Simultaneous Secrecy and Reliability Amplification for a General Channel Model

Russell Impagliazzo[1], Ragesh Jaiswal[2], Valentine Kabanets[3], Bruce M. Kapron[4], Valerie King[4], and Stefano Tessaro[5]

[1] University of California, San Diego russell@cs.ucsd.edu
[2] Indian Institute of Technology Delhi rjaiswal@cse.iitd.ac.in
[3] Simon Fraser University kabanets@cs.sfu.ca
[4] University of Victoria bmkapron@uvic.ca, val@cs.uvic.ca
[5] University of California, Santa Barbara tessaro@cs.ucsb.edu

Abstract. We present a general notion of channel for cryptographic purposes, which can model either a (classical) physical channel or the consequences of a cryptographic protocol, or any hybrid. We consider *simultaneous secrecy and reliability amplification* for such channels. We show that simultaneous secrecy and reliability amplification is not possible for the most general model of channel, but, at least for some values of the parameters, it is possible for a restricted class of channels that still includes both standard information-theoretic channels and keyless cryptographic protocols.

Even in the restricted model, we require that for the original channel, the failure chance for the attacker must be a factor c more than that for the intended receiver. We show that for any $c > 4$, there is a one-way protocol (where the sender sends information to the receiver only) which achieves simultaneous secrecy and reliability. From results of Holenstein and Renner (*CRYPTO'05*), there are no such one-way protocols for $c < 2$. On the other hand, we also show that for $c > 1.5$, there are two-way protocols that achieve simultaneous secrecy and reliability.

We propose using similar models to address other questions in the theory of cryptography, such as using noisy channels for secret agreement, trade-offs between reliability and secrecy, and the equivalence of various notions of oblivious channels and secure computation.

1 Introduction

Modern cryptography has its roots in the work of Shannon [35], using channels as the model of communication where some secrecy is attainable [39, 9]. A cryptographic protocol can also be interpreted as implicitly defining a *computational* channel, where the loss of information is merely computational. For example, consider a channel sending a message m as the pair consisting of a public key pk, and an encryption c of m under pk. If the encryption scheme provides some form of (even weak) security, a computationally bounded adversarial observer of the channel output will only learn partial information about m, even though information-theoretically the channel may well uniquely define its input.

In some circumstances, it may not even be clear whether the limitation is computational or informational. For example, an adversary may not be able to perfectly tune in to a low-power radio broadcast. This might appear an information-theoretic limitation, but improved algorithms to interpolate signals or to predict interference due to atmospheric conditions could also improve the adversary's ability to eavesdrop.

In this work, we introduce a model of computation that combines information-theoretic and computational limitations. Specifically, we present a general notion of channel for cryptographic purposes, which can model either a (classical) physical channel or the consequences of a cryptographic protocol, or any hybrid.

We require our model to satisfy the following properties:

- [**Agnostic**] It should not matter *why* an adversary is limited. Protocols designed exploiting an adversary's weakness should remain secure whether that weakness is due to limited information, computational ability, or any other reason.
- [**Composable**] We should be able to safely combine a protocol that achieves one goal from an assumption, and a second protocol that achieves a second goal from the first, into one that achieves the second goal from the original assumption.
- [**Functional**] The assumptions underlying our protocols should concern what the parties *can do*, rather than concerning what they or the channels through which they communicate *are*. In particular, we should be able to use this to evaluate the danger of side information, and enhanced functionality should not threaten secrecy properties.
- [**Combining reliability and secrecy**] Instead of viewing reliability of a channel and its secrecy as separate issues, our model should combine the two in a seamless way. We want to study how enhancing secrecy might impact reliability, and vice versa. In other words, we view reliability as equally necessary for the overall secrecy.

In this paper, we focus on the *simultaneous secrecy and reliability amplification* for such channels. We start with a channel where the intended receiver gets the transmitted bit except with some probability and the attacker can guess the transmitted bit except with a somewhat higher probability. We wish to use the channel to define one where the receiver gets the transmitted bit almost certainly while only negligible information is leaked to the attacker. We show that simultaneous secrecy and reliability amplification is not possible for the most general model of channel, but, at least for some values of the parameters, it is possible for a restricted class of channels that still includes both standard information-theoretic channels and keyless cryptographic protocols.

Note that, traditionally, error-correction and encryption have been thought of in communications theory as separate layers, with one performed first and then the other on top. However, when one wants to leverage the secrecy of an unreliable channel, it does not seem possible to separate the two. Using an error-correcting code prior to secrecy considerations could totally eliminate even the partial secrecy, and amplifying secrecy could make the channel totally unreliable.

(In some sense, our solution alternates primitive error-correction stages with secrecy amplification stages, but we need several rounds of each nested carefully.)

1.1 Our results

We propose a very general model of channel with state, which makes few assumptions about the way the channel is constructed or the computational resources of the users and attackers. In the present paper, such a channel is used for communication between Alice and Bob, with an active attacker Eve. The channel has certain reliability and secrecy guarantees, ensuring that Bob receives a bit sent to him by Alice with sufficiently higher probability than Eve (see Section 2).

We show (in Section 3) how secrecy and reliability of such channels can be simultaneously amplified with efficient protocols (using one-way communication only), provided that the original channel has a constant-factor gap (at least 4) between its secrecy and reliability (i.e., Eve is 4 times more likely to make a mistake on a random bit sent by Alice across the channel than Bob is on any given bit sent by Alice). We prove (in Section 4) that some constant-factor gap (the factor 2) is necessary for any one-way protocol. Finally, we present (in Section 5) an efficient two-way communication protocol for amplifying secrecy and reliability, assuming the original channel has the factor 1.5 gap between secrecy and reliability.

For our one-way protocol in Section 3, we tighten a result of Halevi and Rabin [16] on the secrecy analysis of a repetition protocol. If the eavesdropper has probability at most $1-\alpha$ of guessing a bit sent across the channel from Alice to Bob, then the eavesdropper has probability at most $1-(2\alpha)^n/2$ of learning the bit, if this bit is sent across the channel n times. This improves upon the analysis of [16], who showed $1-\alpha^n$ probability for the eavesdropper.

Our two-way protocol in Section 5 applies to secret-key agreement between two parties both in the information-theoretic and complexity-theoretic setting, extending the results of Holenstein and Renner [19] on one-way protocols.

1.2 Related work

Our results exhibit both technical and conceptual similarities with the rich line of works on secrecy amplification for cryptographic primitives and protocols. A number of them developed amplification results for both soundness and correctness of specific two-party protocols [1, 32, 33, 37, 16, 15, 17, 4, 20, 5]. Different from our work, however, these consider settings where one of two parties is corrupt, and secrecy for the other party is desired. Here, we envision a scenario with two honest parties, Alice and Bob, communicating in presence of a malicious third party, Eve. Previously, this was only considered in works on secrecy and correctness amplification for public-key encryption and key agreement [11, 18, 19, 26]. We note that our framework is far more general than these previous works.

Following Shannon's impossibility result showing that perfect secrecy requires a secret key as large as the plaintext [35] (see also [10]), there has been a large body of research in information-theoretic cryptography. This line of work shows that perfect secrecy is possible, if one assumes that physical communication channels are noisy. One such model of a noisy communication channel is Wyner's wiretap channel of [39], generalized by [9], and extensively studied since (see [25] for a survey). A number of both possibility and impossibility results were shown for various models of noisy channels, see, e.g., [7, 29, 30, 6, 31, 8, 38, 21, 12].

Different formalizations of secrecy in the information-theoretic setting were studied by [2, 36, 22, 23]. In particular, Bellare et al. [2] consider the wiretap channel and relate the information-theoretic notion of secrecy (traditionally used in information-theoretic cryptography) to the semantic secrecy in the spirit of [14] (used in complexity-theoretic cryptography).

We remark that in the information-theoretic approach to cryptography, the focus is usually on what the channel *is*: for example, a channel between Alice and Bob, with eavesdropper Eve, is modeled as a triple of correlated random variables A, B, E, with certain assumptions on the joint distribution of these variables. Then the question is studied what such a channel can be used for, and how efficiently (e.g., at what rate). In contrast, our main focus is on the *utilization* of the channel, i.e., what the channel can be used for. For example, if a channel can be used for somewhat secret and reliable transmission of information, we would like to know if that channel can be used to construct a new channel for totally secret and reliable transmission.

Below we provide a more detailed comparison between our work and the most closely related previous work.

Comparison with [19]. Perhaps the most closely related to the present paper is the work by Holenstein and Renner [19] that considers the task of secret-key agreement in the information-theoretic setting, where two honest parties, Alice and Bob, have access to some correlated randomness such that the eavesdropper, Eve, has only partial information on that randomness. In particular, [19] consider a special case where the random variables of Alice and Bob, A and B, are binary and have correlation at least α (i.e., A and B are equal with probability at least $(1 + \alpha)/2$), whereas with probability at least $1 - \beta$, the random variable E of Eve contains no information on A. One of the main results of [19] shows that secret key agreement, using one-way communication from Alice to Bob, is possible when $\alpha^2 > \beta$, and impossible otherwise. Holenstein and Renner also observe that one-way secret-key agreement for such random variables is equivalent to the task of black-box circuit polarization, introduced by Sahai and Vadhan [34] in the context of statistical zero knowledge. The impossibility result for one-way secret-key agreement in [19] implies that the parameters for circuit polarization achieved by Sahai and Vadhan [34] are in fact optimal for such black-box protocols.

