control (e.g., [PW03,KS05,DPSW06]). These works consider key-agreement protocols involving at least three parties, in which one or more of the participants wishes to maliciously fix the key or otherwise subvert the security of the protocol. Some of the technical challenges that we encounter are similar to those encountered in the key control literature, and indeed, the simple commitment-based protocol that we present in Section 4.1 can be viewed as a simple instantiation of some of the known (more sophisticated) solutions to the key-control problem (e.g., [DPSW06]). However, since prior work approached this problem from a different perspective—with three or more parties and without reverse firewalls—our more technical solutions presented in Section 5.1 are quite different. In particular, almost all prior work on key agreement focuses on creating protocols that produce “non-malleable” keys, whereas our protocols need some type of malleability specifically to allow the firewall to maul the keys! Perhaps surprisingly, we accomplish this without sacrificing security, and our techniques might therefore be of independent interest.

2 Definitions

2.1 Reverse firewalls

We use the definition of reverse firewalls from [MS15] (and we refer the reader to [MS15] for further discussion of the reverse-firewall framework). A reverse firewall \mathcal{W} is a stateful algorithm that maps messages to messages. For a party A and reverse firewall \mathcal{W} , we define $\mathcal{W} \circ A$ as the “composed” party in which \mathcal{W} is applied to the messages that A receives before A “sees them” and the messages that A sends before they “leave the local network of A .” \mathcal{W} has access to all public parameters, but not to the private input of A or the output of A . (We can think of \mathcal{W} as an “active router” that sits at the boundary between Alice’s private network and the outside world and modifies Alice’s incoming and outgoing messages.) We repeat all relevant definitions from [MS15] below, and we add two new ones.

As in [MS15], we assume that a cryptographic protocol comes with some functionality or correctness requirements \mathcal{F} and security requirements \mathcal{S} . (For example, a functionality requirement \mathcal{F} might require that Alice and Bob output the same thing at the end of the protocol. A security requirement \mathcal{S} might ask that no efficient adversary can distinguish between the transcript of the protocol and a uniformly random string.) Throughout, we use \tilde{A} to represent arbitrary adversarial implementations of party A and \tilde{A} to represent functionality-maintaining implementations of A (i.e., implementations of A that still satisfy the functionality requirements of the protocol). For a protocol \mathcal{P} with party A , we write $\mathcal{P}_{A \rightarrow \tilde{A}}$ to represent the protocol in which the role of party A is replaced by party \tilde{A} .

We are only interested in firewalls that themselves maintain functionality. In other words, the *composed party* $\mathcal{W} \circ A$ should not break the correctness of the protocol. (Equivalently, $\mathcal{P}_{A \rightarrow \mathcal{W} \circ A}$ should satisfy the same functionality requirements as the underlying protocol \mathcal{P} .) We follow [MS15] in requiring something slightly stronger—reverse firewalls should be “stackable”, so that many reverse firewalls composed in series $\mathcal{W} \circ \dots \circ \mathcal{W} \circ A$ still do not break correctness. All of our firewalls will trivially satisfy this notion. (Indeed, the messages sent by our firewalls will be distributed identically to those of an honest party.) Note as well that we are not interested in protocols whose functionality “depends on the presence of the reverse firewall,” so we require that the protocol without the reverse firewall is also functional.

Definition 1 (Reverse firewall). *A reverse firewall \mathcal{W} maintains functionality \mathcal{F} for party A in protocol \mathcal{P} if protocol \mathcal{P} satisfies \mathcal{F} , the protocol $\mathcal{P}_{A \rightarrow \mathcal{W} \circ A}$ satisfies \mathcal{F} , and the protocol $\mathcal{P}_{A \rightarrow \mathcal{W} \circ \dots \circ \mathcal{W} \circ A}$*

also satisfies \mathcal{F} . (I.e., we can compose arbitrarily many reverse firewalls without breaking functionality.)

Of course, a firewall is not interesting unless it provides some benefit. The most natural reason to deploy a reverse firewall is to *preserve* the security of a protocol, even in the presence of compromise. The below definition (which again follows [MS14]) captures this notion by asking that the protocol obtained by replacing party A with $\mathcal{W} \circ \tilde{A}$ for an arbitrary corrupted party \tilde{A} still achieves some notion of security. For example, when we consider message transmission, we will want the firewall to guarantee Alice’s privacy against some adversary, even when Alice’s own computer has been corrupted. (As before, we are only interested in protocols that are secure without the reverse firewall as well.)

Definition 2 (Security preservation). *A reverse firewall strongly preserves security \mathcal{S} for party A in protocol \mathcal{P} if protocol \mathcal{P} satisfies \mathcal{S} , and for any polynomial-time algorithm \tilde{A} , the protocol $\mathcal{P}_{A \rightarrow \mathcal{W} \circ \tilde{A}}$ satisfies \mathcal{S} . (I.e., the firewall can guarantee security even when an adversary has tampered with A .)*

A reverse firewall preserves security \mathcal{S} for party A in protocol \mathcal{P} satisfying functionality requirements \mathcal{F} if protocol \mathcal{P} satisfies \mathcal{S} , and for any polynomial-time algorithm \tilde{A} such that $\mathcal{P}_{A \rightarrow \tilde{A}}$ satisfies \mathcal{F} , the protocol $\mathcal{P}_{A \rightarrow \mathcal{W} \circ \tilde{A}}$ satisfies \mathcal{S} . (I.e., the firewall can guarantee security even when an adversary has tampered with A , provided that the tampered implementation does not break the functionality of the protocol.)

For technical reasons, we will also need a new definition not present in [MS14]. We wish to show generic composition theorems, allowing us to construct a message-transmission protocol with secure reverse firewall from any key-agreement protocol with its own firewalls. In order to accomplish this, we will need the notion of *detectable failure*. Essentially, a protocol fails detectably if we can distinguish between transcripts of valid runs of the protocol and invalid transcripts. We will use this to make sure that a firewall of a larger protocol can test whether a subprotocol has failed. We make this precise below. In order to do so, we need to carefully consider what it means for a transcript to be valid. For simplicity, we assume that honest parties always output \perp when they receive a malformed message (e.g., when a message that should be a pair of group elements is not a pair of group elements). While the general notion of validity is a bit technical, we will use it in very straightforward ways. (E.g., transcripts will be valid if and only if a commitment is properly opened and a certain signature is valid.)

Definition 3 (Valid transcripts). *A sequence of bits r and private input I generate transcript \mathcal{T} in protocol \mathcal{P} if a run of the protocol \mathcal{P} with input I in which the parties’ coin flips are taken from r results in the transcript \mathcal{T} . A transcript \mathcal{T} is a valid transcript for protocol \mathcal{P} if there is a sequence r and private input I generating \mathcal{T} such that no party outputs \perp at the end of the run. (Here, we assume that the public input is part of the transcript.) A protocol has unambiguous transcripts if for any valid transcript \mathcal{T} , there is no possible input I and coins r generating \mathcal{T} that results in a party outputting \perp . (In other words, a valid transcript never results from a failed run of the protocol.)*

Definition 4 (Detectable failure). *A reverse firewall \mathcal{W} detects failure for party A in protocol \mathcal{P} if*

- if $\mathcal{P}_{A \rightarrow \mathcal{W} \circ A}$ has unambiguous transcripts;
- the firewall \mathcal{W} outputs the special symbol \perp when run on any transcript that is not valid for $\mathcal{P}_{A \rightarrow \mathcal{W} \circ A}$; and

- there is a polynomial-time deterministic algorithm that decides whether a transcript \mathcal{T} is valid for $\mathcal{P}_{A \rightarrow \mathcal{W} \circ A}$.

We will also need the notion of *exfiltration resistance*, introduced in [MS15]. Intuitively, a reverse firewall is exfiltration resistant if “no corrupt implementation of Alice can leak information through the firewall.” We say that it is exfiltration resistant for Alice against Bob if Alice cannot leak information to Bob through the firewall, and we say that it is exfiltration resistant against eavesdroppers (or just exfiltration resistant) if Alice cannot leak information through the firewall to an adversary that is only given access to the protocol transcript.

The relationship between security preservation and exfiltration resistance depends on the security notion of the cryptographic primitive in question. E.g., in a message-transmission protocol, a reverse firewall for Alice resists exfiltration if and only if it preserves semantic security. However, it is possible to construct a reverse firewall for a key-agreement protocol that preserves security but does not resist exfiltration (for example, we can add an arbitrary message to the beginning of any key-agreement protocol without changing its security properties, but clearly such a message can be used to leak information). It is also possible to construct a firewall that resists exfiltration but does not preserve key-agreement security. The second definition below (which uses the notion of valid transcripts) is new to this paper and is necessary for our composition theorems.

```

proc. LEAK( $\mathcal{P}, A, B, \mathcal{W}, \lambda$ )
 $(\tilde{A}, \tilde{B}, I) \leftarrow \mathcal{E}(1^\lambda)$ 
 $b \xleftarrow{\$} \{0, 1\}$ 
IF  $b = 1$ ,  $A^* \leftarrow \mathcal{W} \circ \tilde{A}$ 
ELSE,  $A^* \leftarrow \mathcal{W} \circ A$ 
 $\mathcal{T}^* \leftarrow \mathcal{P}_{A \rightarrow A^*, B \rightarrow \tilde{B}}(I)$ 
 $b^* \leftarrow \mathcal{E}(\mathcal{T}^*, S_{\tilde{B}})$ 
OUTPUT  $(b = b^*)$ 

```

Fig. 1: $\text{LEAK}(\mathcal{P}, A, B, \mathcal{W}, \lambda)$, the exfiltration resistance security game for a reverse firewall \mathcal{W} for party A in protocol \mathcal{P} against party B with input I . \mathcal{E} is the adversary, λ the security parameter, $S_{\tilde{B}}$ the state of \tilde{B} after the run of the protocol, I valid input for \mathcal{P} , and \mathcal{T}^* is the transcript resulting from a run of the protocol $\mathcal{P}_{A \rightarrow A^*, B \rightarrow \tilde{B}}$ with input I .

Definition 5 (Exfiltration resistance). A reverse firewall is exfiltration resistant for party A against party B in protocol \mathcal{P} satisfying functionality \mathcal{F} if no PPT algorithm \mathcal{E} with output circuits \tilde{A} and \tilde{B} such that $\mathcal{P}_{A \rightarrow \tilde{A}}$ and $\mathcal{P}_{B \rightarrow \tilde{B}}$ satisfy \mathcal{F} has non-negligible advantage in $\text{LEAK}(\mathcal{P}, A, B, \mathcal{W}, \lambda)$. If B is empty, then we simply say that the firewall is exfiltration resistant.

A reverse firewall is exfiltration resistant for party A against party B in protocol \mathcal{P} with valid transcripts if no PPT algorithm \mathcal{E} with output circuits \tilde{A} and \tilde{B} such that $\mathcal{P}_{A \rightarrow \tilde{A}}$ and $\mathcal{P}_{B \rightarrow \tilde{B}}$ produce valid transcripts for \mathcal{P} has non-negligible advantage in $\text{LEAK}(\mathcal{P}, A, B, \mathcal{W}, \lambda)$. If B is empty, then we simply say that the firewall is exfiltration resistant with valid transcripts.

A reverse firewall is strongly exfiltration resistant for party A against party B in protocol \mathcal{P} if no PPT adversary \mathcal{E} has non-negligible advantage in $\text{LEAK}(\mathcal{P}, A, B, \mathcal{W}, \lambda)$. If B is empty, then we say that the firewall is strongly exfiltration resistant.

2.2 Message-transmission protocols

A *message-transmission protocol* is a two-party protocol in which one party, Alice, is able to communicate a plaintext message to the other party, Bob. (For simplicity, we only formally model the case in which Alice wishes to send a single plaintext to Bob per run of the protocol, but this of course

naturally extends to a more general case in which Alice and Bob wish to exchange many plaintext messages.) We consider two notions of security for such messages. First, we consider *CPA security*, in which the adversary must distinguish between the transcript of a run of the protocol in which Alice communicates the plaintext m_0 to Bob and the transcript with which Alice communicates m_1 to Bob, where m_0 and m_1 are adversarially chosen plaintexts. (Even in this setting, we allow the adversary to start many concurrent runs of the protocol with adaptively chosen plaintexts.) Our strongest notion of security is *CCA security* in which the adversary may “feed” the parties any messages and has access to a decryption oracle. Our security definitions are similar in spirit to [DF14], but adapted for our setting.

Session ids. Throughout this paper, we consider protocols that may be run concurrently many times between the same two parties. In order to distinguish one run of a protocol from another, we therefore “label” each run with a unique session id, denoted sid . We view sid as an implicit part of every message, and we often ignore sid when it is not important. Our parties and firewalls are stateful, and we assume that the parties and the firewall maintain a list of the relevant session ids, together with any information that is relevant to continue the run of the protocol corresponding to sid (such as the number of messages sent so far, any values that need to be used later in the protocol, etc.). We typically suppress explicit reference to these states. In our security games, the adversary may choose the value sid for each run of the protocol, provided that each party has a unique run for each session sid . (In fact, it does not even make sense for the adversary to use the same sid for two different runs of the protocol with the same party, as this party will necessarily view any calls with the same sid as corresponding to a single run of the protocol. However, as is clear from our security games, an active adversary may maintain two separate runs of a protocol with two different parties but the same sid .) In practice, sid can be a simple counter or any other nonce (perhaps together with any practical information necessary for communication, such as IP addresses). We note in passing that, in the setting of reverse firewalls, a counter is preferable to, e.g., a random nonce to avoid providing a channel through sid , but such concerns are outside of our model and the scope of this paper.

The definition below makes the above formal and provides us with some useful terminology.

Definition 6 (Message-transmission protocol). *A message-transmission protocol is a two-party protocol in which one party, Alice, receives as input a plaintext m from some plaintext space \mathcal{M} . The protocol is correct if for any input $m \in \mathcal{M}$, Bob’s output is always m .*

We represent the protocol by four algorithms $\mathcal{P} = (\text{setup}, \text{next}_A, \text{next}_B, \text{return}_B)$. setup takes as input 1^λ , where λ is the security parameter, and returns the starting states for each party, S_A, S_B , which consist of both private input, σ_A and σ_B respectively, and public input π . Each party’s next procedure is a stateful algorithm that takes as input sid and an incoming message, updates the party’s state, and returns an outgoing message. The return_B procedure takes as input Bob’s state S_B and sid and returns Bob’s final output.

We say that a message-transmission protocol is

- unkeyed if setup does not return any private input σ_A or σ_B ;
- singly keyed if setup returns private input σ_B for Bob but none for Alice;
- publicly keyed if setup returns private input for both parties σ_A and σ_B , but these private inputs are independently distributed; and
- privately keyed if setup returns private input for both parties whose distributions are dependent.

When we present protocols, we will often omit the formality of defining explicit functions $\mathcal{P} = (\text{setup}, \text{next}_A, \text{next}_B, \text{return}_B)$ and states for the parties, preferring instead to use diagrams as in Figure 4. But, this formulation is convenient for our security definitions. In particular, we present

the relevant subprocedures for our security games in Figure 2. An adversary plays the game depicted in Figure 2 by first calling `initialize` (receiving as output π) and then making various calls to the other subprocedures. Each time it calls a subprocedure, it receives any output from the procedure. The game ends when the adversary calls `finalize`, and the adversary wins if and only if the output of `finalize` is one.