The setting of binary random variables A, B, E in [19] is similar to the channel model we consider. Their condition on A and B being correlated corresponds to channel's reliability, and the condition on E sometimes having no information on A corresponds to channel's secrecy. We use the impossibility result of [19] (almost directly) to argue the need of a constant-factor (factor 2) separation between reliability and secrecy of channels for the case of one-way protocols. However, our one-way channel protocol (for the case of factor 4 separation between reliability and secrecy) is for a more general, not necessarily information-theoretic, setting. Moreover, we go beyond the one-way communication, and describe an efficient two-way protocol that works for the case where the constant-factor gap between reliability and secrecy of a channel is smaller (factor 1.5) than the gap required by one-way protocols. This yields a new protocol that works both for the information-theoretic setting (as in [19]), and for the complexity-theoretic setting, using the results of [18].

Comparison with [30]. Maurer [30] considered the information-theoretic setting of a channel between Alice and Bob, with eavesdropper Eve, where the channel from Alice to Bob is symmetric noisy channel with the noise parameter ϵ, and the channel from Alice to Eve is an independent symmetric noisy channel with the noise parameter δ. Using the earlier work by [9], Maurer shows that Alice and Bob can securely agree on a secret in this setting, provided $\epsilon < \delta$. Surprisingly, Maurer also shows that secret-key agreement between Alice and Bob is still possible even if $\epsilon \geq \delta$, by using a two-way protocol (where Bob also sends messages to Alice over the public channel)! Like Maurer, we also use a two-way protocol to overcome the limitations of one-way protocols. The difference is that our setting is more general than his information-theoretic setting (of two independent noisy channels). For example, in Maurer's setting, it is easy to see that Eve has less information than Alice about the bit Bob receives, which is not always true in our setting (unless $\alpha > 2\beta$). However, his results raise the question of what additional reasonable conditions on our channel model could be used to reduce the gap between secrecy and reliability that one needs to assume. One natural condition is that Eve has a small probability of learning a random bit sent from Bob to Alice (in addition to the existent secrecy condition that Eve has small probability of learning a random bit sent from Alice to Bob). We leave the study of this channel model with "symmetric secrecy" for future research.

Comparison with [27]. The framework of *constructive cryptography* by Maurer [27] also deals with reductions between channels, using the formalism from the abstract cryptography framework [28]. In constructive cryptography, the main goal is to capture traditional security goals (like secrecy and authenticity) in terms of channel transformations. Contrary to our framework, channels in constructive cryptography are described *exactly* through ideal functionalities, in the same spirit as in Canetti's UC framework [3]. Maurer's framework in fact also allows the definition of *classes* of channels (as we consider here), but this feature appears to be mostly definitional, as we are not aware of any results that would apply to the context of our work.

1.3 Our techniques

We use fairly standard tools such as the direct-product and XOR protocols, relying on the proof techniques in [24, 13]. We also use the repetition protocol, whose secrecy in the cryptographic setting was first analyzed in [16]. We generalize and improve their analysis (see Theorem 14), getting better secrecy ($(2\alpha)^n/2$ instead of α^n), which is crucial for our applications. While the techniques we ended up using in this paper are standard, finding the right techniques to use for our applications was nontrivial, and involved considering many other standard techniques that turned out to be inapplicable to our setting. For example, error-correcting codes are an obvious approach to amplifying reliability. But it is still very unclear how such codes affect secrecy. Also, many of the ways we apply standard techniques are delicate. The XOR protocol we use is standard, but fails dramatically if one reverses the order in which the messages are sent. There seems to be a subtle and intricate interplay between the contradictory requirements of secrecy and reliability that we want to achieve simultaneously.

2 The model and axioms

2.1 Channels

The following is a definition of a one-way channel that communicates information from a user Alice to a user Bob. An attacker Eve is capable of launching possibly active attacks, and can gain some information about communicated messages. We can generalize such a channel to one allowing two-way communication or multi-party channels. Note that while we do capture a variety of classical physical systems with this definition, we do not necessarily capture quantum channels or protocols, because we assume that computation does not change the system's state. We could generalize further, but it's already getting pretty complicated.

Definition 1 (Channel). *A* one-way channel *from user Alice to user Bob with attacker Eve has the following components:*

1. Security parameter: $k \in \mathbb{N}$*;*
2. States: *for each k, a countable set of possible* underlying states*, $\Sigma_k \subseteq \{0,1\}^*$;*
3. Attacks: *for each k, a countable set of possible* attacks *$\Gamma_k \subseteq \{0,1\}^*$;*
4. Transition function: *for each k, a* probabilistic transition function *δ_k which takes as input the current state $s \in \Sigma_k$, an* attack *$\gamma \in \Gamma_k$ from Eve, and a* transmitted bit *b from Alice, and produces a probability distribution $\delta_k(s, \gamma, b)$ on the* updated state *$s' \in \Sigma_k$ and* received message *$b' \in \{0,1\}$;*
5. Eve's view function: *a function $v_E(s)$ from states to strings, giving the* visible part of the state *for Eve;*
6. Resource limits: *a set F of probabilistic functions from strings to strings, computable within the* computational limits *of the adversary. We assume F is closed under polynomial-time (in the lengths of strings and the secrecy parameter) Turing reductions, and under fixing as advice any single bit, visible state or action.*[6]

[6] If a channel is such that the state description rapidly grows (say, squares) after each use, then after very few uses, the adversary that is allowed polynomial time in the

Remark 2. For our application of secret and reliable information transmission from Alice to Bob in the presence of an active evesedropper Eve, we can assume that Alice and Bob, as trusted parties, do not need to keep track of the channel state. This simplifies our definition of channel above. However, for other tasks (e.g., Oblivious Transfer, bit flipping over the phone, secure multiparty computation), we need to include in our model Alice's and Bob's view functions of the channel state, $v_A(s)$ and $v_B(s)$, respectively. This would match the standard information-theoretic view of such a channel as a triple of correlated random variables A (for Alice), B (for Bob), and E (for Eve).

Our main results only apply to limited classes of channels that we call *transparent* and *semi-transparent.*

Definition 3 (Transparency). *A channel of Definition 1 is called* transparent *if it satisfies the following additional properties:*

- $v_E(s) = s$ *(i.e., all of the state is visible to the attacker), and*
- *for every* $k \in \mathbb{N}$, $\delta_k \in F$ *(i.e., the attacker can simulate the channel).*

A channel of Definition 1 is called semi-transparent *if it satisfies the following additional properties:*

- $v_E(s) = s$ *(i.e., all of the state is visible to the attacker), and*
- *for every* $k \in \mathbb{N}$, *computing the new state under* δ_k *is in* F *(i.e., the attacker can simulate the channel as far as the information they get, but not necessarily the output).*

Remark 4. The utility of transparency condition on the channel is that it enables the eavesdropper Eve to simulate the channel forward, by taking control of a virtual Alice. In fact, as was pointed out to us by Daniele Micciancio [personal communication, 2015], given an arbitrary channel that can be simulated forward, one can define a new, equivalent channel that is transparent; the converse is also true. So transparency is equivalent to being simulatable forward.

Transparent channels include any memoryless channel with computationally unbounded (information-theoretic) attackers, and any two-party protocol where there are no secret inputs for either party before the protocol starts.

Definition 5 (α-Secrecy and β-Reliability). *Let* $1/2 > \alpha > \beta \geq 0$ *be constants (or functions of the security parameter). A channel is called* α-secret *and* β-reliable *if it satisfies the following axioms:*

size of the state will get to use exponential-time computation for her attacks. A standard cryptographic channel will unlikely be secure in this case. However, it is up to the designer of the channel to ensure that it remains secure, with respect to polynomial-time adversaries (which will probably force the designer to make sure that the state description does not grow too fast with respect to k).

- **Secrecy Axiom:** *For all but finitely many* $k \in \mathbb{N}$*,* $\forall f \in F$*,* $\forall s \in \Sigma_k$*,* $\forall \gamma \in \Gamma_k$*, and for* $b \in_U \{0,1\}$ *uniformly chosen,*

$$\Pr_{(s',b')=\delta_k(s,\gamma,b)}[f(v_E(s')) = b] \leq 1 - \alpha.$$

- **Reliability Axiom:** $\forall k \in \mathbb{N}$*,* $\forall s \in \Sigma_k$*,* $\forall \gamma \in \Gamma_k$*, and* $\forall b \in \{0,1\}$*,*

$$\Pr_{(s',b')=\delta_k(s,\gamma,b)}[b' = b] \geq 1 - \beta.$$

These conditions are met by the (non-transparent) channel that works as follows. Initially the state is the empty string. The intended receiver always gets the sent bit. The eavesdropper is allowed exponential computation time, and has two attacks: “defer” or “break”. If “defer” is chosen, the eavesdropper learns nothing at the time (the visible state contains no bits), but the current bit sent is appended to the channel state. If “break” is chosen, with probability $1 - 2\alpha$, the channel state is updated as normal but becomes visible to the eavesdropper; with probability 2α, the channel state is erased (becomes the empty string).

The first example provably shows that secrecy amplification cannot be based solely on the above axioms. Consider any protocol to send a bit secretly from Alice to Bob, using the channel above. Eve can use the strategy of using “defer” until the last bit is sent, and attacking the last bit with “break”. With probability $1-2\alpha$, Eve learns the entire conversation between Alice and Bob. By simulating all possible random choices used by Alice and Bob, and seeing which ones are consistent with the conversation, Eve can learn the secret.