<pre> proc. initialize(1^λ) $(\sigma_A, \sigma_B, \pi) \xleftarrow{\\$} \text{setup}(1^\lambda)$ $S_A \leftarrow (\sigma_A, \pi); S_B \leftarrow (\sigma_B, \pi)$ $\text{sid}^* \leftarrow \perp; \text{compromised} \leftarrow \text{false}$ $b \xleftarrow{\\$} \{0, 1\}$ OUTPUT π proc. finalize(b^*) IF $b = b^*$, RETURN 1 ELSE, RETURN 0 proc. start-run(sid, m) IF $\text{sid} \notin S_A$, $S_A.\text{add}(\text{sid}, m)$ proc. start-challenge(sid, m_0, m_1) IF $\text{sid} \notin S_A$ AND $\text{sid}^* = \perp$, $\text{sid}^* \leftarrow \text{sid}$ $S_A.\text{add}(\text{sid}, m_b)$ </pre>	<pre> proc. get-next_A(sid, M) IF compromised, OUTPUT \perp OUTPUT next_A(S_A, sid, M) proc. get-next_B(sid, M) IF compromised, OUTPUT \perp OUTPUT next_B(S_B, sid, M) proc. get-output_B(sid) IF $\text{sid} = \text{sid}^*$ OR compromised, OUTPUT \perp OUTPUT return_B(S_B, sid) proc. get-secrets compromised $\leftarrow \text{true}$ OUTPUT (σ_A, σ_B) </pre>
---	---

Fig. 2: Procedures used to define security for message-transmission protocol $\mathcal{P} = (\text{setup}, \text{next}_A, \text{next}_B)$. An adversary plays this game by first calling `initialize` and then making various oracle calls. The game ends when the adversary calls `finalize`, and the output of `finalize` is one if the adversary wins and zero otherwise.

The below definitions capture formally the intuitive notions of security that we presented above. In particular, the CPA security definition allows the adversary to start arbitrarily many concurrent runs of the protocol with adversarial input, but it does not allow the adversary to change the messages sent by the two parties or to send its own messages. We also define forward secrecy, which requires that security hold even if the parties' secret keys may be leaked to the adversary.

Definition 7 (Message-transmission security). *A message-transmission protocol is called*

- chosen-plaintext secure (*CPA-secure*) if no PPT adversary has non-negligible advantage in the game presented in Figure 2 when `get-nextA(sid, M)` and `get-nextB(sid, M)` output \perp unless this is the first `get-next` call with this `sid` or M is the output from the previous `get-nextA` call with the same `sid` or the previous `get-nextB` with the same `sid` respectively (i.e., the adversary is passive); and
- chosen-ciphertext secure (*CCA-secure*) if no PPT adversary has non-negligible advantage in the game presented in Figure 2 with access to all oracles.

We say that the protocol is chosen-plaintext (resp. chosen-ciphertext) secure without forward secrecy if the above holds without access to the `get-secrets` oracle.

We note that it does not make sense to consider chosen-ciphertext security when Bob may be corrupted. In this case, the output of `get-outputB` could be arbitrary. (Note that the firewall can potentially “sanitize” Bob’s messages, but it of course does not have access to his output.) We therefore only consider firewalls that preserve CPA security for Bob.

2.3 Key agreement

Key-agreement protocols will play a central role in our constructions, so we now provide a definition of key agreement that suffices for our purposes. Our notion of key agreement closely mirrors the definitions from the previous section.

Definition 8 (Key agreement). A key-agreement protocol is represented by five algorithms, $\mathcal{P} = (\text{setup}, \text{next}_A, \text{next}_B, \text{return}_A, \text{return}_B)$. setup takes as input 1^λ , where λ is the security parameter and returns the starting states for each party, S_A, S_B , which consists of public input π and the private input for each party σ_A and σ_B . Each party's next procedure is a stateful algorithm that takes as input sid and an incoming message, updates the party's state, and returns an outgoing message. Each party's return procedure takes as input the relevant party's state and sid and returns the party's final output from some key space \mathcal{K} or \perp . We also allow auxiliary input aux to be added to Alice's state before the first message of a protocol is sent.

The protocol is correct if Alice and Bob always output the same thing at the end of the run of a protocol for any random coins and auxiliary input aux .

We say that a key-agreement protocol is

- unkeyed if setup does not return any private input σ_A or σ_B ;
- singly keyed if setup returns private input σ_B for Bob but no private input σ_A for Alice; and
- publicly keyed if setup returns private input for both parties σ_A and σ_B .

<pre> proc. initialize(1^λ) (σ_A, σ_B, π) $\stackrel{\\$}{\leftarrow}$ setup(1^λ) $S_A \leftarrow (\sigma_A, \pi)$ $S_B \leftarrow (\sigma_B, \pi)$ $\text{sid}^* \leftarrow \perp$ $\text{compromised} \leftarrow \text{false}$ $b \stackrel{\\$}{\leftarrow} \{0, 1\}$ OUTPUT π proc. finalize(b^*) IF $b = b^*$, RETURN 1 ELSE, RETURN 0 </pre>	<pre> proc. get-next$_A$(sid, M) IF NOT compromised, OUTPUT next$_A$(S_A, sid, M) proc. get-next$_B$(sid, M) IF NOT compromised, OUTPUT next$_B$(S_B, sid, M) proc. get-output$_A$(sid) IF compromised, OUTPUT \perp IF $\text{sid} = \text{sid}^*$ AND $b = 0$, IF return$_A$(S_A, sid) = \perp, OUTPUT \perp ELSE, OUTPUT R_{sid} ELSE, OUTPUT return$_A$(S_A, sid) proc. get-output$_B$(sid) IF compromised, OUTPUT \perp IF $\text{sid} = \text{sid}^*$ AND $b = 0$, IF return$_B$(S_B, sid) = \perp, OUTPUT \perp ELSE, OUTPUT R_{sid} ELSE, OUTPUT return$_B$(S_B, sid) proc. get-secrets compromised \leftarrow true OUTPUT (σ_A, σ_B) </pre>
<pre> proc. start-run(sid, aux) IF $\text{sid} \notin S_A$, $S_A.\text{add}(\text{sid}, \text{aux})$ proc. start-challenge(sid, aux) IF $\text{sid}^* = \perp$ AND $\text{sid} \notin S_A$, $\text{sid}^* \leftarrow \text{sid}$ $\mathcal{R}_{\text{sid}^*} \stackrel{\\$}{\leftarrow} \mathcal{K}$ $S_A.\text{add}(\text{sid}, \text{aux})$ </pre>	

Fig. 3: Procedures used to define security for key-agreement protocol $\mathcal{P} = (\text{setup}, \text{next}_A, \text{next}_B, \text{return}_A, \text{return}_B)$. An adversary plays this game by first calling initialize and then making various oracle calls. The game ends when the adversary calls finalize, and the output of finalize is one if the adversary wins and zero otherwise. We suppress the auxiliary input aux when it is irrelevant.

Definition 9 (Key-agreement security). A key-agreement protocol is

- secure against passive adversaries if no probabilistic polynomial-time adversary has non-negligible advantage in the game presented in Figure 3 when $\text{get-next}_A(\text{sid}, M)$ and $\text{get-next}_B(\text{sid}, M)$ output \perp unless this is the first get-next call with this sid or M is the output from the previous get-next_B call with the same sid or the previous get-next_A call with the same sid respectively (i.e., the adversary is passive);
- secure against active adversaries for Alice if no probabilistic polynomial-time algorithm has non-negligible advantage in the game presented in Figure 3 without access to the get-output_A oracle;
- secure against active adversaries for Bob if no probabilistic polynomial-time algorithm has non-negligible advantage in the game presented in Figure 3 without access to the get-output_B oracle; and
- secure against active adversaries if it is secure against active adversaries for both Bob and Alice; and
- authenticated for Bob if no probabilistic polynomial-time algorithm playing the game presented in Figure 3 can output a valid transcript with corresponding session id sid unless $\text{return}_B(S_B, \text{sid}) \neq \perp$ or $\text{compromised} = \text{true}$. (I.e., it is hard to find a valid transcript unless Bob returns a key.) Furthermore, if the transcript is valid and $\text{get-output}_A(\text{sid}) \neq \perp$ then $\text{get-output}_A(\text{sid}) = \text{return}_B(\text{sid})$. (I.e., if the transcript is valid and Alice outputs a key, then Bob outputs the same key.)

Note that these definitions are far from standard. In particular, in the case of active adversaries, we define security for Alice in terms of the keys that *Bob* outputs and security for Bob in terms of the keys that *Alice* outputs. This may seem quite counterintuitive. But, in our setting, we are worried that Alice may be corrupted. In this case, we cannot hope to restrict Alice’s output after she receives invalid messages. (The firewall can modify Alice’s *messages*, but not her *output*.) So, the best we can hope for is that the firewall prevents a tampered implementation of Alice (together with an active adversary) from “tricking” Bob into returning an insecure key.

2.4 Reverse firewalls and message transmission

Now that we have formal definitions, we discuss the relationship between message-transmission protocols, reverse firewalls, and functionality-maintaining adversaries. We start with an obvious observation: a firewall cannot allow Alice to non-trivially communicate with Bob while simultaneously preventing an arbitrarily corrupted implementation of Alice from communicating something nefarious to Bob! In the language of [MS15] and Section 2.1, a reverse firewall for Alice in a message-transmission protocol cannot *strongly* resist exfiltration against Bob. Such basic impossibility results are bound to arise when we consider notions of security for parties that may be corrupted!

However, if we only consider corrupted implementations of Alice that *maintain the functionality* of this message-transmission protocol, then it is technically possible for a firewall to resist exfiltration against Bob. ([MS15] refer to this as exfiltration resistance against Bob, but not *strong* exfiltration resistance.) The functionality of a message-transmission protocol requires Bob to output (i.e., “receive”) the plaintext that Alice is given as input. So, a functionality-maintaining corruption of Alice must “send” *this* message. Intuitively, a firewall that resists exfiltration against Bob for this restricted class of corruptions “allows Alice to send her message to Bob, but nothing else.” In other words, these firewalls prevent steganography. All of our reverse firewalls easily achieve this, and in some sense this is the “best possible” notion of exfiltration resistance against Bob.

Of course, for message-transmission protocols, we are primarily concerned with security against a third party adversary. We say that a firewall *preserves security* if it protects Alice from such an

adversary in some way.⁴ Ideally, we would like to build reverse firewalls for Alice that protect her from such an adversary even if Alice has been arbitrarily corrupted.

In order to do this, we have to worry about an “attack by refusal.” E.g., a corrupted implementation of Alice might simply refuse to communicate at all (in our model, she can send the special “abort” symbol \perp) when the first bit of her bank account number is zero and communicates normally otherwise. Clearly, a firewall can only protect against such attacks if it can “spoof” messages from Alice, so that a third party adversary cannot tell the difference between a legitimate message sent from Alice and a “spoofed” message from the firewall. Now, suppose Alice has a secret key that is used to publicly sign something at some point in the protocol. It should be clear that the strong notion of security preservation discussed above (which [MS15] simply call strong security preservation) is not possible for any such protocols. So, our firewalls that work in the setting in which Alice has a public-key/private-key pair preserve Alice’s security, but they do not *strongly* preserve it. (See Section 5.) In the “mixed” setting in which Bob has a public-key/private-key pair but Alice does not, we show that strong security preservation is equivalent to a strengthened notion of rerandomizable RCCA-secure encryption, a primitive that is currently being actively studied. (See Section 6.) In the unkeyed setting, we do in fact achieve strong security with a simple protocol. (See Section 3.)

Finally, while the barrier discussed above seems inherent, one might hope that this “attack by refusal” is in some sense the only attack possible. We could try to formalize a security notion that captures this, but it is clear that our more advanced protocols would not meet it. In particular, our more advanced protocols use the symmetric-key protocol from [BPR14a] as a subprotocol. The protocol from [BPR14a] is completely deterministic and suffers from an attack in which a corrupted implementation of Alice can simply try to encrypt many different messages until she finds one whose corresponding ciphertext has a specific form (say, one that starts with many zeros). Obviously, this channel can be used to communicate with an eavesdropper, and both of our key-agreement-based protocols inherit this weakness, including our strongest result from Section 5. (We note, however, that this attack only works against our protocols if the corrupted implementation chooses which message to send *after* the key agreement is complete.) We leave it to future work to find a stronger notion of security that is still achievable (but, we stress that we already achieve a remarkably strong notion of security with surprisingly efficient protocols). It seems that something might be possible here, but we might need to sacrifice efficiency by applying (typically very slow) public-key operations to (potentially very long) plaintexts.

3 A two-round protocol from rerandomizable encryption

We first consider the simple case of CPA-secure two-round schemes in which the first message is a public key chosen randomly by Bob and the second message is an encryption of Alice’s plaintext under this public key. Figure 4 shows the protocol. (Note that this also captures the case of a one-round protocol in which Bob has a fixed public key. We consider this slightly stronger setting to kill two birds with one stone.)

In order to provide a reverse firewall for Alice in this protocol, the encryption scheme must be *rerandomizable*. In order to provide a reverse firewall for Bob, the scheme must be *key malleable*. Intuitively, a scheme is key malleable if a third party can “rerandomize” a public key and map ciphertexts under the “rerandomized” public key to ciphertexts under the original public key. For completeness, we include formal definitions below.

⁴ It is easy to see that a firewall for Alice in a message-transmission protocol is exfiltration resistant against an eavesdropper if and only if it preserves semantic security—the weakest notion of security that interests us. So, we may safely ignore exfiltration resistance against eavesdroppers in this context.




Fig. 4: Two round message-transmission protocol using the public-key encryption scheme (KeyGen, Enc, Dec).

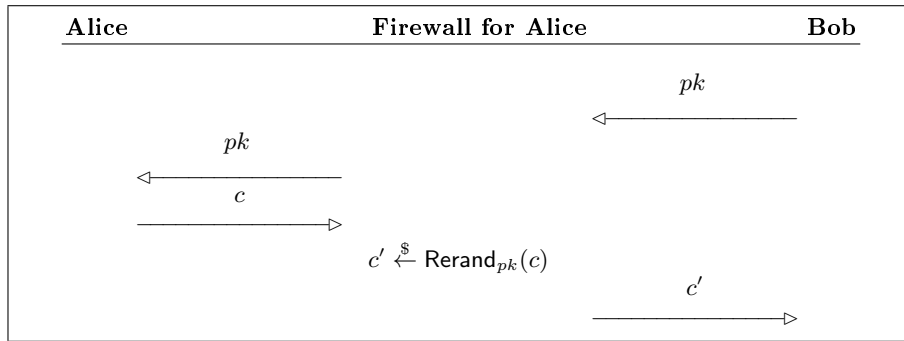


Fig. 5: Reverse firewall for Alice for the protocol shown in Figure 5 that works if the encryption scheme is rerandomizable.