To see where non-transparency could actually prevent secrecy amplification in the cryptographic setting, consider a channel that simulates the following private-key protocol. Alice and Bob share a secret key κ, and to send a message, Alice sends $E_\kappa(m)$ and a weak commitment $C(\kappa)$ to Bob. If an eavesdropper can break the secrecy of the commitment scheme with some small probability α, then no matter how the scheme is used repeatedly and combined, the attacker will learn the key with probability at least α. In general, protocols that assume prior shared information such as a private key will not be transparent, because the attacker cannot simulate a run of the protocol without this shared information.

We will show that for *transparent* channels this problem does *not* arise.

2.2 Examples

We give some examples of both channels in an information-theoretic setting and computational setting. Our results hold for channels that are some hybrid of the two as well, but these two extremes are the most familiar, so will serve as intuition. In general, we'll be using complexity-theoretic methods when proving possibility results, and prove impossibility results using information-theoretic means, so we will be shifting back and forth between the two.

Information-theoretic channels

Noise vs. erasure One interesting channel is a joint symmetric binary noise and erasure channel, where, when Alice sends b, Bob receives the bit b' which is equal to b with probability $1-\beta$ and equal to $1-b$ otherwise. Eve receives (i.e., the new state equals) the bit b with probability $1-2\alpha$ and the message $\perp$ otherwise. [7] There might or might not be correlation between Eve's erasures and Bob's noise. The channel is memoryless, in that the current state does not actually affect the transition function. Any memoryless channel is equivalent to a transparent one in the information-theoretic setting, since we might as well replace the state with the visible state and Eve can always simulate the fixed transition function.

Noise attacks An active Eve might be able to control the noise of the channel, but not gain any information about the bit sent. For example, say attacks are numbers γ between 0 and β. Bob receives a bit b' with binary symmetric noise γ, and Eve receives (i.e., the new state is) $b' \oplus b$, whether or not Bob got the bit sent. This channel gives Eve no information about the bit sent, but allows her to attack reliability. Again it is memoryless, hence transparent.

Arbitrary memoryless channels We can embed conventional results about secrecy capacity of channels in our model. Consider any fixed distribution on triples (A, B, E), where we view a single use of a device as giving Alice information A, Bob, B and the attacker E, and Alice and Bob can communicate in the clear as well. Using the device K times gives a sequence of K values of these variables $A_1, ..., A_K, B_1, ..., B_K$, and $E_1, ..., E_K$ from the same joint distribution. At some point, after using the device and sending some messages, Bob will output a guess as to the bit Alice meant to send him. The new state would be the K tuple of values $E_1, \ldots, E_K$, and the messages sent in the clear. While the sequence A and B are used, and help determine the output, we don't include them in the state (because they will not be used in future transmissions), and since Alice and Bob are trusted participants, there is no reason to keep track of their side information, rather than just the secret they agree on. The system is memoryless, and hence transparent.

Complexity-theoretic channels

Private key encryption If Alice and Bob use a secret key and send messages using a private key encryption, then the state would be both the key and the messages sent in the clear, but the visible state for Eve would just be the messages sent in the clear. So this type of protocol is not transparent, since including the key in the visible state would render it useless.

Noisy trapdoor function with fixed public key Say Bob creates a trapdoor function with probabilistic encryption and noisy decryption, and Alice always sends bits with Bob's fixed public key. Then the state of the channel

[7] Note that Eve can guess the bit with probability 1/2 when she receives $\perp$. So the probability of her knowing the bit b is $1 - 2\alpha + (1/2) \cdot (2\alpha) = 1 - \alpha$.

is the public key and the encryption of the bit sent. This channel is semi-transparent, because Eve can simulate the new state (only the encryption of the bit is changed), but cannot necessarily simulate whether Bob will get the bit correctly without Bob's secret key. If there is feed-back from Bob to Alice, Eve might be able to simulate a chosen cyphertext attack on the encryption function.

Noisy trapdoor function with fresh public keys On the other hand, using the same encryption function but with a fresh key every message, the channel becomes fully transparent. Eve can simulate the channel and Bob's received bit by generating her own keys and using them. Chosen cyphertext attacks become a non-issue, so protocols using feedback are fine.

2.3 Virtual channels and protocol channels

A protocol using a channel defines a new, *virtual channel.* The inputs to this virtual channel are strategies for the participants and attacker, using the old channel. The virtual channel's states accumulate the protocol history, that is the the sequence of observable states during the protocol, together with any messages sent in the clear. The transition function simulates the protocol with the given strategies to obtain the history.

A *protocol channel* fixes the inputs from Alice and Bob in the virtual channel to specific strategies of Alice and Bob.

Definition 6 (Amplifying secrecy and reliability). *For $\alpha' > \alpha > \beta > \beta'$,* secrecy and reliability amplification from (α, β) to (α', β') *means defining a protocol which guarantees that, for any (transparent) channel satisfying α-secrecy and β-reliability, the protocol channel satisfies α'-secrecy and β'-reliability.*

We note that by construction, states of a protocol channel have the same degree of visibility as states of the underlying channel. Furthermore, since transitions of the protocol channel simulate the strategies of the participants, we conclude the following.

Lemma 7. *If a channel is transparent, and the legitimate users' strategies are in F, then the protocol channel is also transparent, regardless of whether the protocol uses one-way or two-way communication. If a channel is semi-transparent, and the legitimate users' strategies are in F, then the protocol channel is also semi-transparent, provided that the protocol uses one-way (from Alice to Bob) communication only.*

Thus, protocol constructions or secrecy and reliability amplifications which assume the axiom of transparency will always be *composable.* In other words, we can have a series of protocols built on top of channels. The protocols will only utilize the channels as black boxes and so not require any knowledge of how the underlying channel works. They will have the property that if the channel is transparent, α-secret and β-reliable, then the protocol is α'-secret and β'-reliable. Then we can use the protocol as the channel in any way of converting α'-secret and β'-reliable channels into α''-secret and β''-reliable ones. The same is true also for *one-way* protocols using *semi-transparent* channels.

3 Secrecy and reliability amplification for one-way protocols

The main result of this section is the following.

Theorem 8. *For any non-negligible ϵ and any $1/2 > \alpha > 4\beta > 0$, there is a one-way protocol for secrecy and reliability amplification from (α, β) to $(1/2-\epsilon, 2^{-k})$.*

The required protocol will rely on the Direct-Product protocol, the Parity protocol, and the Repetition protocol that we discuss next.

3.1 Direct-product protocols

The direct product is one of the fundamental constructions in complexity and the theory of cryptography. Direct product theorems state that if one instance of a problem is unlikely to be solved , then two independent instances are even less likely to be both solved. There are many proofs of direct product theorems that apply to a wide variety of models and circumstances. Modern proofs utilize connections to coding theory, hard-core sets, and so on. However, these proofs do not seem to work in our setting. What does work is one of the oldest techniques in direct products, estimates of conditional probabilities, used, for example, by Levin [24].

Direct product constructions generally decrease reliability but enhance secrecy. The simplest direct product constructions just concatenate the various solutions. We'll analyze such a protocol, but it will not be immediate how to translate the result about concatenating secrets into one where the secrets are combined into a single bit.

Consider the following **Direct-Product Protocol**:

> Alice sends n independent random bits $b_n, \ldots, b_1$ (we number them in reverse order to make an inductive argument cleaner) through the channel.

We compare the probability that Bob receives all n bits with the probability that Eve can guess all n bits. First, for Bob's probability of receiving all n bits, we can use that the reliability axiom holds for each state of the channel. Conditioned on any event for the first i bits, and in particular, conditional on Bob receiving the first i bits correctly, the probability of his receiving the ith bit correctly is at least $1 - \beta$. Therefore, the probability that he receives all n bits correctly is at least $(1 - \beta)^n$.

Next, we use the method of conditional probabilities, due to Levin, to bound the probability that Eve can guess all n bits.

Theorem 9 (Direct-Product Theorem for Channels). *For any non-negligible function ϵ of the secrecy parameter, and any polynomially bounded n, the probability that Eve can guess all n bits is at most $(1 - \alpha)^n + n\epsilon$.*

Proof. Consider the distribution on the information available to Eve by an attack. An attack on the protocol will be determined by two functions A which receives a list of states and determines the next attack a on the channel, and f which after the protocol ends outputs the guess $B_n ... B_1$. The protocol under this strategy will evolve as follows:

1. The protocol starts in some state s_{n+1}. Let the initial history H_{n+1} be the list containing only s_{n+1}.
2. For each i from n to 1:
   (a) Alice picks a random bit $b_i \in \{0, 1\}$.
   (b) Eve picks channel attack $a_i = A(H_{i+1})$.
   (c) The new state s_i and the bit b'_i received by Bob are given by $(s_i, b'_i) = \delta_k(s_{i+1}, a_i, b_i)$.
   (d) Append s_i to H_{i+1} to get an updated history H_i.
3. Eve guesses $B_n, \ldots, B_1 = f(H_1)$.

Note that given any H_i, Eve can simulate the rest of the process to produce H_1 according to the correct conditional distribution, using randomly generated bits $b_{i-1}, \ldots, b_1$ (since $\delta_k \in F$). (This is where we use transparency.)

Let Success_i be the event that $B_i = b_i, \ldots, B_1 = b_1$. The theorem will follow from the next claim for $i = n$.

Claim. For any $1 \leq i \leq n$ and history H_{i+1}, $\Pr[\text{Success}_i \mid H_{i+1}] \leq (1-\alpha)^i + i\epsilon$.

Proof (of Claim). Our proof is by induction on i. The $i = 1$ case is just the secrecy property of the channel at state s_2. Fix H_{i+1}. Consider the following attack on a single bit b_i sent on the channel at state s_{i+1}:

> Eve uses attack a_i, bit b_i is sent by Alice, and the channel arrives in state s_i. Then she repeatedly simulates the conditional distribution on histories starting from H_i as given above, until either Success_{i-1} or the number of simulations reaches $T = (1/\epsilon)\ln(1/\epsilon)$. If the former, she outputs B_i as her guess for b_i, otherwise, the simulations time out without success, she outputs no guess.

By transparency of the channel and its α-secrecy, we get that

$$\Pr[B_i = b_i \mid H_{i+1}] \leq (1 - \alpha). \tag{1}$$

Next, $\Pr[B_i = b_i \mid H_i]$ is $\Pr[\text{Success}_i \mid H_i, \text{Success}_{i-1}]$ times the probability of not timing out, which is $1-(1-\Pr[\text{Success}_{i-1}|H_i])^T$. In particular, if $\Pr[\text{Success}_i \mid H_i] \geq \epsilon$, so is $\Pr[\text{Success}_{i-1} \mid H_i]$ and the probability of not timing out is at least $1 - (1-\epsilon)^T \geq 1 - e^{-\epsilon T} = 1 - \epsilon$ by our choice of T. Then

$$\begin{aligned} \Pr[B_i = b_i \mid H_i] &\geq \frac{\Pr[\text{Success}_i \mid H_i]}{\Pr[\text{Success}_{i-1} \mid H_i]} - \epsilon \\ &\geq \frac{\Pr[\text{Success}_i \mid H_i]}{(1-\alpha)^{i-1} + (i-1)\epsilon} - \epsilon, \end{aligned}$$

where the last inequality is by the induction hypothesis applied to H_{i-1}. So we get

$$\Pr[\text{Success}_i \mid H_i] \leq (1-\alpha)^{i-1} \cdot \Pr[B_i = b_i \mid H_i] + i\epsilon. \tag{2}$$

If $\Pr[\text{Success}_i \mid H_i] < \epsilon$, then Eq. (2) holds for trivial reasons. Finally, averaging over H_i in Eq. (2) and then using the inequality of Eq. (1), concludes the proof.

3.2 Parity protocols

Next we want to use our direct-product protocol to get a single bit message across the channel. Before showing a protocol that works (under some circumstances), we give an illuminating example of a tempting protocol that fails.

Naive parity protocol Consider the naive parity protocol for sending a bit b from Alice to Bob:

> Alice sends random bits $b_n, \ldots, b_1$ as above, and then sends $b \oplus b_n \oplus \cdots \oplus b_1$. Bob's guess at b is the parity of all the bits he receives.

We are not sure whether this protocol boosts secrecy, but it actually fails miserably when it comes to reliability. In fact, there are channels where this protocol is much worse than random guessing from Bob's point of view!

Theorem 10. *For any $1/2 > \beta > 0$, there is a transparent $1/2$-secret and β-reliable channel such that the naive parity protocol above yields the protocol channel with reliability $1-(1-\beta)^n$.*

Proof. Indeed, consider a channel where Eve decides whether each bit is sent with symmetric noise β or with no noise, and learns nothing about the bit sent, only the noise. In other words, the channel has two states, 0 and 1, and there are two attacks, 0 and 1. A coin η of bias β is flipped by the channel, and the new state is η (regardless of the bit sent or the attack). The bit received by Bob is $b \oplus a\eta$, i.e., is flipped if Eve picks attack 1 and the noise is 1, and is not otherwise. One can think of Alice and Bob as communicating by low power radio, and Eve can make the channel noisy by broadcasting at the same time, but can only tell if she disrupted the signal, not what the message was.

This channel has secrecy $1/2$ and β-reliability. But if Alice and Bob use the parity protocol, Eve can use attack 1 (keep the channel noisy) until $\eta = 1$, and then set $a = 0$ after that. Bob only gets the correct bit if η is never 1, so with probability $(1-\beta)^n$.

So the reliability of the naive parity protocol goes totally out the window!

Modified parity protocol Next we show a modification of this protocol that amplifies secrecy of a given channel, albeit at the price of possibly worsening its reliability somewhat. This will be later combined with another protocol that will significantly improve reliability while somewhat worsening secrecy. By carefully choosing the parameters of the protocols in this combination, we will be able to achieve both secrecy and reliability amplification for a given α-secret and β-reliable channel, provided that $\alpha > 4\beta$.

The modified parity protocol sends the parity of a random subset of bits $b_n, \ldots, b_1$, rather than all of them. Consider the **Parity Protocol**:

> To send a given bit b to Bob, Alice uses the channel to send random bits $b_n, \ldots, b_1$, and then, in the clear, sends random bits $r_n, \ldots, r_1$, followed by $b \oplus (\oplus_{i=1}^n b_i r_i)$. Bob receives bits $b'_n, \ldots, b'_1$ through the channel, and outputs $(b \oplus (\oplus_{i=1}^n b_i r_i)) \oplus (\oplus_{i=1}^n b'_i r_i)$.

Theorem 11. *Given any α-secret and β-reliable transparent channel, the Parity Protocol above yields the protocol channel that is α'-secret and β'-reliable for $\alpha' \approx (1 - e^{-\alpha n/2})/2$ and $\beta' \approx (1 - e^{-\beta n})/2$.*

Proof. The probability that Bob receives all n bits is $(1 - \beta)^n$, and then he correctly recovers b with probability 1 over the choice of random bits $r_n, \ldots, r_1$. Otherwise, Bob's string $b'_n \ldots b'_1$ is different from the string $b_n \ldots b_1$, but the two strings have the same inner product modulo 2 with the random string $r_n \ldots r_1$, with probability $1/2$ over the choice of $r_n, \ldots, r_1$. Thus, Bob's overall chance of guessing b correctly is $(1 + (1 - \beta)^n)/2$, which means that the protocol is about $(1/2)(1 - e^{-\beta n})$-reliable.

On the other hand, if Eve can guess b with conditional probability $1/2 + \gamma_{\boldsymbol{b}}$ after $\boldsymbol{b} = b_n, \ldots, b_1$ are sent, using the algorithm of Goldreich and Levin [13], varying over choices of bits $\boldsymbol{r}$, she can guess the entire vector $\boldsymbol{b}$ with probability $c \cdot \gamma_{\boldsymbol{b}}^2$, for some constant $c > 0$. Set $\gamma = \mathbf{Exp}_{\boldsymbol{b}}[\gamma_{\boldsymbol{b}}]$. We conclude that if Eve can guess b with probability $1/2+\gamma$, then she can recover the entire $\boldsymbol{b}$ with probability at least $c \cdot \mathbf{Exp}_{\boldsymbol{b}}[\gamma_{\boldsymbol{b}}^2]$, which by Jensen's Inequality is at least $c \cdot (\mathbf{Exp}_{\boldsymbol{b}}[\gamma_{\boldsymbol{b}}])^2 = c \cdot \gamma^2$.

Finally, using the Direct-Product Theorem for Channels, Theorem 9, we must have $c \cdot \gamma^2 \leq (1 - \alpha)^n + n\epsilon$ for any non-negligible ϵ, or $\gamma \leq \sqrt{c} \cdot (1 - \alpha)^{n/2} + \epsilon'$ for any such ϵ'. So secrecy is roughly $1/2(1 - e^{-\alpha n/2})$.

While both secrecy and reliability in the above protocol are close to $1/2$, a multiplicative difference in α vs. β has become an exponent in the advantage over random guessing, with the factor of 2 lost in the process.

Remark 12. Note that order matters in the protocol. Although sending $b_n, \ldots, b_1$ then $r_n, \ldots, r_1$ is the same information as sending r first then b, the reverse order would be subject to the same attack as the naive parity protocol above.

3.3 Repetition protocol

Here we get a protocol for improving reliability. It is the following **Repetition Protocol**:

> To transmit a given bit b to Bob, Alice sends this b over the channel n times. Bob takes the majority value of the received bits.

This protocol is somewhat dual to direct product: here reliability is enhanced at the price of secrecy dropping substantially. In fact, it is not clear that any secrecy would remain. In the cryptographic setting, Halevi and Rabin [16] showed that at least α^n secrecy remains. We generalize and improve their result, showing that the repetition protocol has at least $(2\alpha)^n/2$ secrecy.

First, we analyze reliability using familiar probabilistic tools.

Theorem 13. *The Repetition Protocol applied to a β-reliable channel yields a channel with reliability $\beta' \leq e^{-(1-2\beta)^2 n/8}$.*

Proof. We need to show that, for any attack on the Repetition Protocol over a β-reliable channel, the probability that Bob fails to output b is at most $e^{-(1-2\beta)^2 n/8}$. Let $b'_n, \ldots, b'_1$ be the bits received by Bob. Look at the quantity that adds β each time bit $b'_i = b$ and subtracts $(1-\beta)$ if the bit received is incorrect. By the definition of β-reliability, this quantity is a sub-martingale, with the difference bounded by 1. Bob only returns the wrong bit if there are more incorrect bits received than correct bits, in which case this quantity is at most $\beta n/2 - (1-\beta)n/2 = -(1-2\beta)n/2$. By Azuma's inequality, the probability of this is at most $e^{-((1-2\beta)n/2)^2/(2n))}$, as claimed.