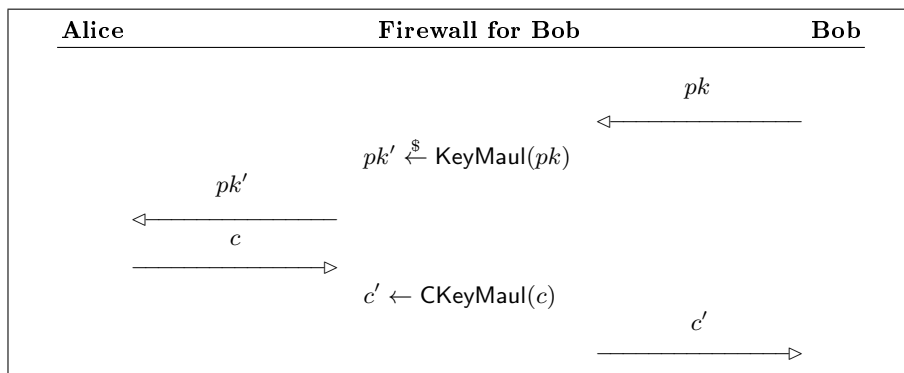


Fig. 6: Reverse firewall for Bob for the protocol shown in Figure 4 that works if the encryption scheme is key malleable. We suppress the randomness r used as input to KeyMaul and CKeyMaul.

Definition 10 (Public-key encryption). A public-key encryption scheme is a triple of efficient algorithms $(\text{KeyGen}, \text{Enc}, \text{Dec})$. KeyGen takes as input 1^λ where λ is the security parameter and outputs a public-key/private-key pair, (pk, sk) . Enc takes as input the public key and a plaintext m from some plaintext space \mathcal{M} and outputs a ciphertext c from some ciphertext space \mathcal{C} . Dec takes as input a ciphertext and the private key and outputs a plaintext or the special symbol \perp . We sometimes omit the keys from the input to Enc and Dec and the security parameter input to KeyGen . The scheme is correct if $\text{Dec}(\text{Enc}(m)) = m$ for all $m \in \mathcal{M}$. The scheme is semantically secure if for any adversarially chosen pair of plaintexts (m_0, m_1) , $\text{Enc}(m_0)$ is computationally indistinguishable from $\text{Enc}(m_1)$.

Definition 11 (Rerandomizable encryption). A semantically secure public-key encryption scheme is rerandomizable if there is an efficient algorithm Rerand (with access to the public key) such that for any ciphertext c such that $\text{Dec}(c) \neq \perp$, we have $\text{Rerand}(\text{Dec}(c)) = \text{Dec}(c)$, and the pair $(c, \text{Rerand}(c))$ is computationally indistinguishable from $(c, \text{Rerand}(\text{Enc}(0)))$. We say that it is strongly rerandomizable if the previous property holds even when $\text{Dec}(c) = \perp$.

Definition 12 (Key malleability). A public-key encryption scheme is key malleable if (1) the output of KeyGen is distributed uniformly over the space of valid keys; (2) for each public key pk there is a unique associated private key sk ; and (3) there is a pair of efficient algorithms KeyMaul and CKeyMaul that behave as follows. KeyMaul is a randomized algorithm that takes as input a public key pk and returns a new public key pk' whose distribution is uniformly random over the public key space and independent of pk . Let (sk, pk) be a private key/public key pair. Let (sk', pk') be the unique pair associated with randomness r such that $pk' = \text{KeyMaul}(pk; r)$. Then, CKeyMaul takes as input a ciphertext c and randomness r and returns c' such that $\text{Dec}_{sk'}(c) = \text{Dec}_{sk}(c')$. We suppress the input r when it is understood. Furthermore, we require that KeyMaul outputs a uniformly random key pk' if called on input that is not in the public-key space.

Example. It is well-known that ElGamal encryption [ElG85] is both key malleable and strongly rerandomizable. In particular, given an ElGamal public-key (g, h) over a group of order p , a ciphertext (u, v) can be rerandomized by applying the operation $(u, v) \rightarrow (g^r u, h^r v)$ where $r \stackrel{\$}{\leftarrow} \mathbb{Z}_p^*$ is chosen uniformly at random. The public key can be mauled by applying the operation $(g, h) \rightarrow (g^\alpha, h^\beta)$ where $(\alpha, \beta) \stackrel{\$}{\leftarrow} (\mathbb{Z}_p^*)^2$ are chosen uniformly and independently at random. Finally, a ciphertext (u, v) under key (g^α, h^β) can be converted into a ciphertext under (g, h) by applying the operation $(u, v) \rightarrow (u^{\beta/\alpha}, v)$.

If the underlying encryption scheme in Figure 4 is rerandomizable, then we can build a reverse firewall for Alice as in Figure 5. If it is key malleable, then we can build a reverse firewall for Bob as in Figure 6. The following theorem shows that this protocol and its reverse firewalls are secure.

Theorem 1. *The unkeyed message-transmission protocol shown in Figure 4 is CPA-secure if the underlying encryption scheme is semantically secure. If the scheme is also (strongly) rerandomizable, then the reverse firewall shown in Figure 5 (strongly) preserves security for Alice and (strongly) resists exfiltration for Alice. If the scheme is key malleable, then the reverse firewall shown in Figure 6 maintains functionality for Bob, strongly preserves Bob's security, and strongly resists exfiltration for Bob against Alice.*

Proof. It is a common folklore result that the underlying protocol (i.e., the protocol without reverse firewalls) is CPA-secure. It is clear that the two firewalls maintain functionality.

Security and exfiltration resistance of Bob's firewall follows from the definition of key malleability. In particular, for any tampered implementation \bar{B} of Bob, after the post-processing by the reverse

firewall, the key pk' is uniformly random by the definition of key malleability, regardless of the behavior of \bar{B} . This implies exfiltration resistance, and security then follows from the fact that the underlying protocol is secure when the key is chosen legitimately.

Consider a tampered implementation of Alice \tilde{A} that maintains functionality, and let c be the output of Alice. Since Alice maintains functionality, $\text{Dec}(c) \neq \perp$. So, by the definition of rerandomizability, the output of $\text{Rerand}(c)$ is indistinguishable from $\text{Enc}_{pk}(0)$. The security preservation and exfiltration resistance of Alice’s firewall follows. A similar argument shows that strong rerandomizability implies strong security preservation and strong exfiltration resistance. \square

3.1 Hybrid encryption fails.

A major drawback of the above scheme is that it requires public-key operations of potentially very long plaintexts, which can be very inefficient in practice. A common solution in the classical setting is to use *hybrid encryption*, in which $\text{Enc}_{pk}(m)$ is replaced by $(\text{Enc}_{pk}(rk), \text{SEnc}_{rk}(m))$, where SEnc is some suitable symmetric-key encryption scheme and rk is a freshly chosen uniformly random key for SEnc . However, if we simply replace the public-key encryption in Figure 4 with the corresponding hybrid-key encryption scheme, then this fails spectacularly. For example, a tampered version of Alice \tilde{A} can choose some *fixed* secret key rk^* and send the message $(\text{Enc}_{pk}(rk^*), \text{SEnc}_{rk^*}(m))$. If rk^* is a valid key, then \tilde{A} maintains functionality, but an adversary that knows rk^* can of course read any messages that Alice sends.

So, in order for such a protocol to have a secure reverse firewall, the RF must be able to maul the encrypted key $\text{Enc}_{pk}(rk)$ into $\text{Enc}_{pk}(rk')$ for some rk' and then convert the encrypted plaintext $\text{SEnc}_{rk}(m)$ into an encryption under this new key $\text{SEnc}_{rk'}(m)$. In particular, the symmetric-key scheme must be “key malleable.” Unfortunately, such a scheme implies public-key encryption. Therefore, our supposed “symmetric-key” scheme actually must use public-key primitives. So, hybrid encryption buys us nothing. Below, we give a proof sketch.

Remark (Informal). *Any key-malleable symmetric-key encryption scheme implies public-key encryption.*

Proof. Let $(\text{SEnc}, \text{SDec})$ be the symmetric-key scheme. We assume that we have some mauling function Maul that takes as input a ciphertext c and a description of a key-modifying function σ such that $\text{Maul}(\text{Enc}_{rk}(m), \sigma) = \text{Enc}_{\sigma(rk)}(m)$. (E.g., if the keys form a group, then σ may be a group element, and we may have $\text{Maul}(\text{Enc}_{rk}(m), \sigma) = \text{Enc}_{\sigma \cdot rk}(m)$.) Then, we construct a public-key encryption scheme $(\text{PEnc}, \text{PDec})$ as follows. The public key is $pk := \text{SEnc}_{sk}(m)$ for a uniformly random message m from the message space. To encrypt zero, we sample a random σ and set the ciphertext to $c := (\sigma, \text{Maul}(pk, \sigma))$. To encrypt one, we sample σ, σ' with $\sigma \neq \sigma'$ and send $c := (\sigma', \text{Maul}(pk, \sigma))$. I.e., an encryption of zero is an “honest mauling” of the public key, while an encryption of one is a “dishonest mauling.” It is easy to see that such a scheme is correct and secure, assuming the appropriate notions of correctness and security for the key-malleable symmetric-key scheme. \square

4 A solution using key agreement

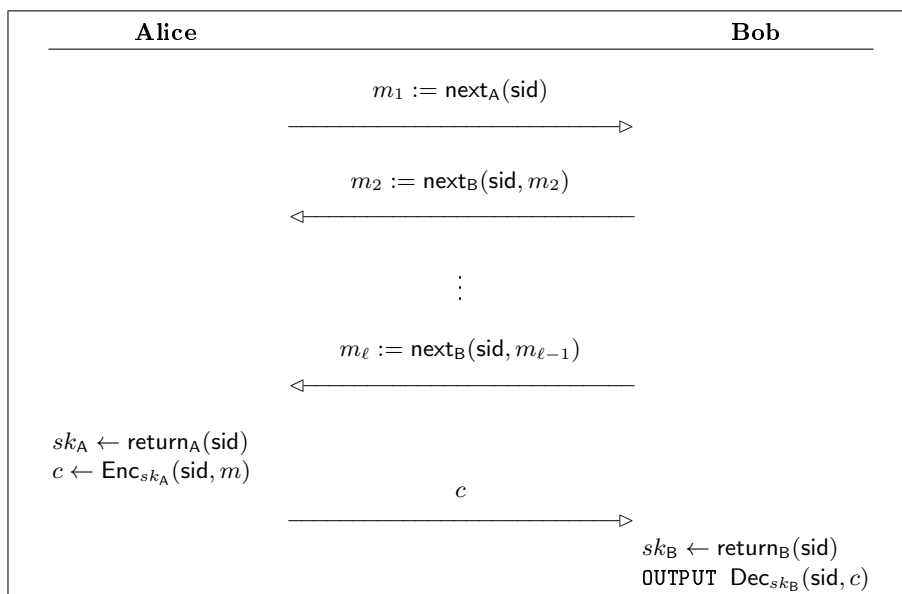


Fig. 7: The message-transfer protocol obtained by combining a key-agreement scheme (setup , next_A , next_B , return_A , return_B) and a nonce-based encryption scheme, (Enc , Dec).

In this section, we remain in the setting in which neither Alice nor Bob has a public key, so we are still interested in CPA security. (We address CCA security in the next section.) The protocol from Section 3 works, but it requires a public-key operation on the plaintext, which may be very long. In practice, this can be very inefficient. And, we showed in Section 3.1 that one common solution to this problem in the classical setting, hybrid encryption, seems hopeless with reverse firewalls because it allows Alice to *choose* a key rk that will be used to encrypt the plaintext—thus allowing a tampered version of Alice to “choose a bad key.”

So, we instead consider an alternative common solution to this efficiency problem: key agreement followed by symmetric-key encryption. (See Figure 7.) As in Appendix A, we use a nonce-based encryption scheme with unique ciphertexts. We can view this as a modification of hybrid encryption in which “Alice and Bob together choose the key rk ” that will be used to encrypt the plaintext. More importantly from our perspective, the messages that determine the key will go through the firewall. As an added benefit, once a key is established, Alice can use it to efficiently send multiple messages, not just one, without any additional public-key operations (though we do not model this here).⁵

The composition theorem below shows that this protocol can in fact have a reverse firewall for both parties, provided that the key-agreement protocol itself has a reverse firewall that satisfies some suitable security requirements. See Appendix C for the proof. In the next section, we construct such a protocol.

Theorem 2 (Composition theorem for CPA security). *Let \mathcal{W}_A and \mathcal{W}_B be reverse firewalls in the underlying key-agreement protocol in Figure 7 for Alice and Bob respectively. Let \mathcal{W}_A^* be the firewall for Alice in the full protocol in Figure 7 obtained by applying \mathcal{W}_A to the key-agreement*

⁵ To keep our definitions relatively simple, we only formally model the case in which Alice wishes to send a single message to Bob per key agreement, and we only formally prove security in this model. We note in passing that this weaker model actually only requires semantically secure symmetric-key encryption. The same is true of Theorem 4.

messages and then letting the last message through if \mathcal{W}_A does not output \perp and replacing the last message by \perp otherwise. Let \mathcal{W}_B^* be the firewall for Bob in the full protocol in Figure 7 obtained by applying \mathcal{W}_B to the key-agreement messages and simply letting the last message through. Then,

1. the protocol in Figure 7 is CPA-secure if the underlying key-agreement protocol is secure against passive adversaries and the underlying nonce-based encryption scheme is CPA-secure;
2. \mathcal{W}_B^* preserves CPA security if \mathcal{W}_B preserves security of the key-agreement protocol; and
3. \mathcal{W}_A^* preserves CPA security if the encryption scheme has unique ciphertexts and \mathcal{W}_A preserves semantic security and is exfiltration resistant against Bob.

Finally, we note that strong security preservation is not possible for this protocol (at least for Alice).

Remark 1 (Informal). There is no reverse firewall for Alice in the protocol illustrated in Figure 7 that maintains functionality and strongly preserves Alice’s security.

Proof. Consider the tampered implementation of Alice \bar{A} that goes through the key-agreement protocol as normal and then sends as its last message $c := \text{Enc}_{sk'}(\text{sid}, m)$ for some fixed key sk' chosen by the adversary by simulating a run of the key-agreement protocol. Since the firewall maintains functionality, it must be the case that the message sent by the firewall to Bob c' satisfies $\text{Dec}_{sk'}(\text{sid}, c') = m$. So, the adversary can decrypt the message itself. Clearly, this protocol is not secure (by any reasonable definition). \square

4.1 Key agreement secure against passive adversaries




Fig. 8: Diffie-Hellman key agreement over a group G of prime order p with generator g .

Theorem 2 motivates the study of unkeyed key-agreement protocols with reverse firewalls that preserve security against passive adversaries. In the classical setting (i.e., without reverse firewalls), the canonical example is the elegant key-agreement protocol of Diffie and Hellman [DH06], shown in Figure 8, whose security follows immediately from the hardness of DDH over the base group G . We use this as an example to illustrate the basic idea of a reverse firewall in the key-agreement setting.

Diffie-Hellman key agreement has a simple reverse firewall for Alice, which raises both messages to a single random power, $\alpha \in \mathbb{Z}_p^*$. We present this reverse firewall in Figure 9. Note that this firewall effectively replaces Alice’s message with a uniformly random message. Security then follows from the security of the underlying protocol, since the transcript and resulting key in the two cases are distributed identically. This very simple idea of rerandomizing Diffie-Hellman key agreement is part of all of our key-agreement protocols.

But, this protocol cannot have a reverse firewall that maintains correctness and preserves security for Bob, as we described in the introduction. In particular, we run the risk that one party has the ability to selectively reject keys.