Next we show:

Theorem 14. *For any parameters α and n (with n polynomially bounded in the security parameter, and $(2\alpha)^n$ non-negligible), the n-bit Repetition Protocol over an α-secret transparent channel has secrecy at least $(2\alpha)^n/2$.*

Proof. As in the proof of Theorem 9, fixing functions A and f that describe Eve's attack, the process can be described as follows:

1. Alice picks a random bit r (to be sent over the channel n times).
2. The protocol starts in some state s_{n+1}. Let the initial history H_{n+1} be the list containing only s_{n+1}.
3. For each i from n to 1:
   (a) Eve picks channel attack $a_i = A(H_{i+1})$.
   (b) The new state and bit Bob receives is $(s_i, b'_i) = \delta_k(s_{i+1}, a_i, r)$.
   (c) Append s_i to H_{i+1} to get an updated history H_i.
4. Eve guesses $R = f(H_1)$.

Consider starting from partial history H_{i+1}, picking a new random bit r_1 and simulating the protocol from then on sending r_1 for the i remaining bits to be sent. The theorem will follow from the next claim when $i = n$.

Claim. For every $1 \leq i \leq n$, $\Pr[R \neq r_1 \mid H_{i+1}] \geq (2\alpha)^i/2$.

Proof. The proof is by induction on i. For $i = 1$, this is exactly the definition of α-secrecy. Consider the following attack on a single bit r_1 sent on the channel at state s_{i+1}:

> Eve uses attack a_i and r_1 is sent by Alice, and the channel arrives in state s_i. Then she picks a new random bit r_2 and simulates the repetition protocol starting from H_i, with Alice sending r_2 each time. If the simulation returns an $R \neq r_2$, Eve guesses R. Otherwise, Eve repeats the simulation for a fresh random bit r_2. (Note that the expected number of repetitions is at most $2(2\alpha)^{-i}$, by the induction hypothesis, which is feasible by assumption).

By α-secrecy, the described strategy must fail with probability at least α, i.e.,

$$\Pr[R \neq r_1 \mid R \neq r_2, H_{i+1}] \geq \alpha. \tag{3}$$

Now fix any history H_i and bit r_1. For the R returned by Eve in the above strategy, the probability that $R \neq r_1$ is the conditional probability

$$\Pr[R \neq r_1 \mid R \neq r_2, H_i] = \frac{\Pr[R = \neg r_1 = \neg r_2 \mid H_i]}{\Pr[R \neq r_2 | H_i]}.$$

By induction, for each H_i the denominator of this expression is at least $(2\alpha)^{i-1}/2$. So for each H_i and r_1, we have

$$((2\alpha)^{i-1}/2) \cdot \Pr[R \neq r_1 \mid R \neq r_2, H_i] \leq \Pr[r_1 = r_2, R \neq r_1 \mid H_i].$$

Averaging both sides over H_i, we get

$$((2\alpha)^{i-1}/2) \cdot \Pr[R \neq r_1 \mid R \neq r_2, H_{i+1}] \leq \Pr[r_2 = r_1, R \neq r_1 \mid H_{i+1}]. \tag{4}$$

Finally, applying Eq. (3) to the left-hand side of Eq. (4), we get

$$\begin{aligned}((2\alpha)^{i-1}/2) \cdot \alpha &\leq \Pr[r_2 = r_1, R \neq r_1 \mid H_{i+1}] \\ &= \Pr[r_2 = r_1] \cdot \Pr[R \neq r_1 \mid r_2 = r_1, H_{i+1}] \\ &= (1/2) \cdot \Pr[R \neq r_1 \mid r_2 = r_1, H_{i+1}],\end{aligned}$$

and so $\Pr[R \neq r_1 \mid r_2 = r_1, H_{i+1}] \geq (2\alpha)^{i-1}(2\alpha)/2 = (2\alpha)^i/2$. Observe that the last probability is for the process where, starting at H_{i+1}, the same bit r_1 is sent i times. This is exactly the probability in the statement of our claim (for the repetition protocol starting at H_{i+1}).

This completes the proof of the theorem.

3.4 Assembling the pieces for one-way protocols

Here we show how to combine the two building blocks we just used: the Parity protocol and the repetition protocol. Let $\alpha > 4(1 + 2\delta)\beta$. We re-state the main theorem of this section.

Theorem 15. *For any non-negligible ϵ and any $1/2 > \alpha > 4\beta > 0$, there is a one-way protocol for secrecy and reliability amplification from (α, β) to $(1/2 - \epsilon, 2^{-k})$.*

Proof. First, we can use the following protocol to make α and β suitably small without changing their ratios:

> With probability p, Alice uses the channel to send a random bit b, otherwise she sends b in the clear. This protocol is $\alpha' = p\alpha$ secret and $\beta' = p\beta$ reliable.

Since $1 - \alpha' \approx e^{-\alpha'}$ for small α', we can pick p small enough so that $(1 - \alpha') < e^{-\alpha(1-\delta)}$. Then we use the Parity protocol of Theorem 11 with $n = \log k$ to define a channel that has secrecy at least

$$\begin{aligned}(1/2) \cdot \left(1 - (1 - \alpha')^{n/2}\right) &\geq (1/2) \cdot \left(1 - k^{-(\alpha/2)(1-\delta)}\right) \\ &\geq (1/2) \cdot \left(1 - k^{-2\beta(1+\delta)}\right),\end{aligned}$$

and reliability at least $(1/2) \cdot \left(1 - e^{-\beta n}\right) = (1/2) \cdot \left(1 - k^{-\beta}\right)$.

We use the repetition protocol on this channel for $N = k^{2\beta(1+\delta/2)}$ repetitions. By Theorem 14, the resulting channel has secrecy at least $(1/2) \cdot (1 - k^{-\beta\delta})$ and, by Theorem 14, reliability at most $e^{-k^{-2\beta}N/8} = e^{-(1/8)k^{\beta\delta}}$, which tends to 0 exponentially fast with k. We can use the Parity protocol with $n = k$ on this protocol, to get one that is $(1/2 - \epsilon)$-secret for arbitrary non-negligible ϵ, and still has exponentially small reliability. If we want, we can then use repetition on this protocol for any polynomial number of times to keep the advantage of an adversary negligible, while making the reliability as good as desired.

Remark 16. The above shows a one-way protocol when $\alpha > 4\beta$. The factor of 4 can be thought of as two factors of two. The first one is due to the quadratic dependence of list size on the advantage when list decoding the Hadamard code (cf. the proof of Theorem 11 above). The second factor of 2 is because repeating a message through a symmetric channel takes quadratic time in the advantage, whereas for an erasure channel, the advantage grows linearly (cf. the proof of Theorem 17 below).

4 Impossibility results for one-way protocols

Here, we show that a constant factor difference of two between α and β is *necessary*. To get our negative result, we will look at a particular channel; of course, it follows that if no protocol exists for this channel, then no protocol exists for an unknown channel. Our particular channel is stateless, and is

- Symmetric β-Noise Channel for Bob: each bit sent over the channel is flipped with probability β, and is unchanged with probability $1 - \beta$,
- 2α-Erasure Channel for Eve: each bit sent over the channel is erased with probability 2α (with Eve getting a special symbol '?'), and is unchanged with probability $1 - 2\alpha$.

In addition, we allow Eve to have unlimited computational power.

We prove the following result, using the techniques of Holenstein and Renner [19].

Theorem 17. *If $\alpha \leq 2\beta - 2\beta^2$, then no one-way protocol for the above channel has reliability* .01 *and secrecy* .49*.*

Proof. We use the techniques of Holenstein and Renner [19] who showed that the same relationship between secrecy and reliability parameters is necessary for any information-theoretic one-way protocol for secret key agreement. Let a random variable B denote the bit to be sent. Let $X_1, \ldots, X_n$ be the distribution on bits Alice sends through the channel, and let V be the distribution on messages she sends in the clear. Let $Y_1, \ldots, Y_n$ be the bits Bob receives, and $Z_1, \ldots, Z_n$ be the information Eve receives.

Let H be the entropy function. Let B' bet the Boolean random variable that is 1 iff Bob correctly guesses the bit B, given V and $Y_1, \ldots, Y_n$. Since, given $V, Y_1, \ldots, Y_n$, Bob guesses B correctly with probability at least .99, we get $H(B' \mid V, Y_1, \ldots, Y_n) \leq H(.99)$. On the other hand, note that V and $Y_1, \ldots, Y_n$ determine Bob's guess at B, and so if we know B, then we also know B', and vice versa. It follows that $H(B \mid V, Y_1, \ldots, Y_n) = H(B' \mid V, Y_1, \ldots, Y_n) \leq H(.99) \approx 0$. By a similar reasoning for Eve, we get that $H(B \mid V, Z_1, \ldots, Z_n) \geq H(.49) \approx 1$.

Consider $H(B \mid V, Y_1, ..Y_i, Z_{i+1}...Z_n)$. When $i = n$, this is close to 0, and when $i = 0$, close to 1. So there must exist an index i, $0 \leq i \leq n$, such that

$$H(B \mid V, Y_1, \ldots, Y_i, Z_{i+1}, \ldots, Z_n) < H(B \mid V, Y_1, \ldots, Y_{i-1}, Z_i, \ldots, Z_n).$$

Then by an averaging argument, there must exist values for V, $Y_1, \ldots, Y_{i-1}$ and $Z_{i+1}, \ldots, Z_n$, so that in the conditional distribution, we have

$$H(B \mid Y_i) < H(B \mid Z_i). \tag{5}$$

Note that, because the protocol is one-way, conditioning on these values does not change the conditional distributions of Y_i or Z_i as functions of X_i (the bit sent)[8]. It will possibly change both the distributions of B and X_i to arbitrary distributions.

By Eq. (5), and using the entropy chain rule twice, we get

$$\begin{aligned} 0 &> H(B \mid Y_i) - H(B \mid Z_i) \\ &= H(B, Y_i) - H(Y_i) - H(B, Z_i) + H(Z_i) \\ &= H(B) + H(Y_i \mid B) - H(Y_i) - H(B) - H(Z_i|B) + H(Z_i) \\ &= H(Y_i \mid B) - H(Y_i) - H(Z_i \mid B) + H(Z_i). \end{aligned}$$

[8] In contrast, consider a 2-way protocol where Bob, after receiving his n bits over the channel, sends Alice a message in the clear stating whether all his received bits are the same. Then fixing the value of Bob's message to Alice *will change* the distribution of Y_i as a function of X_i. So the argument in the present theorem does not apply to this 2-way protocol. (In fact, we use such a 2-way protocol in Section 5 in order to overcome the "factor-2 barrier" for one-way protocols given by the present theorem.)

Next we analyze each of the four summands in the last equation above.

Let q be the conditional probability that $B=1$, and let p_1 be the conditional probability that $X_i=1$ if $B=1$, and p_0 be the conditional probability that $X_i=1$ if $B=0$. Then the overall probability that $X_i=1$ is

$$p := qp_1 + (1-q)p_0.$$

Note that Y_i is equal to X_i with probability $1-\beta$, and to $\neg X_i$ otherwise. It follows that

$$H(Y_i) = H(p(1-2\beta)+\beta). \tag{6}$$

Next, given $B=1$, Y_i is distributed as first flipping a coin with probability p_1 to determine X_1, then a coin with probability β, and finally taking the parity. So we have

$$H(Y_i \mid B=1) = H(p_1(1-2\beta)+\beta),$$

and similarly,

$$H(Y_i \mid B=0) = H(p_0(1-2\beta)+\beta).$$

Combining the two conditional entropies, we conclude

$$H(Y_i \mid B) = q\cdot H(p_1(1-2\beta)+\beta) + (1-q)\cdot H(p_0(1-2\beta)+\beta), \tag{7}$$

Finally, Z_i reveals whether the bit is erased, a random event with probability 2α no matter what, and then, with probability $1-2\alpha$, it reveals the value of X_i. Thus, $H(Z_i) = H(2\alpha) + (1-2\alpha)\cdot H(X_i)$, and the same for any conditional distribution. So we get

$$H(Z_i) = H(2\alpha) + (1-2\alpha)\cdot H(p), \tag{8}$$

and

$$H(Z_i \mid B) = H(2\alpha) + (1-2\alpha)\cdot (q\cdot H(p_1) + (1-q)\cdot H(p_0)). \tag{9}$$

Combining Eqs. (6)–(9), we get

$$\begin{aligned}
0 &> (H(Y_i \mid B) - H(Y_i)) - (H(Z_i \mid B) - H(Z_i))\\
&= q\cdot H(p_1(1-2\beta)+\beta) + (1-q)\cdot H(p_0(1-2\beta)+\beta) - H(p(1-2\beta)+\beta) -\\
&\quad (H(2\alpha) + (1-2\alpha)\cdot(q\cdot H(p_1) + (1-q)\cdot H(p_0)) - H(2\alpha) - (1-2\alpha)\cdot H(p)).
\end{aligned}$$

Rearranging the terms in the last expression, we can write it as

$$\begin{aligned}
&q\cdot (H(p_1(1-2\beta)+\beta) - (1-2\alpha)\cdot H(p_1))\\
&+ (1-q)\cdot (H(p_0(1-2\beta)+\beta) - (1-2\alpha)\cdot H(p_0))\\
&- (H(p(1-2\beta)+\beta) - (1-2\alpha)\cdot H(p))\\
&= q\cdot F(p_1) + (1-q)\cdot F(p_0) - F(p),
\end{aligned}$$

for the function $F(x) := H(x\cdot(1-2\beta)+\beta) - (1-2\alpha)\cdot H(x)$. Thus, we have

$$q\cdot F(p_1) + (1-q)\cdot F(p_0) - F(p) < 0,$$

which is equivalent (recalling that $p = qp_1 + (1-q)p_0$) to

$$F(qp_1 + (1-q)p_0) > q \cdot F(p_1) + (1-q) \cdot F(p_0). \tag{10}$$

Observe that Eq. (10) states that the function F at a convex combination of two points is greater than the convex combination of its values at those two points. This condition is violated if F is a convex function on the interval $[0,1]$. So, to complete our proof by contradiction, it suffices to show

Claim. The function $F(x)$ defined above is convex on $[0,1]$.

Proof (of Claim). We use the convexity criterion for twice differentiable functions: such a function is convex over an interval iff its second derivative is nonnegative on that interval. We can change the binary logs to natural logs, since that just multiplies F by a positive constant factor. For the ln-based entropy function $h(x) = -x\ln x - (1-x)\ln(1-x)$, its first derivative is $h'(x) = -\ln x + \ln(1-x)$, and its second derivative is $h''(x) = -1/x - 1/(1-x)$.

Similarly, for the linear function $L(x) := x(1-2\beta)+\beta$, one can easily verify that

$$(h(L(x)))' = (1-2\beta)\cdot(\ln(1-L(x)) - \ln(L(x))),$$

and

$$(h(L(x)))'' = (1-2\beta)^2 \cdot \left(-\frac{1}{1-L(x)} - \frac{1}{L(x)}\right).$$

Using these expressions for the second derivatives of $h(x)$ and $h(L(x))$, we get

$$\begin{aligned} F''(x) &= (H(L(x)))'' - (1-2\alpha)\cdot H''(x) \\ &= (1-2\beta)^2 \cdot \left(-\frac{1}{1-L(x)} - \frac{1}{L(x)}\right) + (1-2\alpha)\cdot\left(\frac{1}{x} + \frac{1}{1-x}\right) \\ &= -(1-2\beta)^2 \cdot \frac{1}{L(x)\cdot(1-L(x))} + (1-2\alpha)\cdot\frac{1}{x(1-x)}. \end{aligned}$$

We want to show that $F''(x) \geq 0$ for all $x \in [0,1]$, i.e., that

$$\frac{1-2\alpha}{x(1-x)} \geq \frac{(1-2\beta)^2}{L(x)\cdot(1-L(x))}.$$

Note that $L(x) = x(1-2\beta) + (1/2)(2\beta)$, and so $L(x)$ is always between x and $1/2$ (no matter which side of $1/2$ the point x is). Since the function $x(1-x)$ is symmetric around $1/2$, and achieves its maximum at the point $1/2$, we conclude that $L(x)(1-L(x)) \geq x(1-x)$. Thus it suffices to show

$$\frac{1-2\alpha}{x(1-x)} \geq \frac{(1-2\beta)^2}{x(1-x)},$$

equivalent to $1-2\alpha \geq (1-2\beta)^2$. The latter is equivalent to $\alpha \leq 2\beta - 2\beta^2$, which is our assumption on the α and β.

This completes the proof of the theorem.

5 Breaking the factor of two barrier with two-way protocols

By the lower bound of Theorem 17, we know that it is impossible to amplify secrecy and reliability of a given α-secret and β-reliable channel when $\alpha < 2\beta$, if we use *one-way* communication only. Here we show that a *two-way* communication protocol exists that works even for $\alpha < 2\beta$, as long as $\alpha > (3/2)\beta$.

Our main result of the section is the following.

Theorem 18. *For any non-negligible ϵ and for any $1/2 > \alpha > 1.5 \cdot \beta > 0$, there is a two-way protocol for secrecy and reliability amplification from (α, β) to $(1/2 - \epsilon, 2^{-k})$.*

We will need a simple variant on the repetition protocol where Bob communicates one bit in the clear. Like the repetition protocol, this variant will reduce both secrecy and reliability exponentially. But, if $\alpha > 1.5\beta$, the exponent that secrecy decreases by will be larger than that for Bob's failure chance. So the ratio between them will improve with the number of repetitions. We can then pick the number of repetitions to be such that the ratio is greater than 4, and use this protocol as the channel in the one-way protocol from Theorem 15.

The variant protocol is **Repetition with Feedback**:

1. Alice uses the channel to send b to Bob n times.
2. If Bob receives the same bit b' each time, he sends the message "Consistent" to Alice in the clear and uses b' as his output. Otherwise he sends the message "Inconsistent" to Alice in the clear.
3. If Bob sends "Inconsistent", Alice sends b in the clear, and Bob uses that as his output.

We show the following.