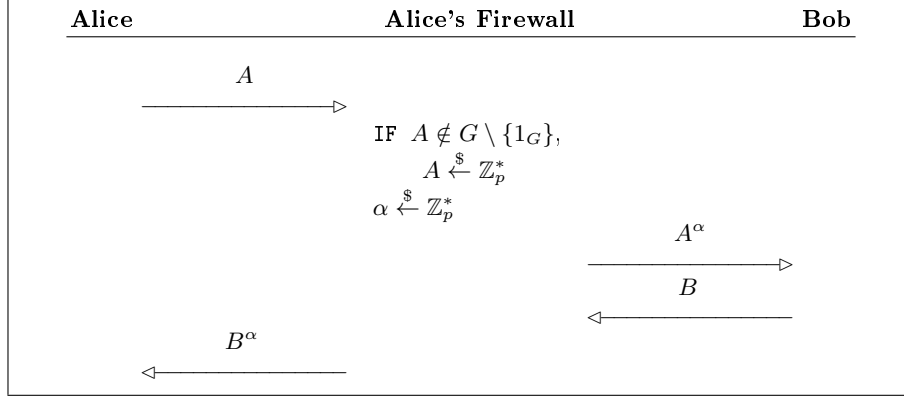


Fig. 9: Reverse firewall for Alice in the protocol from Figure 8.

To solve this problem, we add an additional message to the beginning of the protocol in which Bob commits to the message that he will send later. Of course, in order to permit a secure firewall, the commitment scheme itself must be both malleable (so that the firewall can rerandomize the underlying message that Bob has committed to, mapping a commitment of B to a commitment of B^α) and rerandomizable (so that the randomness used by Bob to commit and open will not leak any information about his message). To achieve our strongest level of security, we also need the scheme to be statistically hiding and for each commitment to have a unique opening for a given message. (These requirements are easily met in practice. For example, a simple variant of the Pedersen commitment suffices [Ped92]. For completeness, we present such a scheme in Appendix D.) The protocol is shown in Figure 10. In Figure 11, we present a single reverse firewall for this protocol that happens to work for either party. (Each party would need to deploy its own version of this firewall to guarantee its own security. It just happens that each party's firewall would have the same "code.")




Fig. 10: A variant of Diffie-Hellman key agreement over a group G of prime order p with public generator g . (Com, Open) is a commitment scheme.

Theorem 3. *The protocol in Figure 10 is secure against passive adversaries if DDH is hard in G . The reverse firewall \mathcal{W} in Figure 11 is functionality maintaining. If the commitment scheme is statistically hiding, then \mathcal{W} preserves security for Alice and is strongly exfiltration resistant against Bob. If the commitment scheme is computationally binding, then \mathcal{W} is exfiltration resistant for Bob against Alice and preserves security for Bob. \mathcal{W} also detects failure for both parties.*




Fig. 11: Reverse firewall for either Alice or Bob in the protocol from Figure 8. $\text{Maul}(C, \alpha)$ takes a commitment $C = \text{Com}(B)$ and converts it into a commitment of B^α . $\text{Rerand}(C)$ takes a commitment $C = \text{Com}(B)$ and converts it into a uniformly random commitment of B . We assume that a rerandomized and mauled commitment can be opened with access to an opening of the original commitment (and the randomness used in the rerandomization and mauling).

Proof. It is clear that the underlying protocol is secure provided that DDH is hard in G . It is also clear that \mathcal{W} maintains functionality and fails detectably for both parties. (Here, we assume that the firewall outputs \perp if it ever receives a malformed message.)

Note that, after rerandomization and mauling, the commitment C' is a uniformly random commitment of a uniformly random group element, independent of the original commitment C . Since Bob is functionality maintaining, his second message is fixed unless he can find an alternative opening for the commitment, which by assumption is computationally hard. It follows that \mathcal{W} is exfiltration resistant for Bob against Alice and preserves security for Bob.

To prove strong exfiltration resistance for Alice against Bob and security preservation, we again note that Bob's first message is a uniformly random commitment of a uniformly random group element. Since the commitment is statistically binding, it is statistically close to independent from α , regardless of Bob's choice of C . Therefore, Alice's message A is statistically close to independent from α , and A^α is statistically close to uniform. The result follows. \square

5 CCA-security using key agreement

In the setting of the previous section, with no public-key infrastructure, it is trivially impossible to achieve CCA-security. (An adversary can simply “pretend to be Bob” and read Alice's plaintext.) In this section, we show that a CCA-secure message-transmission protocol with reverse firewalls does in fact exist in the publicly keyed setting. In particular, below, we give the CCA analogue of Theorem 2, showing that a key-agreement protocol that is secure against active adversaries and has sufficiently secure reverse firewalls together with a symmetric-key encryption scheme suffices. As in the previous section, this key-agreement-based protocol has the additional benefit that it is efficient, in the sense that it does not apply public-key operations to the plaintext. In Section 5.1, we construct a key-agreement protocol that suffices.

We now present our stronger composition theorem. (Recall that Bob’s reverse firewall can only preserve CPA security. Such a firewall is already given by Theorem 2, so we do not repeat this here.) See Appendix E for the proof.

Theorem 4 (Composition theorem for CCA security). *Define \mathcal{W}_A and \mathcal{W}_A^* as in Theorem 2. Then,*

1. *the protocol in Figure 7 is CCA-secure if the underlying key-agreement protocol is secure against active adversaries for Alice and the underlying nonce-based encryption scheme is CCA-secure; and*
2. *\mathcal{W}_A^* preserves CCA-security if the encryption scheme has unique ciphertexts, the key-agreement protocol is authenticated for Bob, and \mathcal{W}_A preserves security for Alice, is exfiltration resistant against Bob with valid transcripts, and detects failure.*

5.1 Key agreement secure against active adversaries

Theorem 4 motivates the study of key-agreement protocols with reverse firewalls that preserve security against active adversaries. In the classical setting, the common solution is essentially for each of the parties to sign the transcript of this run of the protocol. However, this solution does not work in our setting because it is important for us that messages *can* be altered without breaking functionality, so that the firewall can rerandomize messages when necessary.

Of course, while we want to allow for the possibility that Alice and Bob see a different transcript but still output a key, we still want them to agree on the key itself. This leads to the idea of signing some deterministic function of the key, so that the signatures can be used to verify that the parties share the same key without necessarily requiring them to share the same transcript. This is the heart of our solution.

We also have to worry that the signatures themselves can provide channels, allowing tampered versions of the parties to leak some information. We solve this by using a unique signature scheme, as defined by [MRV99]. (See [AMV15] for a thorough analysis of signatures in the context of reverse firewalls and corrupted implementations, including alternative ways to implement signatures that would suffice for our purposes.)

Furthermore, in order for our firewall to fail detectably, it has to be able to check the signature itself—so that it can distinguish a valid transcript from an invalid one. So, we would like the parties to sign a deterministic function of g^{ab} that is efficiently computable given only access to g^a and g^b . This leads naturally to the use of a symmetric bilinear map $e : G \times G \rightarrow G_T$. The parties then sign $e(g^a, g^b)$. Of course, g^{ab} is no longer indistinguishable from random in the presence of a bilinear map. But, it can be hard to compute. So, we apply a cryptographic hash function H to g^{ab} in order to extract the final key $H(g^{ab})$.

We now provide two definitions to make this precise.

Definition 13 (Unique Signatures). *A unique signature scheme is a triple of algorithms $(\text{KeyGen}, \text{USig}, \text{Ver})$. KeyGen takes as input 1^λ where λ is the security parameter and outputs a public key pk and a private key sk . Sig takes as input the secret key sk and a plaintext m and outputs a signature τ . Ver takes as input the public key pk , a signature τ and a message m and outputs either `true` or `false`. A signature scheme is correct if $\text{Ver}_{pk}(\text{USig}_{sk}(m), m) = \text{true}$. It is unique if for each plaintext m and public key pk , there is a unique signature τ such that $\text{Ver}_{pk}(\tau, m) = \text{true}$.*

A signature scheme is secure against adaptive chosen-message existential-forgery attacks if no adversary with access to the public key and a signature oracle can produce a valid signature not returned by the oracle.

We will need to use a group with a symmetric bilinear map in which the following variant of the computational Diffie-Hellman assumption holds.

Definition 14 (Inverse CDH). For a group G of order p , we say that inverse CDH is hard in G if no probabilistic polynomial-time adversary taking input (g, g^a, g^b) where $g \xleftarrow{\$} G$ and $(a, b) \xleftarrow{\$} \mathbb{Z}_p^2$ has non-negligible probability of returning $(h, h^{1/a}, h^{1/b})$ for some element $h \in G \setminus \{1_G\}$.

Note that inverse CDH is stronger than the standard CDH assumption, as (g^{ab}, g^b, g^a) is a solution to inverse CDH.

We now present our protocol in Figure 12 with a reverse firewall for both parties in Figure 13. It requires a unique signature scheme $(\text{USig}, \text{Ver})$ with public keys pk_A for Alice and pk_B for Bob and corresponding secret keys sk_A and sk_B respectively, a base group G with generator g in which inverse CDH is hard, a target group G_T , and a non-trivial bilinear map between the two groups $e : G \times G \rightarrow G_T$. We also need a function $H : G \rightarrow \{0, 1\}^\ell$ for some polynomially large ℓ that extracts hardcore bits from CDH. Presumably a standard cryptographic hash function will work. For simplicity, we model H as a random oracle, but we note that the proof can be modified to apply to any function H such that $(g^a, g^b, H(g^{ab}))$ is indistinguishable from random. (See [Kil07] for a discussion of such functions.) We stress again that this protocol is remarkably efficient, and we think that it can and should be used in practice.

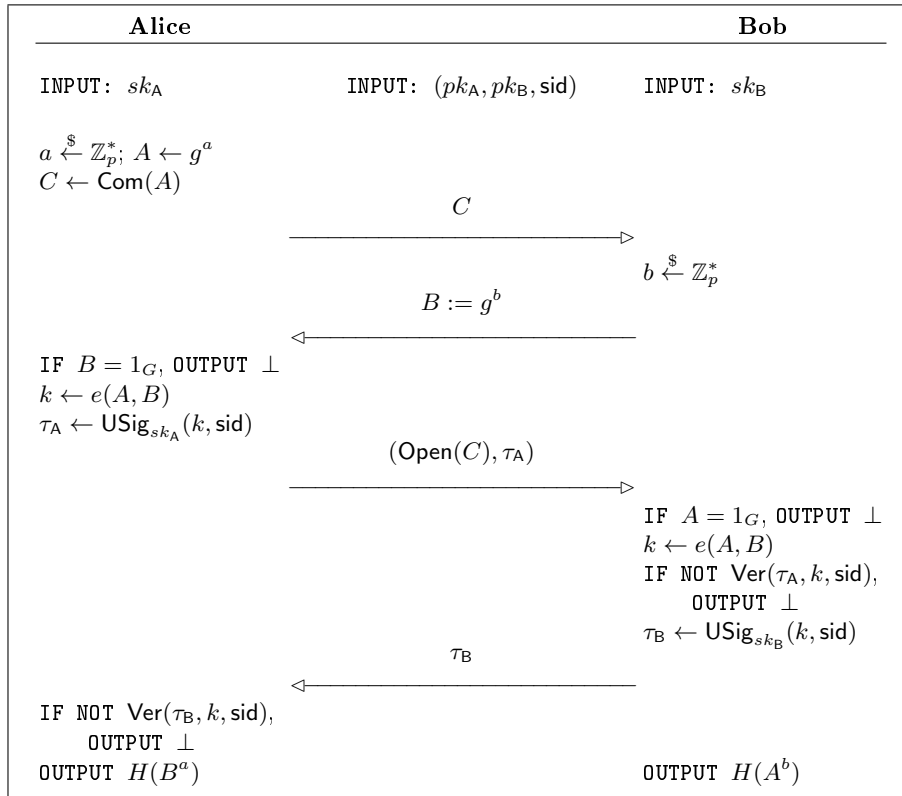


Fig. 12: Authenticated key agreement with a firewall for both parties. USig is a unique signature.

Theorem 5. The protocol shown in Figure 12 is authenticated for Bob and secure against active adversaries if the signature scheme is secure and inverse CDH is hard in G . The reverse firewall \mathcal{W} shown in Figure 13 preserves security against active adversaries for Alice, preserves authenticity, is exfiltration resistant for Alice against Bob with valid transcripts, and detects failure for Alice. \mathcal{W}




Fig. 13: Reverse firewall for either Alice or Bob in the protocol from Figure 12. C is a commitment of the group element A . $\text{Maul}(C, \alpha)$ takes a commitment $C = \text{Com}(A)$ and converts it into a commitment of A^α . $\text{Rerand}(C)$ takes a commitment $C = \text{Com}(A)$ and converts it into a uniformly random commitment of A . We assume that a rerandomized and mauled commitment can be opened with access to an opening of the original commitment.

also preserves security against active adversaries for Bob, is exfiltration resistant for Bob against Alice with valid transcripts, and detects failure for Bob.

Proof. The fact that the protocol is authenticated for Bob is clear. Informally, any valid transcript must have a valid signature from Bob. By the security of the signature scheme, except with negligible probability, the adversary cannot produce a signature with a given sid without Bob returning a key. The same proof shows that the firewall preserves authenticity. The fact that the firewall detects failure for both parties is also clear.

We prove that \mathcal{W} preserves security against active adversaries for Alice. Let \tilde{A} be a functionality-maintaining tampered implementation of Alice. Consider the following sequence of games.

- **Game 1** is the key-agreement security game against active adversaries for Alice in the protocol shown in Figure 12 with A replaced by $\mathcal{W} \circ \tilde{A}$.
- **Game 2** is **Game 1** in which the adversary never provides a signature τ_A or τ_B unless it came from a response from Alice or Bob respectively.
- **Game 3** is **Game 2** in which Alice’s commitment C is replaced by A (the value that her commitment opens to), and the firewall’s commitment C' is replaced by A^α in all runs of the protocol. We also remove Alice’s second message is simply removed and join Bob’s last two messages into a single message, (B, τ_B) .
- **Game 4** is **Game 3** in which Alice’s message A is replaced by a uniformly random group element, and the firewall takes $\alpha := 1$ (i.e., the firewall leaves all messages alone, besides checking signatures) in all runs of the protocol.

The following claim follows from the definition of adaptive chosen-message existential-forgery security.

Claim 5.1. *If the signature scheme is secure, then for any PPT adversary \mathcal{E} , $|\text{Adv}^{\text{Game } 1}(\mathcal{E}) - \text{Adv}^{\text{Game } 2}(\mathcal{E})|$ is negligible.*

It is trivial to take any adversary in **Game 2** and convert it to an adversary in **Game 3** with the same advantage. Similarly, **Game 4** only differs from **Game 3** syntactically. Indeed, $A' := A^\alpha$ is a uniformly random group element independent of everything else, and $B' := B^\alpha$ is uniformly random and independent of A' and everything else. (Recall that Bob is honest.) So, since Alice's next message is deterministic (because the signature is unique and Alice maintains functionality) and no output from Alice is given to the adversary other than these two messages (recall that the get-output_A oracle is not allowed in this game), this is merely a change of variables.

Claim 5.2. *No PPT adversary \mathcal{E} has non-negligible advantage in **Game 4** if inverse CDH is hard in G .*

Proof. We assume without loss of generality that \mathcal{E} makes no “trivial” oracle calls, whose output can be easily predicted based on prior calls, such as calls to get-output_B before the relevant run of the protocol is finished. We build an adversary \mathcal{E}' in the inverse CDH game as follows. On input (g_1, g_2, g_3) , \mathcal{E}' first generates the keys for the signature scheme, (pk_A, pk_B, sk_A, sk_B) , and passes (g_1, pk_A, pk_B) to \mathcal{E} . As \mathcal{E}' will be simulating many runs of the protocol defined by **Game 4**, for convenience, we assume that \mathcal{E} “oracle” calls to the protocol as in the key-agreement security game. \mathcal{E}' simply passes the random oracle calls of \mathcal{E} to its own random oracle (or simulates a random oracle), keeping a list of all calls. Finally, \mathcal{E}' responds to other oracle calls of \mathcal{E} as follows.

- When \mathcal{E} calls $\text{start-run}(\text{sid})$, \mathcal{E}' calls its own “oracle” $\text{start-run}(\text{sid})$.
- When \mathcal{E} calls $\text{start-challenge}(\text{sid}^*)$, \mathcal{E}' stores sid^* .
- When \mathcal{E} calls $\text{get-next}_A(\text{sid}, M)$, if $\text{sid} \neq \text{sid}^*$, \mathcal{E}' calls its own “oracle” $\text{get-next}_A(\text{sid}, M)$ and passes the response to \mathcal{E} . Otherwise, it sets $h_1 \leftarrow M$ and passes $(g_3, \text{Sig}_{sk}(e(h_2, g_3), \text{sid}))$ to \mathcal{E} .
- When \mathcal{E} calls $\text{get-next}_B(\text{sid}, M)$, if $\text{sid} \neq \text{sid}^*$, \mathcal{E}' calls its own “oracle” $\text{get-next}_B(\text{sid}, M)$ and passes the response to \mathcal{E} . Otherwise, if this is the first message of the run, it passes g_2 to \mathcal{E} . Otherwise, it sets $(\tau, h_1) \leftarrow M$.
- When \mathcal{E} calls $\text{get-output}_B(\text{sid})$, if $\text{sid} \neq \text{sid}^*$, \mathcal{E}' calls its own “oracle” $\text{get-output}_B(\text{sid})$. If $\text{sid} = \text{sid}^*$, and $e(h_1, g_3) \neq e(g_2, h_2)$, then it returns \perp . Otherwise, \mathcal{E}' responds with a uniformly random string.
- When \mathcal{E} calls get-secrets , \mathcal{E}' responds with sk .
- When \mathcal{E} calls $\text{finalize}(b^*)$, \mathcal{E}' stops simulating and proceeds as below.

Note that the view of \mathcal{E} is identical to its view in **Game 4** unless it calls the random oracle on the actual key, $h_3 := \text{DH}(h_1, g_3) = \text{DH}(g_2, h_2)$. Note as well that, because a random oracle is extractable, if \mathcal{E} has non-negligible advantage, it must call the random oracle on the key. Therefore, \mathcal{E}' can search through its random oracle queries until it finds h_3 satisfying $e(g_1, h_3) = e(h_1, g_3)$. It then returns (h_3, h_2, h_1) . If it finds nothing, it returns \perp . The result follows from noting that (h_3, h_2, h_1) is a valid solution to CDH. (5.2) ■

It follows that \mathcal{W} preserves security against active adversaries for Alice. The proof for Bob is essentially identical. And, essentially the same proof shows exfiltration resistance. □

6 Less efficient one-round protocols

Finally, we show one-round protocols in the singly keyed setting, in which Bob has a public-key/private-key pair, but Alice does not. These are essentially the single-round analogue of the two-round protocol presented in Figure 4 in Section 3. They can also be thought of as the natural extension of public-key encryption schemes to the reverse firewall setting. (In particular, these protocols are not efficient, in the sense that they use public-key operations on the plaintext, which may be very long.) Indeed, we show that the existence of such a protocol is equivalent to the existence of rerandomizable encryption, and we show how to achieve CCA-security (though not forward secrecy).

6.1 One-round CPA-secure protocols

The next theorem shows that one-round CPA-secure protocols with reverse firewalls are equivalent to rerandomizable public-key encryption.

Theorem 6. *Any (strongly) rerandomizable semantically secure public-key encryption scheme implies a one-round CPA-secure singly keyed message-transmission protocol without forward security with a reverse firewall that (strongly) preserves security and (strongly) resists exfiltration. Conversely, any one-round CPA-secure message-transmission protocol with a reverse firewall that (strongly) preserves security implies a (strongly) rerandomizable semantically secure public-key encryption scheme.*

Proof. To prove the first statement, we consider the protocol in which Alice simply sends Bob an encryption of the plaintext under Bob’s public key. Alice’s firewall applies the **Rerand** algorithm to the plaintext. (Bob does not need a firewall, since he does not send any messages.) Security of this protocol follows immediately from the security of the encryption scheme, and the fact that the firewall preserves security follows immediately from the rerandomizability of the encryption scheme.

To prove the second statement, consider the following encryption scheme. The key generation algorithm runs the **setup** algorithm of the underlying MTP protocol, receiving as output σ_A , σ_B , and π . The public key is then π and σ_A , and the private key is σ_B . The encryption algorithm first uses σ_A and an arbitrarily chosen **sid** to compute Alice’s single message in the protocol, given the plaintext. It then applies the reverse firewall to this message; the result is the ciphertext. The rerandomization algorithm simply applies the reverse firewall to this message. The security and rerandomizability of the scheme are immediate from the security of the underlying MTP and the security of the firewall respectively. \square

6.2 A one-round CCA-secure protocol

To extend this idea to stronger notions of security, we need the underlying encryption scheme to satisfy stronger notions of security. A natural candidate is CCA security. However, CCA-secure encryption schemes cannot be rerandomizable, so we need a slightly weaker notion. RCCA security, as defined by [CKN03], suffices, and rerandomizable RCCA-secure schemes do exist (see, e.g., [Gro04, PR07]), though they are relatively inefficient. (They are not *strongly* rerandomizable; their rerandomization procedures do not work on invalid ciphertexts.) We present the RCCA security game in Figure 14. In addition, we need a rerandomized ciphertext to be indistinguishable from a valid encryption *even with access to a decryption oracle*. Figure 15 and the definition below makes this precise.

<pre> proc. IND-RCCA(λ) (pk, sk) $\xleftarrow{\\$}$ KeyGen(1^λ) (m_0, m_1) $\leftarrow \mathcal{E}^{\mathcal{O}_1}(pk)$ $b \xleftarrow{\\$} \{0, 1\}$; $C^* \xleftarrow{\\$} \text{Enc}_{pk}(m_b)$ $b^* \leftarrow \mathcal{E}^{\mathcal{O}_2}(\sigma, c^*)$ OUTPUT ($b = b^*$) </pre>	<pre> proc. $\mathcal{O}_1(c)$ OUTPUT Dec$_{sk}(c)$ proc. $\mathcal{O}_2(c)$ $m \leftarrow \text{Dec}_{sk}(c)$ IF $m = m_0$ OR $m = m_1$, OUTPUT Challenge ELSE, OUTPUT m </pre>
--	--

Fig. 14: The RCCA security game.

<pre> proc. IND-RCCA(λ) (pk, sk) $\xleftarrow{\\$}$ KeyGen(1^λ) (c_0, c_1) $\leftarrow \mathcal{E}^{\mathcal{O}_1}(pk)$ $b \xleftarrow{\\$} \{0, 1\}$; $c^* \xleftarrow{\\$}$ Rerand$_{pk}(c_b)$ $b^* \leftarrow \mathcal{E}^{\mathcal{O}_2}(c^*)$ OUTPUT ($b = b^*$) </pre>	<pre> proc. $\mathcal{O}_1(c)$ OUTPUT Dec$_{sk}(c)$ proc. $\mathcal{O}_2(c)$ $m \leftarrow \text{Dec}_{sk}(c)$ IF $m = \text{Dec}_{sk}(c_0)$ OR $m = \text{Dec}_{sk}(c_1)$, OUTPUT Challenge ELSE, OUTPUT m </pre>
---	--

Fig. 15: The RCCA rerandomization game.

Definition 15. An encryption scheme is RCCA secure if no probabilistic polynomial-time adversary \mathcal{E} has non-negligible advantage in the game presented in Figure 14. It is RCCA rerandomizable if there exists an algorithm Rerand with access to the public key such that for any ciphertext c with $\text{Dec}(c) \neq \perp$, $\text{Dec}(\text{Rerand}(c)) = \text{Dec}(c)$ and no probabilistic polynomial-time adversary \mathcal{E} has non-negligible advantage in the game presented in Figure 15 when we require that $\text{Dec}(c_i) \neq \perp$. It is strongly RCCA-rerandomizable if the previous statement holds even if $\text{Dec}(c_i) = \perp$.

The below theorem is the CCA analogue of Theorem 6.

Theorem 7. Any (strongly) RCCA-rerandomizable, RCCA-secure encryption scheme implies a one-round CCA-secure singly keyed message-transmission protocol without forward security with a reverse firewall that (strongly) preserves security and (strongly) resists exfiltration.

Proof. Consider the protocol in which Alice, given as input a plaintext m and a session id sid , simply sends the message $\text{Enc}_{pk}(\text{sid}, m)$. On input c , Bob's return function computes $(\text{sid}^\dagger, m) \leftarrow \text{Dec}_{sk}(c)$. If $\text{sid}^\dagger = \text{sid}$, it outputs m . Otherwise, it outputs \perp . Alice's firewall simply applies the Rerand algorithm to Alice's message.

For any PPT adversary \mathcal{E} in the CCA-security game against this message-transmission protocol (without forward security), we construct \mathcal{E}' in the RCCA-security game against the underlying encryption scheme with $\text{Adv}(\mathcal{E}') = \text{Adv}(\mathcal{E})$. We assume without loss of generality that \mathcal{E} never makes a "useless" oracle call. I.e., \mathcal{E} never calls `start-run` with an already used `sid`, never calls `get-nextA` with a never used `sid`, never calls `get-output` on the challenge `sid`, etc. The adversary \mathcal{E}' behaves as follows in response to the oracle calls of \mathcal{E} .

- When \mathcal{E} calls the `start-run` oracle with input (sid, m) , \mathcal{E}' sets $c_{\text{sid},A} \leftarrow \text{Enc}_{pk}(\text{sid}, m)$.
- When \mathcal{E} calls `start-challenge`(sid, m_0, m_1), \mathcal{E}' sets $\text{sid}^* \leftarrow \text{sid}$. It then returns (sid^*, m_0^*) and (sid^*, m_1^*) as its challenge plaintexts, receiving in response the challenge ciphertext c^* . It sets $c_{\text{sid}^*} = c^*$.
- When \mathcal{E} calls `get-nextA`(sid), \mathcal{E}' replies with $c_{\text{sid},A}$.

- When \mathcal{E} calls $\text{get-next}_{\mathbb{B}}(\text{sid}, c)$, \mathcal{E}' sets $c_{\text{sid}, \mathbb{B}} \leftarrow c$.
- When \mathcal{E} calls $\text{get-output}_{\mathbb{B}}(\text{sid})$, \mathcal{E}' calls $\mathcal{O}_1(c_{\text{sid}, \mathbb{B}})$. If the output is not of the form (sid, m) , \mathcal{E}' responds with \perp . Otherwise, it responds with m .
- When \mathcal{E} calls $\text{finalize}(b^*)$, \mathcal{E}' simply returns b^* .

Note that the view of \mathcal{E} is identical to its view in the CCA-security game against the message-transmission protocol. Furthermore, \mathcal{E}' wins the RCCA-security game if and only if \mathcal{E} wins its simulated game. The result follows.

It follows that the protocol is CCA secure. The proof that the firewall preserves security is nearly identical to the above proof. \square

6.3 Achieving forward secrecy and efficiency?

Note that the protocols described above suffer from two problems: they do not have forward secrecy, and they are inefficient (i.e., they require public-key operations on the entire plaintext). Ideally, we would like a composition theorem in the singly keyed setting to match Theorems 2 and 4. Such a theorem would solve both problems, allowing us to achieve an efficient CCA-secure and forward-secret message-transmission protocol in the singly keyed setting. We leave this as an open question.

As a potential alternative direction to achieving forward secrecy, we note that the protocol from Theorem 7 can be converted into a two-round CCA-secure and forward-secret protocol with a reverse firewall for Alice *but no reverse firewall for Bob*. In particular, in the first round of the protocol, Bob generates a fresh pair of keys (pk^\dagger, sk^\dagger) for an RCCA-secure RCCA-rerandomizable encryption scheme and sends Alice pk^\dagger together with a signature τ of (sid, pk^\dagger) , where the signature is under Bob's signature key. Alice checks the signature and, if it is valid, sends Bob an encryption of her plaintext under the secret key. This is essentially the CCA analogue of the protocol in Section 3. As in that case, Alice's reverse firewall can simply check the signature and rerandomize the ciphertext. This protocol is in fact CCA secure, and Alice's firewall does preserve this security. However, we do not know how to construct a reverse firewall for Bob in this setting. In analogy with Section 3, a possible method would be to find an RCCA-rerandomizable encryption scheme that is also key malleable. We know of no such scheme, but even if such a scheme were constructed, the signature scheme would have to be similarly malleable, while still achieving an appropriate notion of security. So, we also leave this as an open question.

7 Conclusion and open questions

We consider the problem of message-transmission protocols in the cryptographic reverse firewalls framework of [MS15]. We show that this problem has a rich structure, in analogy with the classical setting, with a variety of solutions that require different setup assumptions, achieve different levels of security, and provide different levels of efficiency. Perhaps surprisingly, we show that it is possible to achieve concurrent, interactive CCA security in the presence of compromise against functionality-maintaining adversaries and CPA security against arbitrary adversaries. Many of our protocols (including those that provide the strongest notions of security) are very efficient and relatively simple, so that they can (and should) be implemented in practice.

Therefore, the most important work that we leave open is the implementation of our protocols. The most important theoretical question that we leave open is whether there is a non-trivial composition theorem in the singly keyed case, in analogy with Theorem 4. Note that Theorem 2 naturally extends to the singly keyed case, but we see no inherent reason why CCA-security should not be achievable from key agreement followed by symmetric-key encryption in this setting.

In addition, our work brings new attention to the question of rerandomizable RCCA-secure schemes. In particular, in Section 6, we show that such schemes give a one-round CCA-secure message-transmission protocol with a reverse firewall (without forward security). However, we do not know of such schemes that are “strongly rerandomizable” (as defined in Section 6). If such schemes existed, then we show that they would immediately imply one-round CCA-secure message transmission with a reverse firewall that *strongly* preserves security, a very powerful primitive. (Even just more efficient versions of current schemes would be quite interesting.)