Theorem 19. *Let α, β, n be any parameters such that n is poly-bounded in the security parameter, and $(2(\alpha - \beta))^n$ is non-negligible. The n-bit Repetition with Feedback protocol applied to an α-secret and β-reliable transparent channel yields a new α'-secret and β'-reliable channel, for $\alpha' \geq (2(\alpha - \beta))^n/2$ and $\beta' \leq \beta^n$.*

Proof. Reliability: First we argue reliability of the new channel. We need to show that for any attack on the Repetition with Feedback Protocol over a β-reliable channel, the probability that Bob fails to output b is at most β^n. Indeed, Bob gets b unless he receives the same bit b' each of n times, and $b' \neq b$. Thus, the protocol only fails if the channel fails n times in a row, which happens with probability at most β^n.

Secrecy: Next we argue secrecy of the new channel. We need to show that no attack on the n-bit Repetition with Feedback protocol using an α-secret and β-reliable transparent channel can predict a random bit b sent by the protocol with better than $1 - (2(\alpha - \beta))^n/2$ probability of success. As before, fixing functions A and f that describe Eve's attack, the process can be described as:

1. Alice picks a random bit r.
2. The protocol starts in some state s_{n+1}. Let the initial history H_{n+1} be the list containing only s_{n+1}.
3. For each i from n to 1:
   (a) Eve picks channel attack $a_i = A(H_{i+1})$.
   (b) The new state and bit Bob receives is $(s_i, b'_i) = \delta_k(s_{i+1}, a_i, r)$.
   (c) Append s_i to H_{i+1} to get an updated history H_i.
4. If all b'_i are equal (according to Bob's message in the clear), Eve guesses $R = f(H_1, \text{“Consistent”})$. Otherwise she learns b when it is sent in the clear.

The intuition is that, even if we revealed the secret to Eve whenever Bob fails to get the secret, the channel would remain $(\alpha - \beta)$-secret, because failure happens with probability at most β . We could then apply the analysis of the repetition protocol to this altered channel.

Define random variable $R = f(H_1, \text{“Consistent”})$, even if the bits received are possibly inconsistent. Consider starting from partial history H_{i+1}, picking a new random bit r_1 and simulating the protocol from then on sending r_1 for the i remaining bits to be sent, and verifying that $b'_i = r_1$ each time. The theorem will follow form the next claim for $i = n$ (which shows that with probability at least $(2(\alpha - \beta))^n/2$, Bob gets b all n times, sends “Consistent”, and Eve outputs $R \neq b$).

Claim. For each $1 \leq i \leq n$, $\Pr[R \neq r_1, \wedge_{1 \leq j \leq i}(b'_j = r_1) \mid H_{i+1}] \geq (2(\alpha - \beta))^i/2$.

Proof (of Claim). Our proof is by induction on i. For $i = 1$, this follows from α-secrecy and β-reliability: the probability that $R \neq r_1$ is at least α, and the probability that $b'_1 \neq r_1$ is at most β, so the probability that $R \neq r_1 = b'_1$ is at least $\alpha - \beta$. Consider the following strategy for Eve to predict a single bit r_1 sent on the channel at state s_{i+1}:

> Eve uses a_i as her attack when Alice sends r_1, and the channel arrives in state s_i. Then she picks a new random bit r_2 and simulates the repetition protocol with feedback starting from H_i, with Alice sending r_2 each time (including simulating the bit Bob receives) . If the simulation returns an $R \neq r_2$ and Bob receives r_2 each time, Eve guesses R. Otherwise, Eve repeats the simulation for a fresh random bit r_2. (Note that the expected number of repetitions is at most $2(2(\alpha-\beta))^{-i}$, by the induction hypothesis, which is feasible by assumption).

Denote by Success_i the event that Bob receives r_2 each of the last i times. Fix any history H_i, together with r_1. The probability that, for the R returned by Eve in the above strategy, $R \neq r_1$ is

$$\Pr[R \neq r_1 \mid R \neq r_2, H_i, \text{Success}_{i-1}] = \frac{\Pr[R = \neg r_1 = \neg r_2, \text{Success}_{i-1} \mid H_i]}{\Pr[R \neq r_2, \text{Success}_{i-1} \mid H_i]}.$$

By induction, for each such H_i, the denominator of this expression is at least $(2(\alpha-\beta))^{i-1}/2$. So for each H_i where $b'_i = r_1$,

$$\frac{(2(\alpha-\beta))^{i-1}}{2} \cdot \Pr[R \neq r_1 \mid R \neq r_2, H_i, \mathrm{Success}_{i-1}] \\ \leq \Pr[r_2 = r_1, R \neq r_1, \mathrm{Success}_{i-1} | H_i].$$

Note that H_i already determines (although Eve doesn't know which way) whether Bob received r_1, i.e., whether $b'_i = r_1$. For those histories where this did happen, the conditional probability that $R \neq r_1$ and Bob receives r_1 is the same as just the first clause, and for the others, it is 0. So either way we get

$$\frac{1}{2} \cdot (2(\alpha-\beta))^{i-1} \cdot \Pr[R \neq r_1, b'_i = r_1 \mid R \neq r_2, H_i, \mathrm{Success}_{i-1}] \\ \leq \Pr[r_2 = r_1, R \neq r_1, b'_i = r_1, \mathrm{Success}_{i-1} \mid H_i].$$

Then we can average both sides over all H_i, to get

$$\frac{1}{2} \cdot (2(\alpha-\beta))^{i-1} \cdot \Pr[R \neq r_1, b'_i = r_1 \mid R \neq r_2, H_{i+1}, \mathrm{Success}_{i-1}] \\ \leq \Pr[r_2 = r_1, R \neq r_1, b'_i = r_1, \mathrm{Success}_{i-1} \mid H_{i+1}].$$

By α-secrecy and β-reliability, the probability on the left-hand side of the inequality above is at least $\alpha - \beta$. The probability on the right-hand side is $1/2$ (the probability that $r_2 = r_1$), times the probability that $R \neq r_1$ and $\mathrm{Success}_i$ when r_1 is sent i times starting at H_{i+1}. The latter probability is exactly the probability in the statement of the claim. Thus, we get

$$\Pr[R \neq r_1, \wedge_{1 \leq j \leq i}(b'_j = r_1) \mid H_{i+1}] \geq \frac{1}{2} \cdot (2(\alpha-\beta))(2(\alpha-\beta))^{i-1}.$$

This completes the proof of the theorem.

As a corollary, we get the desired proof of the main result of this section.

Proof (of Theorem 18). Given $\alpha > 1.5\beta$, we first use the Repetition with Feedback protocol for an appropriate number of times to get a new protocol channel with α'-secrecy and β'-reliability for $\alpha' > 4\beta'$. Then we use the protocol of Theorem 15 on this protocol channel.

Tightness of the analysis of the Repetition with Feedback protocol. In our analysis of the Repetition with Feedback protocol, the ratio of secrecy to reliability improves with n when $2(\alpha-\beta) > \beta$, i.e., when $\alpha > 1.5\beta$. In other cases, it makes things worse, rather than better. We now show this analysis is actually tight.

Consider the channel where, with probability 2β, Eve and Bob both receive a random bit b'. In addition, Eve receives A, denoting that this is the case in question. With probability $2(\alpha-\beta)$, Bob receives the correct bit b, and Eve receives just the message B, saying that this is the case. With the remaining probability $1 - 2\alpha$, Bob receives the correct bit b, and Eve also receives b and the message C.

In the repetition with feedback, if the messages Bob receives are consistent, and C has occurred, Eve knows with certainty one bit Bob received and hence that bit must have been received all n times. If the messages Bob receives are consistent, and A occurred, then Eve and Bob get the same random bit b' all n times.

If Bob's messages are inconsistent, the secret is sent in the clear and Eve gets it. Eve fails to get the secret when either *(i)* case B happens all n times, and thereafter Eve does not guess the random bit sent by Alice, or *(ii)* case A happens all n times, and the random bit b' is different from Alice's bit. Thus the overall failure probability for Eve is at most $(2(\alpha - \beta))^n/2 + \beta^n$.

6 Conclusions and open problems

In this paper, we considered just the simplest issue in secure communication, the transmission of secret information from one party to another. Even here, there are unexpected complications arising from the joint consideration of secrecy and reliability. We gave non-trivial constructions of secure protocols that under some circumstances are guaranteed to amplify both secrecy and reliability to within negligible amounts of the ideal.

However, our results raise more questions than they answer. We hope that these will be addressed in future work, and that future work will consider similar models for more complex issues in secure communications. We suggest the following tasks to consider for the case of trusted parties: authentication, covert channels (steganography), and traffic analysis. For the case of untrusted parties, it will be interesting to use an appropriate channel model to argue about: coin flipping, oblivious transfer, multi-party computation, and broadcast.

It would also be very interesting to study channel models with weaker restrictions on transparency. For example, can one generalize our channel model to include the quantum-computational setting?

Acknowledgments We thank Yevgeny Dodis, Noah Stevens-Davidowitz, Giovanni di Crescenzo, Daniele Micciancio, Thomas Holenstein and Steven Rudich for helpful comments and discussions. Russell Impagliazzo's work was partially supported by the Simons Foundation and NSF grant CCF-121351; this work was done [in part] while Russell Impagliazzo was visiting the Simons Institute for the Theory of Computing, supported by the Simons Foundation and by the DIMACS/Simons Collaboration in Cryptography through NSF grant #CNS-1523467. Valentine Kabanets was partially supported by the NSERC Discovery grant. Bruce Kapron's work was supported in part by the NSERC Discovery Grant "Foundational Studies in Privacy and Security". Stefano Tessaro was partially supported by NSF grants CNS-1423566, CNS-1553758, CNS-1528178, IIS-1528041 and the Glen and Susanne Culler Chair.