References

- [AMV15] Giuseppe Ateniese, Bernardo Magri, and Daniele Venturi. Subversion-resilient signature schemes. In *CCS*, 2015.
- [AsV08] Joël Alwen, abhi shelat, and Ivan Visconti. Collusion-free protocols in the mediated model. In *CRYPTO*, 2008.
- [BBG13] James Ball, Julian Borger, and Glenn Greenwald. Revealed: how US and UK spy agencies defeat internet privacy and security. *Guardian Weekly*, September 2013.
- [BBS98] Matt Blaze, Gerrit Bleumer, and Martin Strauss. Divertible protocols and atomic proxy cryptography. In *EUROCRYPT*, 1998.
- [BD91] Mike Burmester and Yvo Desmedt. All languages in NP have divertible zero-knowledge proofs and arguments under cryptographic assumptions. In *EUROCRYPT*, 1991.
- [BDPR98] Mihir Bellare, Anand Desai, David Pointcheval, and Phillip Rogaway. Relations among notions of security for public-key encryption schemes. In *CRYPTO*, 1998.
- [BH15] Mihir Bellare and Viet Tung Hoang. Resisting randomness subversion: Fast deterministic and hedged public-key encryption in the standard model. In *CRYPTO*, 2015.
- [BJK15] Mihir Bellare, Joseph Jaeger, and Daniel Kane. Mass-surveillance without the state: Strongly undetectable algorithm-substitution attacks. In *CCS*, 2015.
- [BPR14a] Mihir Bellare, Kenneth G. Paterson, and Phillip Rogaway. Security of symmetric encryption against mass surveillance. In *CRYPTO*, 2014. Full version: [BPR14b].
- [BPR14b] Mihir Bellare, Kenneth G. Paterson, and Phillip Rogaway. Security of symmetric encryption against mass surveillance. Cryptology ePrint Archive, Report 2014/438, 2014. <http://eprint.iacr.org/>.
- [CKN03] Ran Canetti, Hugo Krawczyk, and Jesper Buus Nielsen. Relaxing chosen-ciphertext security. In *CRYPTO*, 2003.
- [CVE14a] Vulnerability summary for CVE-2014-1260 (‘Heartbleed’). <http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2014-1260>, April 2014.
- [CVE14b] Vulnerability summary for CVE-2014-1266 (‘goto fail’). <http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2014-1266>, February 2014.
- [CVE14c] Vulnerability summary for CVE-2014-6271 (‘Shellshock’). <http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2014-6271>, September 2014.
- [Des90] Yvo Desmedt. Abuses in cryptography and how to fight them. In *CRYPTO*, 1990.
- [Des94] Y. Desmedt. Subliminal-free sharing schemes. In *Information Theory*, 1994.
- [DF14] Yevgeniy Dodis and Dario Fiore. Interactive encryption and message authentication. In *Security and Cryptography for Networks*, 2014.
- [DFP15] Jean Paul Degabriele, Pooya Farshim, and Bertram Poettering. A more cautious approach to security against mass surveillance. In *Fast Software Encryption*, 2015.
- [DGG⁺15] Yevgeniy Dodis, Chaya Ganesh, Alexander Golovnev, Ari Juels, and Thomas Ristenpart. A formal treatment of backdoored pseudorandom generators. In *EUROCRYPT*, 2015.
- [DH06] W. Diffie and M. Hellman. New directions in cryptography. *IEEE Trans. Inf. Theor.*, 22(6):644–654, September 2006.
- [DPSW06] Yvo Desmedt, Josef Pieprzyk, Ron Steinfeld, and Huaxiong Wang. A non-malleable group key exchange protocol robust against active insiders. In *Information Security*, 2006.
- [EIG85] Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. In *CRYPTO*, 1985.
- [Gre14] Glenn Greenwald. *No Place to Hide: Edward Snowden, the NSA, and the U.S. Surveillance State*. Metropolitan Books, May 2014.
- [Gro04] Jens Groth. Rerandomizable and replayable adaptive chosen ciphertext attack secure cryptosystems. In *TCC*, 2004.
- [Jun15] Juniper vulnerability. <https://kb.juniper.net/InfoCenter/index?page=content&id=JSA10713>, 2015.

- [Kil07] Eike Kiltz. Chosen-ciphertext secure key-encapsulation based on Gap Hashed Diffie-Hellman. In *PKC*, 2007.
- [KS05] Jonathan Katz and Ji Sun Shin. Modeling insider attacks on group key-exchange protocols. In *CCS*, 2005.
- [LHA⁺12] Arjen K. Lenstra, James P. Hughes, Maxime Augier, Joppe W. Bos, Thorsten Kleinjung, and Christophe Wachter. Public keys. In *CRYPTO*, 2012.
- [LMs05] Matt Lepinski, Silvio Micali, and abhi shelat. Collusion-free protocols. In *STOC*, 2005.
- [MRV99] Silvio Micali, Michael Rabin, and Salil Vadhan. Verifiable random functions. In *FOCS*, 1999.
- [MS14] Ilya Mironov and Noah Stephens-Davidowitz. Cryptographic reverse firewalls. Cryptology ePrint Archive, Report 2014/758, 2014. <http://eprint.iacr.org/2014/758>.
- [MS15] Ilya Mironov and Noah Stephens-Davidowitz. Cryptographic reverse firewalls. In *EUROCRYPT*, 2015. Available from [MS14].
- [OO90] Tatsuaki Okamoto and Kazuo Ohta. Divertible zero knowledge interactive proofs and commutative random self-reducibility. In *EUROCRYPT*, 1990.
- [Ped92] TorbenPryds Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In *CRYPTO*, 1992.
- [PLS13] Nicole Perlroth, Jeff Larson, and Scott Shane. N.S.A. able to foil basic safeguards of privacy on Web. *The New York Times*, September 2013.
- [PR07] Manoj Prabhakaran and Mike Rosulek. Rerandomizable RCCA encryption. In *CRYPTO*, 2007.
- [PW03] Josef Pieprzyk and Huaxiong Wang. Key control in multi-party key agreement protocols. Workshop on Coding, Cryptography and Combinatorics, 2003.
- [RTYZ15] Alexander Russell, Qiang Tang, Moti Yung, and Hong-Sheng Zhou. Cliptography: Clipping the power of kleptographic attacks. Cryptology ePrint Archive, Report 2015/695, 2015. <https://eprint.iacr.org/2015/695>.
- [SFKR15] Bruce Schneier, Matthew Fredrikson, Tadayoshi Kohno, and Thomas Ristenpart. Surreptitiously weakening cryptographic systems. Technical report, IACR Cryptology ePrint Archive, 2015: 97, 2015. <http://eprint.iacr.org/2015/97>.
- [Sim84] Gustavus Simmons. The prisoners' problem and the subliminal channel. In David Chaum, editor, *CRYPTO*, 1984.
- [Sup15] <https://www.us-cert.gov/ncas/alerts/TA15-051A>, February 2015.

A The symmetric-key setting

As a warmup, we first consider the setting in which Alice and Bob share a private key. We observe that an elegant one-round protocol due to Bellare, Paterson, and Rogaway provides a solution that does not even need a reverse firewall [BPR14a]. We will also use this scheme in the sequel to build protocols that do not rely on shared private keys. We first define nonce-based encryption.

Definition 16 (Nonce-based encryption). A nonce-based symmetric-key encryption scheme is a pair of deterministic algorithms (Enc, Dec) . Enc takes as input a key from a key space \mathcal{K} , a nonce from a nonce space \mathcal{N} , and a plaintext from a plaintext space \mathcal{M} and outputs a ciphertext from a ciphertext space \mathcal{C} . Dec takes as input a key, a nonce, and a ciphertext and returns a plaintext or the special symbol \perp . The scheme is correct if for any key sk , nonce r , and plaintext m , $\text{Dec}(r, \text{Enc}_{sk}(r, m)) = m$.

Such a scheme is CPA secure if no probabilistic polynomial-time adversary can distinguish between $\text{Enc}_{sk}(r^*, m_0)$ and $\text{Enc}_{sk}(r^*, m_1)$ with non-negligible advantage where r^* , m_0 , and m_1 are adversarially chosen when given access to an encryption oracle that outputs \perp unless given a unique nonce r . It is CCA-secure if no probabilistic polynomial-time has non-negligible advantage when also given access to a decryption oracle that outputs \perp if $r = r^*$.

Such a scheme (Enc, Dec) has unique ciphertexts if for any key sk , message m , and nonce r , there is exactly one ciphertext c such that $\text{Dec}(r, c) = m$.

Theorem 8. Let (Enc, Dec) be a nonce-based symmetric-key encryption scheme. Then, if the encryption scheme is CPA-secure (resp. CCA-secure) the one-round protocol in which Alice sends

$\text{Enc}_{sk}(\text{sid}, m)$ and Bob returns $\text{Dec}_{sk}(\text{sid}, m)$ is a CPA-secure (resp. CCA-secure) one-round privately keyed message-transmission protocol without forward secrecy. If the encryption scheme has unique ciphertexts, then the “trivial” reverse firewall that simply passes Alice’s messages to Bob unchanged preserves security and is exfiltration resistant against Bob.

See, e.g., [BPR14a] for formal analysis and the construction of such a scheme. The key thing to note from our perspective is that, as Bellare et al. observe, the fact that the encryption scheme has unique ciphertexts implies that any tampered version of Alice that maintains functionality necessarily behaves identically to honest Alice. The next theorem shows that we essentially cannot do any better for a one-round protocol without using public-key primitives. The proof is in appendix F.

Theorem 9. *There is a black-box reduction from semantically secure public-key encryption to CPA-secure symmetric-key encryption with at least four possible plaintexts and a reverse firewall that strongly preserves CPA security.*

Of course, the primary drawbacks of this approach are that it requires Alice and Bob to share a secret key and that it does not offer forward secrecy.

B Rerandomizable and key-malleable public-key encryption

Definition 17 (Public-key encryption). A public-key encryption scheme is a triple of efficient algorithms $(\text{KeyGen}, \text{Enc}, \text{Dec})$. KeyGen takes as input 1^λ where λ is the security parameter and outputs a public-key/private-key pair, (pk, sk) . Enc takes as input the public key and a plaintext m from some plaintext space \mathcal{M} and outputs a ciphertext c from some ciphertext space \mathcal{C} . Dec takes as input a ciphertext and the private key and outputs a plaintext or the special symbol \perp . We sometimes omit the keys from the input to Enc and Dec and the security parameter input to KeyGen . The scheme is correct if $\text{Dec}(\text{Enc}(m)) = m$ for all $m \in \mathcal{M}$. The scheme is semantically secure if for any adversarially chosen pair of plaintexts (m_0, m_1) , $\text{Enc}(m_0)$ is computationally indistinguishable from $\text{Enc}(m_1)$.

Definition 18 (Rerandomizable encryption). A semantically secure public-key encryption scheme is rerandomizable if there is an efficient algorithm Rerand (with access to the public key) such that for any ciphertext c such that $\text{Dec}(c) \neq \perp$, we have $\text{Rerand}(\text{Dec}(c)) = \text{Dec}(c)$, and the pair $(c, \text{Rerand}(c))$ is computationally indistinguishable from $(c, \text{Rerand}(\text{Enc}(0)))$. We say that it is strongly rerandomizable if the previous property holds even when $\text{Dec}(c) = \perp$.

Definition 19 (Key malleability). A public-key encryption scheme is key malleable if (1) the output of KeyGen is distributed uniformly over the space of valid keys; (2) for each public key pk there is a unique associated private key sk ; and (3) there is a pair of efficient algorithms KeyMaul and CKeyMaul that behave as follows. KeyMaul is a randomized algorithm that takes as input a public key pk and returns a new public key pk' whose distribution is uniformly random over the public key space and independent of pk . Let (sk, pk) be a private key/public key pair. Let (sk', pk') be the unique pair associated with randomness r such that $pk' = \text{KeyMaul}(pk; r)$. Then, CKeyMaul takes as input a ciphertext c and randomness r and returns c' such that $\text{Dec}_{sk'}(c) = \text{Dec}_{sk}(c')$. We suppress the input r when it is understood. Furthermore, we require that KeyMaul outputs a uniformly random key pk' if called on input that is not in the public-key space.

Example. It is well-known that ElGamal encryption [ElG85] is both key malleable and strongly rerandomizable. In particular, given an ElGamal public-key (g, h) over a group of order p , a ciphertext (u, v) can be rerandomized by applying the operation $(u, v) \rightarrow (g^r u, h^r v)$ where $r \xleftarrow{\$} \mathbb{Z}_p^*$ is chosen

uniformly at random. The public key can be mauald by applying the operation $(g, h) \rightarrow (g^\alpha, h^\beta)$ where $(\alpha, \beta) \xleftarrow{\$} (\mathbb{Z}_p^*)^2$ are chosen uniformly and independently at random. Finally, a ciphertext (u, v) under key (g^α, h^β) can be converted into a ciphertext under (g, h) by applying the operation $(u, v) \rightarrow (u^{\beta/\alpha}, v)$.

C Proof of Theorem 2

Proof. The security of the underlying protocol (i.e., without firewalls) follows by a folklore composition theorem. We assume that the last message of the key-agreement protocol is sent by Bob. (This is without loss of generality, as we can always add an empty message from Bob at the end of the protocol.) We prove Item 2. The proof of Item 3 is a simplified version of the proof of Theorem 4.

Let \tilde{B} be some functionality-maintaining tampered implementation of Bob. Let \tilde{B}^* be the tampered implementation of Bob in the key-agreement protocol obtained by simply “truncating” \tilde{B} . (This is merely a syntactic change.) We assume without loss of generality that the adversary makes no “trivial” calls whose output can be trivially predicted from its previous oracle calls. E.g., it makes no calls to `get-output` (which always either returns \perp or the plaintext that the adversary provided for the corresponding `sid` in the CPA-security game), no `get-next` calls with modified messages, etc. Consider the following sequence of games.

- **Game 1** is the CPA-security game against the message-transfer protocol with Bob replaced by $\mathcal{W}_B^* \circ \tilde{B}$.
- **Game 2** is **Game 1** in which the final message of the challenge run of the protocol is replaced by an encryption of the challenge plaintext m_b under a uniformly random key sk .
- **Game 3** is **Game 2** in which the final message of the challenge run is replaced by an encryption of 0 under a uniformly random key.

Note that no adversary can have non-zero advantage in the last game, as none of the messages depend on the challenge bit.

Claim 2.1. *If the encryption scheme has unique ciphertexts, the key-agreement protocol is secure against passive adversaries, and \mathcal{W}_B preserves this security, then for any PPT adversary \mathcal{E} , $|\text{Adv}^{\text{Game 1}}(\mathcal{E}) - \text{Adv}^{\text{Game 2}}(\mathcal{E})|$ is negligible.*

Proof. We construct a passive adversary \mathcal{E}' in the security game against the key-agreement protocol with Bob replaced by $\mathcal{W}_B \circ \tilde{B}^*$ as follows. \mathcal{E}' receives as input the public parameters π and passes them to \mathcal{E} . \mathcal{E}' then selects $b^\dagger \xleftarrow{\$} \{0, 1\}$ uniformly at random and sets $S_A^\dagger \leftarrow \emptyset$. It responds to the (non-trivial) oracle calls of \mathcal{E} as follows.

- When \mathcal{E} calls `start-run`(`sid`, m), \mathcal{E}' adds (`sid`, m) to S_A^\dagger and calls its own oracle `start-run`(`sid`).
- When \mathcal{E} calls `start-challenge`(`sid`^{*}, m_0, m_1), \mathcal{E}' adds (`sid`^{*}, m_{b^\dagger}) to S_A^\dagger and calls its own oracle `start-challenge`(`sid`^{*}).
- When \mathcal{E} calls `get-next`_A(`sid`, M), \mathcal{E}' calls its own oracle `get-next`_A(`sid`, M). If this is not the last message of the protocol, \mathcal{E}' then simply passes the resulting message to \mathcal{E} . If it is the last message, it \mathcal{E}' calls its own oracle `get-output`_A(`sid`), receiving as output some key sk . It responds with `Enc` _{sk} (`sid`, m) where m is the unique plaintext such that $(\text{sid}, m) \in S_A^\dagger$. (Note that since the oracle calls of \mathcal{E} are non-trivial, it must have made a unique call to either `start-run` or `start-challenge` with this `sid`.)
- When \mathcal{E} calls `get-next`_B(`sid`, M), if this is not the last message of the protocol, \mathcal{E}' calls its own oracle `get-next`_B(`sid`, M) and passes the resulting message to \mathcal{E} . Otherwise, \mathcal{E}' does nothing.

- When \mathcal{E} calls `get-secrets`, \mathcal{E}' calls its own oracle `get-secrets` and passes the result to \mathcal{E} .
- When \mathcal{E} calls `finalize(b^*)`, \mathcal{E}' returns 1 if $b^\dagger = b^*$ and 0 otherwise.

Suppose the challenge bit b in the key-agreement security game is 0, so that the challenge key sk^* corresponding to the challenge session id sid^* was selected uniformly at random. Then, clearly the view of \mathcal{E} is identical to its view in **Game 2**, and \mathcal{E}' correctly outputs 0 if and only if \mathcal{E} “loses” its simulated game by returning $b^\dagger = b^*$. If, on the other hand, the challenge bit b is 1, then the view of \mathcal{E} is identical to its view in **Game 1**, and \mathcal{E}' correctly outputs 1 if and only if \mathcal{E} “wins” its simulated game. The result follows. (2.1) ■

Claim 2.2. *If the encryption scheme is CPA-secure, then for any PPT adversary \mathcal{E} , $|\text{Adv}^{\text{Game } 2}(\mathcal{E}) - \text{Adv}^{\text{Game } 3}(\mathcal{E})|$ is negligible.*

Proof. We construct an adversary \mathcal{E}' in the CPA-security game against the encryption scheme as follows. \mathcal{E}' first runs the `setup` procedure of the key-agreement protocol, receiving as output σ_A, σ_B , and π . \mathcal{E}' will simulate many runs of the key-agreement protocol with Bob replaced by $\mathcal{W}_B \circ \tilde{\mathbf{B}}^*$ and input $(\sigma_A, \sigma_B, \pi)$. For convenience, we give \mathcal{E}' an “oracle interface” to these simulated runs with “oracle” calls `start-run`, `get-nextA`, `get-nextB`, and `get-outputA` as in the key-agreement security game (Figure 3). \mathcal{E}' sends π to \mathcal{E} , selects $b^\dagger \xleftarrow{\$} \{0, 1\}$ uniformly at random, and sets $S_A^\dagger \leftarrow \emptyset$. It then responds to the oracle queries of \mathcal{E} as follows.

- When \mathcal{E} calls `start-run(sid, m)`, \mathcal{E}' adds (sid, m) to S_A^\dagger . It calls its own “oracle” `start-run(sid)`.
- When \mathcal{E} calls `start-challenge(sid^*, m_0, m_1)`, \mathcal{E}' , if $b^\dagger = 0$, it sends the challenge $(\text{sid}^*, m_0, 0)$ to its challenger. Otherwise, it sends the challenge $(\text{sid}^*, 0, m_1)$. It stores the resulting challenge ciphertext c^* and calls its “oracle” `start-run(sid)`.
- When \mathcal{E} calls `get-nextA(sid, M)`, \mathcal{E}' calls its own “oracle” `get-nextA(sid, M)`. If this call does not correspond to the last message of the relevant run of the message-transfer protocol, \mathcal{E}' simply passes the response to \mathcal{E} . If this is the last message and $\text{sid} = \text{sid}^*$, \mathcal{E}' sends c^* to \mathcal{E} . Otherwise, \mathcal{E}' sets $sk \leftarrow \text{get-output}_A(\text{sid})$, finds the unique m such that $(\text{sid}, m) \in S_A^\dagger$, and sends $\text{Enc}_{sk}(\text{sid}, m)$ to \mathcal{E} .
- When \mathcal{E} calls `get-nextB(sid, M)`, if this is not the last message of the protocol, \mathcal{E}' calls its own “oracle” `get-nextB(sid, M)` and passes the response to \mathcal{E} . If this is the last message, \mathcal{E}' does nothing.
- When \mathcal{E} calls `get-secrets`, \mathcal{E}' responds with (σ_A, σ_B) .
- When \mathcal{E} calls `finalize(b^*)`, \mathcal{E}' returns b^* .

Let b be the challenge bit of the CPA-security game. If $b = b^\dagger$, then the challenge ciphertext is an encryption of m_b , as in **Game 2** and \mathcal{E}' correctly outputs b if and only if \mathcal{E} “wins” and $b^* = b^\dagger$. Otherwise, the challenge ciphertext is an encryption of 0 as in **Game 3** and \mathcal{E}' correctly outputs b if and only if \mathcal{E} “loses.” The result follows. (2.2) ■

The result follows. □

D A suitable malleable commitment scheme

We briefly describe a simple commitment scheme that is statistically hiding, computationally binding, and malleable in the way that we need. It is a basic variant of the Pedersen commitment [Ped92]. We first provide the relevant definitions. For our application, we require that the committed plaintext is a group element and that the commitment can be mauled in such a way that a commitment

to group element B can be converted into a commitment to group element B^α . As such, we define a malleable commitment scheme in this specific setting.

Definition 20 (Commitment scheme). A commitment scheme is a collection of four efficient algorithms (KeyGen , Com , Open , Ver). KeyGen takes as input 1^λ , where λ is the security parameter, and outputs public parameters ρ . Com takes as input the public parameters ρ , a plaintext m , and randomness r , and outputs a commitment C . Open takes as input the public parameters ρ , a commitment C , and randomness r and outputs an opening x . Finally, Ver takes as input a commitment C and opening x and outputs either a plaintext m or the special symbol \perp . For simplicity, we often omit references to the public parameters and the randomness.

The scheme is correct if a commitment opens to the committed message, i.e., $\text{Ver}(C, x) = m$ whenever $C = \text{Com}(m)$ and $x = \text{Open}(C)$. The scheme is perfectly hiding if for any two plaintexts m_1, m_2 , the distribution of $\text{Com}(m_1)$ is identical to that of $\text{Com}(m_2)$. The scheme is computationally binding if no efficient adversary can produce a commitment C and two openings x, x' that verify to different plaintexts (and not \perp).

Definition 21 (Additional commitment properties). A commitment scheme is tight if for each commitment C and plaintext m , there is a unique opening x such that $\text{Ver}(C, x) = m$.

A commitment scheme is rerandomizable if there exists a pair of efficient algorithms (Rerand , OpenRerand) such that

1. for any commitment C and uniformly random r , $\text{Rerand}(C, r)$ is uniformly random over the commitment space; and
2. for any commitment C and opening x , $\text{Ver}(C, x) = \text{Ver}(C', x')$, where $C' = \text{Rerand}(C, r)$ and $x' = \text{OpenRerand}(x, r)$.

Similarly, a commitment scheme is malleable if the plaintext space is a group of order p (written multiplicatively) and there exists a pair of efficient algorithms (Maul , OpenMaul) such that for any commitment C and opening x , $\text{Ver}(C, x)^\alpha = \text{Ver}(C', x')$, where $C' = \text{Maul}(C, \alpha)$ and $x' = \text{OpenMaul}(x, \alpha)$ for $\alpha \in \mathbb{Z}_p$. In the main body, we omit explicit reference to the functions OpenRerand and OpenMaul .

We now describe a commitment scheme that suffices for our purposes.

- KeyGen takes as input a group G of order p and returns two uniformly random non-identity group elements g, h .
- Com takes as input a group element $m \in G$ and randomness $(r, s) \in \mathbb{Z}_p^2$ and returns $C := (g^r h^s, h^s m)$.
- Open simply outputs the randomness (r, s) .
- Ver takes as input $C = (c_1, c_2)$ and (r, s) . It first checks that $c_1 = g^r h^s$. If so, it returns $h^{-s} c_2$.
- Rerand takes as input $C = (c_1, c_2)$ and randomness (r', s') and returns $C' := (g^{r'} h^{s'} c_1, h^{s'} c_2)$.
- OpenRerand takes as input an opening (r, s) and randomness (r', s') and returns $(r + r', s + s')$.
- Maul takes as input $C = (c_1, c_2)$ and $\alpha \in \mathbb{Z}_p$ and returns $C' := (c_1^\alpha, c_2^\alpha)$.
- OpenMaul takes as input (r, s) and α and returns $(\alpha r, \alpha s)$.

The following proposition is immediate from inspection and the security of the standard Pedersen commitment scheme [Ped92].

Proposition 1. *The above commitment scheme is correct, statistically hiding, tight, rerandomizable, and malleable. If the discrete log is hard over G , it is also computationally binding.*

E Proof of Theorem 4

Proof. Item 1 (the security of the underlying protocol without reverse firewalls) follows by a folklore composition theorem. We assume that the last message of the key-agreement protocol is sent by Bob. (This is without loss of generality, as we can always add an empty message from Bob at the end of the protocol.)

Let \tilde{A} be some functionality-maintaining tampered implementation of Alice in the protocol from Figure 7. Let \tilde{A}^* be the “truncation” of \tilde{A} to the key-agreement protocol. Note that \tilde{A}^* produces valid transcripts (though it may not preserve functionality). Let q be some polynomial bound on the number of oracle calls made by the adversary in the CCA-security game. We assume without loss of generality that the adversary makes no “trivial” calls whose output can be trivially predicted from its previous oracle calls. E.g., it makes no calls to $\text{get-output}_B(\text{sid})$ where sid does not correspond to a completed run of the protocol, no $\text{get-next}_A(\text{sid}, \cdot)$ without first calling $\text{start-run}(\text{sid}, m)$ or $\text{start-challenge}(\text{sid}, m_0, m_1)$, etc. Consider the following sequence of games.

- **Game 1** is the CCA-security game against the message-transmission protocol with Alice replaced by $\mathcal{W}_A^* \circ \tilde{A}$.
- **Game 2** is **Game 1** in which Alice never sends the final message in a run of the protocol unless Bob’s key sk_B is well-defined. (I.e., a call to the $\text{return}_B(\text{sid})$ procedure of the underlying key-agreement protocol for the relevant sid does not return \perp .)
- For $i = 1, \dots, q$, **Game $i + 2$** is **Game $i + 1$** in which the final message of the i th run of the protocol is replaced by an encryption of the relevant plaintext m under a uniformly random key.
- **Game $q + 3$** is **Game $q + 2$** in which the final message of the challenge run is replaced by an encryption of 0 under a uniformly random key.
- **Game $q + 4$** is **Game $q + 3$** in which the final message is removed from each run of the protocol and the oracle get-output_B is removed.
- For $i = 1, \dots, q$, **Game $q + 4 + i$** is **Game $q + 3 + i$** in which $\mathcal{W}_A^* \circ \tilde{A}$ is replaced by $\mathcal{W}_A^* \circ A$ (the honest implementation of Alice composed with the firewall) in the i th run of the protocol.

Note that no adversary can have any advantage in the last game because none of the responses to any of the adversary’s queries depend on the challenge bit.

Claim 4.1. *If the encryption scheme has unique ciphertexts, the key-agreement protocol is authenticated for Bob, and \mathcal{W}_A fails detectably and preserves authenticity, then for any PPT adversary \mathcal{E} , $|\text{Adv}^{\text{Game 1}}(\mathcal{E}) - \text{Adv}^{\text{Game 2}}(\mathcal{E})|$ is negligible.*

Proof. We construct an adversary \mathcal{E}' in the authentication game against the key-agreement protocol with Alice replaced by $\mathcal{W}_A \circ \tilde{A}^*$ as follows. \mathcal{E}' receives as input the public parameters π and passes them to \mathcal{E} . \mathcal{E}' then selects $b^\dagger \xleftarrow{\$} \{0, 1\}$ uniformly at random and sets $S_A^\dagger, S_B^\dagger, \text{keys} \leftarrow \emptyset$. \mathcal{E}' then responds to oracle calls as follows.

- When \mathcal{E} calls $\text{start-run}(\text{sid}, m)$, \mathcal{E}' adds (sid, m) to S_A^\dagger and calls its own oracle $\text{start-run}(\text{sid}, m)$.
- When \mathcal{E} calls $\text{start-challenge}(\text{sid}, m_0, m_1)$, \mathcal{E}' adds $(\text{sid}, m_{b^\dagger})$ to S_A^\dagger and calls its own oracle $\text{start-run}(\text{sid}, m_{b^\dagger})$.
- When \mathcal{E} calls $\text{get-next}_A(\text{sid}, M)$, \mathcal{E}' calls its own oracle $\text{get-next}_A(\text{sid}, M)$. If this is not the last message of this run of the protocol, \mathcal{E}' then simply passes the resulting message to \mathcal{E} . If it is the last message, it checks if the transcript of the underlying key-agreement protocol is valid for $\mathcal{W}_A \circ A$ (using the efficient algorithm guaranteed by detectable failure). If it is invalid, \mathcal{E}' responds to \mathcal{E} with the special symbol \perp . If it is valid, let m be the unique plaintext and index

such that $(\text{sid}, m) \in S_A^\dagger$ and let $sk \leftarrow \text{get-output}_B(\text{sid})$. If $sk = \perp$, then \mathcal{E}' returns the transcript of this run of the key-agreement protocol (and wins the authentication game). Otherwise, it adds (sid, sk) to keys and responds with $\text{Enc}_{sk}(\text{sid}, m)$.

- When \mathcal{E} calls $\text{get-next}_B(\text{sid}, M)$, if this is not the last message of the protocol, \mathcal{E}' calls its own oracle $\text{get-next}_B(\text{sid}, M)$ and passes the resulting message to \mathcal{E} . If this is the last message of the protocol, it adds (sid, M) to S_B^\dagger and sends nothing to \mathcal{E} .
- When \mathcal{E} calls $\text{get-output}_B(\text{sid})$, \mathcal{E}' finds the unique M and sk such that $(\text{sid}, M) \in S_B$ and $(\text{sid}, sk) \in \text{keys}$. It computes $\text{Dec}_{sk}(\text{sid}, M)$ and responds with the result.
- When \mathcal{E} calls get-secrets , \mathcal{E}' calls its own oracle get-secrets and passes the result to \mathcal{E} .
- When \mathcal{E} calls $\text{finalize}(b^*)$, \mathcal{E}' simply terminates.

Note that the view of \mathcal{E} is identical to its view in both **Game 1** and **Game 2** unless at some point it constructs a valid transcript such that sk_B is not well-defined. If it does construct such a transcript, then \mathcal{E}' wins the authentication game. The result follows. (4.1) ■

Claim 4.2. *If the key-agreement protocol is secure against active adversaries for Alice, the encryption scheme has unique ciphertexts, and \mathcal{W}_A preserves Alice's security and fails detectably, then for any PPT adversary \mathcal{E} , $|\text{Adv}^{\text{Game } i+1}(\mathcal{E}) - \text{Adv}^{\text{Game } i+2}(\mathcal{E})|$ is negligible.*

Proof. We construct an adversary \mathcal{E}' in the security game against the key-agreement protocol with Alice replaced by $\mathcal{W}_A \circ \tilde{A}^*$ as follows. \mathcal{E}' receives as input the public parameters π and passes them to \mathcal{E} . \mathcal{E}' then selects $b^\dagger \xleftarrow{\$} \{0, 1\}$ uniformly at random and sets $S_A^\dagger, S_B^\dagger, \text{keys} \leftarrow \emptyset$ and $j \leftarrow 0$. It responds to the oracle calls of \mathcal{E} as follows.