Bibliography

[1] Mihir Bellare, Russell Impagliazzo, and Moni Naor. Does parallel repetition lower the error in computationally sound protocols? In *FOCS '97: Proceedings of the 38th IEEE Annual Symposium on Foundations of Computer Science*, pages 374–383, 1997. 3

[2] Mihir Bellare, Stefano Tessaro, and Alexander Vardy. Semantic security for the wiretap channel. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology – CRYPTO 2012*, volume 7417 of *Lecture Notes in Computer Science*, pages 294–311. Springer Berlin Heidelberg, 2012. 4

[3] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14-17 October 2001, Las Vegas, Nevada, USA*, pages 136–145. IEEE Computer Society, 2001. 5

[4] Kai-Min Chung and Feng-Hao Liu. Parallel repetition theorems for interactive arguments. In *Theory of Cryptography — TCC 2010*, volume 5978 of *Lecture Notes in Computer Science*, pages 19–36, 2010. 3

[5] Kai-Min Chung and Rafael Pass. Tight parallel repetition theorems for public-coin arguments using kl-divergence. In *Theory of Cryptography — TCC 2010*, volume 9015 of *Lecture Notes in Computer Science*, pages 229–246, 2015. 3

[6] Claude Crépeau. Efficient cryptographic protocols based on noisy channels. In Walter Fumy, editor, *Advances in Cryptology EUROCRYPT 97*, volume 1233 of *Lecture Notes in Computer Science*, pages 306–317. Springer Berlin Heidelberg, 1997. 4

[7] Claude Crépeau and Joe Kilian. Achieving oblivious transfer using weakened security assumptions. In *Foundations of Computer Science, 1988., 29th Annual Symposium on*, pages 42–52, Oct 1988. 4

[8] Claude Crépeau, Kirill Morozov, and Stefan Wolf. Efficient unconditional oblivious transfer from almost any noisy channel. In *Security in Communication Networks, 4th International Conference, SCN 2004, Amalfi, Italy, September 8-10, 2004, Revised Selected Papers*, pages 47–59, 2004. 4

[9] I. Csiszar and J. Körner. Broadcast channels with confidential messages. *Information Theory, IEEE Transactions on*, 24(3):339–348, May 1978. 1, 4, 5

[10] Yevgeniy Dodis. Shannon impossibility, revisited. In *Proceedings of the 6th International Conference on Information Theoretic Security*, ICITS'12, pages 100–110, Berlin, Heidelberg, 2012. Springer-Verlag. 3

[11] Cynthia Dwork, Moni Naor, and Omer Reingold. Immunizing encryption schemes from decryption errors. In *Advances in Cryptology — EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 342–360, 2004. 3

[12] Sanjam Garg, Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Cryptography with one-way communication. In Rosario Gennaro

and Matthew Robshaw, editors, *Advances in Cryptology – CRYPTO 2015*, volume 9216 of *Lecture Notes in Computer Science*, pages 191–208. Springer Berlin Heidelberg, 2015. 4

[13] Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing*, pages 25–32, 1989. 5, 14

[14] Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270 – 299, 1984. 4

[15] Iftach Haitner. A parallel repetition theorem for any interactive argument. In *FOCS '09: Proceedings of the 50th IEEE Annual Symposium on Foundations of Computer Science*, pages 241–250, 2009. 3

[16] Shai Halevi and Tal Rabin. Degradation and amplification of computational hardness. In *Theory of Cryptography — TCC 2008*, volume 4948 of *Lecture Notes in Computer Science*, pages 626–643, 2008. 3, 6, 15

[17] Johan Håstad, Rafael Pass, Douglas Wikström, and Krzysztof Pietrzak. An efficient parallel repetition theorem. In *Theory of Cryptography — TCC 2010*, volume 5978 of *Lecture Notes in Computer Science*, pages 1–18, 2010. 3

[18] Thomas Holenstein. Key agreement from weak bit agreement. In *STOC '05: Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 664–673, 2005. 3, 5

[19] Thomas Holenstein and Renato Renner. One-way secret-key agreement and applications to circuit polarization and immunization of public-key encryption. In Victor Shoup, editor, *Advances in Cryptology - CRYPTO 2005: 25th Annual International Cryptology Conference, Santa Barbara, California, USA, August 14-18, 2005, Proceedings*, volume 3621 of *Lecture Notes in Computer Science*, pages 478–493. Springer, 2005. 3, 4, 5, 18

[20] Thomas Holenstein and Grant Schoenebeck. General hardness amplification of predicates and puzzles - (extended abstract). In *Theory of Cryptography — TCC 2010*, volume 6597 of *Lecture Notes in Computer Science*, pages 19–36, 2011. 3

[21] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, Manoj Prabhakaran, Amit Sahai, and Jürg Wullschleger. Constant-rate oblivious transfer from noisy channels. In Phillip Rogaway, editor, *Advances in Cryptology CRYPTO 2011*, volume 6841 of *Lecture Notes in Computer Science*, pages 667–684. Springer Berlin Heidelberg, 2011. 4

[22] M. Iwamoto and K. Ohta. Security notions for information theoretically secure encryptions. In *Information Theory Proceedings (ISIT), 2011 IEEE International Symposium on*, pages 1777–1781, July 2011. 4

[23] Mitsugu Iwamoto, Kazuo Ohta, and Junji Shikata. Security formalizations and their relationships for encryption and key agreement in information-theoretic cryptography. *CoRR*, abs/1410.1120, 2014. 4

[24] Leonid A. Levin. One-way functions and pseudorandom generators. *Combinatorica*, 7(4):357–363, 1987. 5, 11

[25] Yingbin Liang, H. Vincent Poor, and Shlomo Shamai (Shitz). Information theoretic security. *Foundations and Trends in Communications and Information Theory*, 5(45):355–580, 2008. 4

[26] Huijia Lin and Stefano Tessaro. Amplification of chosen-ciphertext security. In *Advances in Cryptology — EUROCRYPT 2009*, volume 7881 of *Lecture Notes in Computer Science*, pages 503–519, 2013. 3

[27] Ueli Maurer. Constructive cryptography - A new paradigm for security definitions and proofs. In Sebastian Mödersheim and Catuscia Palamidessi, editors, *Theory of Security and Applications - Joint Workshop, TOSCA 2011, Saarbrücken, Germany, March 31 - April 1, 2011, Revised Selected Papers*, volume 6993 of *Lecture Notes in Computer Science*, pages 33–56. Springer, 2011. 5

[28] Ueli Maurer and Renato Renner. Abstract cryptography. In *ICS*, pages 1–21. Tsinghua University Press, 2011. 5

[29] Ueli M. Maurer. Perfect cryptographic security from partially independent channels. In *Proceedings of the Twenty-third Annual ACM Symposium on Theory of Computing*, STOC '91, pages 561–571, New York, NY, USA, 1991. ACM. 4

[30] Ueli M. Maurer. Secret key agreement by public discussion from common information. *Information Theory, IEEE Transactions on*, 39(3):733–742, May 1993. 4, 5

[31] Ueli M. Maurer. Information-theoretic cryptography. In Michael Wiener, editor, *Advances in Cryptology CRYPTO 99*, volume 1666 of *Lecture Notes in Computer Science*, pages 47–65. Springer Berlin Heidelberg, 1999. 4

[32] Rafael Pass and Muthuramakrishnan Venkitasubramaniam. An efficient parallel repetition theorem for Arthur-Merlin games. In *STOC '07: Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pages 420–429, 2007. 3

[33] Krzysztof Pietrzak and Douglas Wikström. Parallel repetition of computationally sound protocols revisited. In *Theory of Cryptography — TCC 2007*, volume 4392 of *Lecture Notes in Computer Science*, pages 86–102, 2007. 3

[34] Amit Sahai and Salil P. Vadhan. A complete promise problem for statistical zero-knowledge. In *38th Annual Symposium on Foundations of Computer Science, FOCS '97, Miami Beach, Florida, USA, October 19-22, 1997*, pages 448–457. IEEE Computer Society, 1997. 4

[35] Claude E. Shannon. Communication theory of secrecy systems. *Bell System Technical Journal*, 28:656–715, 1949. 1, 3

[36] J. Shikata. Formalization of information-theoretic security for key agreement, revisited. In *Information Theory Proceedings (ISIT), 2013 IEEE International Symposium on*, pages 2720–2724, July 2013. 4

[37] Jürg Wullschleger. Oblivious-transfer amplification. In *Advances in Cryptology — EUROCRYPT 2007*, volume 4515 of *Lecture Notes in Computer Science*, pages 555–572, 2007. 3

[38] Jürg Wullschleger. Oblivious transfer from weak noisy channels. In *Proceedings of the 6th Theory of Cryptography Conference on Theory of Cryptography*, TCC '09, pages 332–349, Berlin, Heidelberg, 2009. Springer-Verlag. 4

[39] Aaron D. Wyner. The wire-tap channel. *Bell System Technical Journal*, 54:1355–1387, 1975. 1, 4