- When \mathcal{E} calls $\text{start-run}(\text{sid}, m)$, \mathcal{E}' adds (sid, j, m) to S_A^\dagger . If $j = i$, it calls its own oracle $\text{start-challenge}(\text{sid}^*, m)$; otherwise it calls $\text{start-run}(\text{sid}, m)$. Finally, it increments j .
- When \mathcal{E} calls $\text{start-challenge}(\text{sid}, m_0, m_1)$, \mathcal{E}' adds $(\text{sid}, m_{b^\dagger})$ to S_A^\dagger and (sid, j) to S_B^\dagger . If $j = i$, it calls its own oracle $\text{start-challenge}(\text{sid}^*, m_{b^\dagger})$; otherwise it calls $\text{start-run}(\text{sid}, m_{b^\dagger})$. Finally, it increments j .
- When \mathcal{E} calls $\text{get-next}_A(\text{sid}, M)$, \mathcal{E}' calls its own oracle $\text{get-next}_A(\text{sid}, M)$. If this is not the last message of this run of the protocol, \mathcal{E}' then simply passes the resulting message to \mathcal{E} . If it is the last message, it checks if the transcript of the underlying key-agreement protocol is valid for $\mathcal{W}_A \circ A$ (using the efficient algorithm guaranteed by detectable failure). If it is invalid, \mathcal{E}' responds to \mathcal{E} with the special symbol \perp . If it is valid, let m, k be the unique plaintext and index such that $(\text{sid}, k, m) \in S_A^\dagger$. If $k < i$, \mathcal{E}' selects $sk \xleftarrow{\$} \mathcal{K}$. If $k \geq i$, \mathcal{E}' sets $sk \leftarrow \text{get-output}_B(\text{sid})$. Finally, it responds to \mathcal{E} with the message $\text{Enc}_{sk}(\text{sid}, m)$.
- When \mathcal{E} calls $\text{get-next}_B(\text{sid}, M)$, \mathcal{E}' calls its own oracle $\text{get-next}_B(\text{sid}, M)$ and passes the resulting message to \mathcal{E} . If this is the last message of the protocol, it also adds (sid, M) to S_B^\dagger .
- When \mathcal{E} calls $\text{get-output}_B(\text{sid})$, \mathcal{E}' finds the unique M and sk such that $(\text{sid}, M) \in S_B^\dagger$ and $(\text{sid}, sk) \in \text{keys}$. It computes $\text{Dec}_{sk}(\text{sid}, M)$ and responds with the result.
- When \mathcal{E} calls get-secrets , \mathcal{E}' calls its own oracle get-secrets and passes the result to \mathcal{E} .
- When \mathcal{E} calls $\text{finalize}(b^*)$, \mathcal{E}' returns 1 if $b^\dagger = b^*$ and 0 otherwise.

Suppose the challenge bit b in the key-agreement security game is 0 so that the challenge key sk^* corresponding to the i th run of the protocol was selected uniformly at random or is the special symbol \perp . Then, clearly the view of \mathcal{E} is identical to its view in **Game $i + 2$** , and \mathcal{E}' correctly outputs 0 if and only if \mathcal{E} “loses” its simulated game by returning $b^\dagger = b^*$. If, on the other hand, the challenge bit b is 1, then the view of \mathcal{E} is identical to its view in **Game $i + 1$** , and \mathcal{E}' correctly outputs 1 if and only if \mathcal{E} “wins” its simulated game. The result follows. (4.2) ■

Claim 4.3. *If the encryption scheme is CCA-secure, then for any PPT adversary \mathcal{E} , $|\text{Adv}^{\text{Game } q+2}(\mathcal{E}) - \text{Adv}^{\text{Game } q+3}(\mathcal{E})|$ is negligible.*

Proof. We construct an adversary \mathcal{E}' in the CCA-security game against the encryption scheme as follows. \mathcal{E}' first runs the **setup** procedure of the key-agreement protocol, receiving as output σ_A, σ_B , and π . \mathcal{E}' will simulate many runs of the key-agreement protocol with Alice replaced by $\mathcal{W}_A \circ \mathbf{A}^*$ and input $(\sigma_A, \sigma_B, \pi)$. For convenience, we give \mathcal{E}' an “oracle interface” to these simulated runs with “oracle” calls **start-run**, **get-next_A**, and **get-next_B** as in the key-agreement security game (Figure 3). \mathcal{E}' sends π to \mathcal{E} , selects $b^\dagger \xleftarrow{\$} \{0, 1\}$ uniformly at random, and sets $S_A, S_B, \text{keys} \leftarrow \emptyset$. It then responds to the oracle queries of \mathcal{E} as follows.

- When \mathcal{E} calls **start-run**(sid, m), \mathcal{E}' adds (sid, m) to S_A . It calls its own “oracle” **start-run**(sid, m).
- When \mathcal{E} calls **start-challenge**(sid^*, m_0, m_1), if $b^\dagger = 0$, it \mathcal{E}' sends the challenge $(\text{sid}^*, m_0, 0)$ to its challenger. Otherwise, it sends the challenge $(\text{sid}^*, 0, m_1)$. It stores the resulting challenge ciphertext c^* . and calls its “oracle” **start-run**($\text{sid}, m_{b^\dagger}$).
- When \mathcal{E} calls **get-next_A**(sid, M), if this call does not correspond to the last message of the relevant run of the message-transmission protocol, \mathcal{E}' calls its own “oracle” **get-next_A**(sid, M) and passes the response to \mathcal{E} . If the transcript of the underlying key-agreement protocol is invalid, then \mathcal{E}' sends \perp to \mathcal{E} . If it is valid and $\text{sid} = \text{sid}^*$, it sends c^* to \mathcal{E} . Otherwise, \mathcal{E}' selects a key $sk \xleftarrow{\$} \mathcal{K}$ uniformly at random, adds (sid, sk) to **keys**, finds the unique m such that $(\text{sid}, m) \in S_A$, and sends $\text{Enc}_{sk}(\text{sid}, m)$ to \mathcal{E} .
- When \mathcal{E} calls **get-next_B**(sid, M), if this is not the last message of the protocol \mathcal{E}' calls its own “oracle” **get-next_B**(sid, M) and passes the response to \mathcal{E} . If this is the last message of the protocol, it adds (sid, M) to S_B .
- When \mathcal{E} calls **get-output_B**(sid), \mathcal{E}' finds the unique M and sk such that $(\text{sid}, M) \in S_B$ and $(\text{sid}, sk) \in \text{keys}$. It computes $\text{Dec}_{sk}(\text{sid}, M)$ and responds with the result.
- When \mathcal{E} calls **get-secrets**, \mathcal{E}' responds with (σ_A, σ_B) .
- When \mathcal{E} calls **finalize**(b^*), \mathcal{E}' returns b^* .

Let b be the challenge bit in the CPA-security game against \mathcal{E}' . If $b^\dagger = b$, the view of \mathcal{E} is identical to its view in **Game 2**. In this case, \mathcal{E}' wins if and only if \mathcal{E} wins the simulated **Game 2**. If $b^\dagger \neq b$, the view of \mathcal{E} is identical to its view in **Game 3** and \mathcal{E}' wins if and only if \mathcal{E} loses the simulated game. The result follows. (4.3) ■

It should be clear that any adversary in game **Game** $q + 3$ can be easily converted into an adversary in **Game** $q + 4$ with same advantage.

Claim 4.4. *If \mathcal{W}_A is exfiltration resistant against \mathbf{B} with valid transcripts then for any PPT adversary \mathcal{E} , $|\text{Adv}^{\text{Game } q+3+i}(\mathcal{E}) - \text{Adv}^{\text{Game } q+4+i}(\mathcal{E})|$ is negligible.*

Proof. Let $\tilde{\mathbf{A}}^{(m)}$ be $\tilde{\mathbf{A}}^*$ with input plaintext fixed to m .

We construct an adversary \mathcal{E}' in LEAK (Figure 1) as follows. \mathcal{E}' first runs the **setup** procedure of the key-agreement protocol, receiving as output σ_A, σ_B , and π . As above, \mathcal{E}' will simulate many runs of the key-agreement protocol with Alice replaced by $\mathcal{W}_A \circ \mathbf{A}$ and input $(\sigma_A, \sigma_B, \pi)$. So, for convenience, we give \mathcal{E}' an “oracle interface” to these simulated runs with “oracle” calls **start-run**, **get-next_A** and **get-next_B** as in the key-agreement security game (Figure 3). \mathcal{E}' selects $b^\dagger \xleftarrow{\$} \{0, 1\}$ uniformly at random, sets $j \leftarrow 1$ and $\text{ids} \leftarrow \emptyset$, and simulates a run of \mathcal{E} , responding to oracle calls as follows.

- When \mathcal{E} calls $\text{start-run}(\text{sid}, m)$, \mathcal{E}' adds (sid, j) to ids . If $j \neq i$, it calls its own “oracle” $\text{start-run}(\text{sid}, m)$ and increments j . If $j = i$, \mathcal{E}' increments j and constructs the circuit \tilde{B} described below with sid^* , $(\sigma_A, \sigma_B, \pi)$, b^\dagger , j , ids , the state of \mathcal{E} , and the state of the various “oracles” hard-coded into it. It then returns $(\tilde{A}^{(m_{b^\dagger})}, \tilde{B}, (\sigma_A, \sigma_B, \pi))$.
- When \mathcal{E} calls $\text{get-next}_A(\text{sid}, M)$, \mathcal{E}' calls its own “oracle” $\text{get-next}_A(\text{sid}, M)$ and passes the response to \mathcal{E} .
- When \mathcal{E} calls $\text{get-next}_B(\text{sid}, M)$, \mathcal{E}' calls its own “oracle” $\text{get-next}_B(\text{sid}, M)$ and passes the response to \mathcal{E} .
- When \mathcal{E} calls $\text{start-challenge}(\text{sid}^*, m_0, m_1)$,
- When \mathcal{E} calls get-secrets , \mathcal{E}' responds with (σ_A, σ_B) .

\tilde{B} will play the role of Bob in the key-agreement protocol, and it has the state of \mathcal{E} and the “oracles” hard-coded into it. It can make “oracle” calls to simulated protocols with Alice replaced by $\mathcal{W}_A \circ A$. It also starts its own “oracle” simulations with Alice replaced by $\mathcal{W}_A \circ \tilde{A}^*$ instead. To distinguish these oracles, we write, e.g., $\text{get-next}_{\mathcal{W}_A \circ \tilde{A}^*}$ and $\text{get-next}_{\mathcal{W}_A \circ A}$. Note that \tilde{B} is itself playing a game in which it exchanges its own messages with the challenge party A^* in **LEAK** (Figure 1). \tilde{B} continues to simulate \mathcal{E} from its current state, responding to oracle calls as follows.

- When \mathcal{E} calls $\text{start-run}(\text{sid}, m)$, \tilde{B} adds (sid, j) to ids , increments j , and calls its own “oracle” $\text{start-run}(\text{sid}, m)$.
- When \mathcal{E} calls $\text{get-next}_A(\text{sid}, M)$, \tilde{B} finds the unique k such that $(\text{sid}, k) \in \text{ids}$. If $k < i$, \tilde{B} calls its “oracle” $\text{get-next}_{\mathcal{W}_A \circ A}(\text{sid}, M)$ and passes the response to \mathcal{E} . If $k = i$, then \tilde{B} sends the message M to the challenge party A^* and passes the response to \mathcal{E} . If $k > i$, then \tilde{B} calls its “oracle” $\text{get-next}_{\mathcal{W}_A \circ \tilde{A}^*}(\text{sid}, M)$ and passes the response to \mathcal{E} .
- When \mathcal{E} calls $\text{get-next}_B(\text{sid}, M)$, \tilde{B} calls its own “oracle” $\text{get-next}_B(\text{sid}, M)$ and passes the response to \mathcal{E} .
- When \mathcal{E} calls $\text{finalize}(b^*)$, \tilde{B} sets its state to 0 if $b^* = b^\dagger$ and to 1 otherwise.
- When \mathcal{E} calls get-secrets , \tilde{B} responds with (σ_A, σ_B) .

Finally, \mathcal{E}' receives the state of \tilde{B} and simply returns its value.

Let b be the challenge bit in **LEAK**. Then, if $b = 0$ so that the challenge party is honest, the view of \mathcal{E} is identical to its view in **Game** $q + 5$. Then, the final state of \tilde{B} matches b if and only if \mathcal{E} wins this simulated game. If, on the other hand, $b = 1$, then the view of \mathcal{E} is identical to its view in **Game** $q + 4$, and the final state of \tilde{B} matches b if and only if \mathcal{E} loses this simulated game. The result follows. (4.4) ■

□

F Proof of Theorem 9

Proof of Theorem 9. Let $(\text{KeyGen}, \text{Enc}, \text{Dec})$ be a CPA-secure encryption scheme with some reverse firewall \mathcal{W} . Note that we can view \mathcal{W} as a map between ciphertexts.

We present a one-bit PKE scheme as follows. Let m_0, m_1 be distinct plaintexts. The public key is then $(e_0 = \mathcal{W}(\text{Enc}_{s_k}(m_0)), e_1 = \mathcal{W}(\text{Enc}_{s_k}(m_1)))$, and the secret key is just the secret key of the underlying scheme. To encrypt a bit b , we compute $\mathcal{W}(e_b)$. The decryption algorithm of the public-key scheme runs the decryption algorithm of the symmetric-key scheme Dec and outputs 0 if the result is m_0 , 1 if it is m_1 , and \perp otherwise.

Let \mathcal{E} be a PPT adversary in the semantic-security game against the above scheme. We assume without loss of generality that \mathcal{E} never outputs a pair of identical challenge plaintexts. We construct

an efficient tampered encryption algorithm $\bar{\text{Enc}}$ and an efficient adversary \mathcal{E}' in the CPA-security game against $\mathcal{W} \circ \bar{\text{Enc}}$. Choose m_0^\dagger and m_1^\dagger uniformly at random from the plaintext space. Simulate \mathcal{E} polynomially many times and let i such in at least polynomial many of these runs, the challenge plaintexts chosen by \mathcal{E} differ in the i th bit. Fix $c_0 = \text{Enc}_{sk}(m_0)$ and $c_1 = \text{Enc}_{sk}(m_1)$. Then, we define $\bar{\text{Enc}}_{sk}(m_b^\dagger) = c_b$ and for all other plaintexts m , $\bar{\text{Enc}}(m) = \mathcal{W}(c_b)$ where b is the i th bit of m . Then, \mathcal{E}' behaves as follows.

- It calls the encryption oracle on input m_0^\dagger and m_1^\dagger . Call the results e_0 and e_1 .
- It runs \mathcal{E} with input $pk = (e_0, e_1)$, receiving as output two challenge plaintexts, (m_0^*, m_1^*) . \mathcal{E}' then outputs these as its own challenge plaintexts.
- On input c^* , a challenge ciphertext, \mathcal{E}' passes c^* to \mathcal{E} , receiving as output a bit b .
- If m_0^* and m_1^* differ in their i th bit and are distinct from m_0^\dagger and m_1^\dagger , output the bit corresponding to the plaintext whose i th bit is b . Otherwise, return a uniformly random bit.

Note that the view of \mathcal{E} is identical to its view in the semantic security game against the public-key scheme. Furthermore, with non-negligible probability, we have that m_0^\dagger , m_1^\dagger , m_0^* , and m_1^* are distinct and m_0^* and m_1^* differ in their i th bit. When both of these conditions are satisfied, \mathcal{E}' guesses correctly if and only if \mathcal{E} guesses correctly. The result follows. \square