Towards Multiparty Computation Withstanding Coercion of All Parties

Ran Canetti*   Oxana Poburinnaya†

November 15, 2020

Abstract

Incoercible multi-party computation (Canetti-Gennaro '96) allows parties to engage in secure computation with the additional guarantee that the public transcript of the computation cannot be used by a coercive outsider to verify representations made by the parties regarding their inputs, outputs, and local random choices. That is, it is guaranteed that the only deductions regarding the truthfulness of such representations, made by an outsider who has witnessed the communication among the parties, are the ones that can be drawn just from the represented inputs and outputs alone. To date, all incoercible secure computation protocols withstand coercion of only a fraction of the parties, or else assume that all parties use an execution environment that makes some crucial parts of their local states physically inaccessible even to themselves.

We consider, for the first time, the setting where *all parties* are coerced, and the coercer expects to see *the entire history of the computation.* We allow both protocol participants and external attackers to access a common reference string which is generated once and for all by an uncorruptable trusted party. In this setting we construct:

- A general multi-party function evaluation protocol, for any number of parties, that withstands coercion of all parties, as long as all parties use the prescribed "faking algorithm" upon coercion. This holds even if the inputs and outputs represented by coerced parties are globally inconsistent with the evaluated function.
- A general two-party function evaluation protocol that withstands even the case where some of the coerced parties do follow the prescribed faking algorithm. (For instance, these parties might disclose their true local states.) This protocol is limited to functions where the input of at least one of the parties is taken from a small (poly-size) domain. It uses fully deniable encryption with public deniability for one of the parties; when instantiated using the fully deniable encryption of Canetti, Park, and Poburinnaya [Crypto'20], it takes 3 rounds of communication.

Both protocols operate in the common reference string model, and use fully bideniable encryption (Canetti Park and Poburinnaya, Crypto'20) and sub-sexponential indistiguishability obfuscation. Finally, we show that protocols with certain communication patterns cannot be incoercible, even in a weaker setting where only some parties are coerced.

Contents

1 Introduction

Consider a tight-knit society whose members regularly meet behind closed doors and run their society's business with complete privacy. An external entity might be able to deduce information on the nature of the interactions that take place in the society's meetings from the external behavior of the society members, but no direct information on what really takes place at the meetings can be obtained. As long as the meetings are not directly monitored by the external entity, this continues to be the case even if the external entity has coercive power over the society's members and demand that they fully disclose the contents of the meetings: All that the coercive entity can obtain is the word of the members, which may or may not be true.

Can we reproduce this situation online, where the society members communicate over public channels that are accessible to the external entity? That is, can the society members engage in a multiparty computation that allows them to limit the power of the external coercive entity to the power that it had when they met behind closed doors? Furthermore, can they do so even in the case where *all* members are coerced, and the coercive entity now expects to have the complete history of the interaction and the local states of all parties, including all the local randomness used? Indeed, doing so essentially results in rewriting the entire history of a system, in a way that's undetectable to anyone that did not directly witness the events at the time and location where they took place.

This is a special case of the *incoercible multiparty computation* problem, first studied in [CG96]. In a nutshell, a multiparty protocol is *incoercible* if it enables the participants to preserve the privacy of their inputs and outputs even against a coercive adversary who demands to see the entire internal state of the coerced parties. Towards this end, each party is equipped with a "faking algorithm" that enables it to run the protocol as prescribed on the given input x, obtain an output y, and then, given arbitrary values x', y', generate a "fake random string" r', such that:

Local consistency: The public communication transcript of the party is consistent with input x', output y' and local random string r' for the party.

Global consistency: As long as all coerced parties follow their prescribed faking algorithms, and the inputs and outputs claimed by all the coerced parties are consistent with the evaluated function, the entire information reported by the parties should look like an honest execution of the protocol with these inputs and outputs; this should hold regardless of whether the inputs and outputs are true, or fake, or partially true and partially fake.

Furthermore, even if the inputs and outputs claimed by the individual parties, put together, are not consistent with the evaluated function, the coercer should not be able to deduce any information which it cannot deduce given the inputs and outputs alone. In other words, even when the coercer notices that the inputs and outputs reported by the parties are inconsistent, the internal randomness r' provided by the coerced parties does not help the adversary to determine the source of the inconsistency (ie, which parties are faking their inputs or outputs).

The above requirements relate to parties that follow their prescribed faking algorithm. (Essentially, this model captures the case where the parties share a common interest to thwart an external coercer, as in the tight-knit society example described above.) In addition, one may consider situations where parties deviate from the prescribed faking algorithm, potentially in order to "collaborate" with the coercer, or to "frame" another party. Here would like provide the following adtional guarantee:

Unframeability: The global consistency guarantees continue to hold with respect to the parties that follow the protocol during execution and the faking algorithm at coercion. This is the case even if some of the players deviate arbitrarily from the prescribed faking algorithm.

Note that the above properties refer to the case where the coercer contacts the parties only after the protocol has ended. We will return to this point later on. (Following [AOZZ15], we refer to this case as *semi-honest incoercibility*.)

These requirements might indeed appear unobtainable at first. Still, assuming sender-deniable encryption [CDNO96, SW14], the works of [CG96] construct an incoercible general multi-party function evaluation protocol, for the case where only a minority of the parties are coerced. The protocols of [CGP15], [DKR14], originally devised as adaptively secure protocols withstanding corruption of all parties, can withstand coercion of a single party. The works of [MN06, AOZZ15] consider the case where all parties are coerced - in fact they consider an even more adversarial setting of *active coercions*, where the coercer may force parties to deviate from the protocol, to make it harder for them to deceive the coercer. To provide incoercibility against such a stronger adversary, they consider a model where the parties have access to secure hardware whose internals are not available *even to the parties themselves.*

Whether incoercibility is at all possible in a setting where all parties are coerced, the communication is public, and the parties have full access to the transcripts of their own internal computations, has remains open. Indeed, in this case the adversary obtains an entire transcript of a computation, which can be verified step by step. Still, it should be unable to tell a fake transcript from a real one.

Our results. We consider the case where the parties have full access to the computing devices they use, all communication is public, and all parties are eventually coerced. We allow parties to have access to a common reference string, generated once and for all by a trusted party out of adversary's reach. We concentrate on the case where all corruptions and coercions take place only at the end of the protocol. (Indeed, when parties are corrupted during protocol execution, and the adversary has access to all communication and all internal states of the coerced or corupted parties, incoercibility becomes essentially unobtainable.) We consider two main settings, or levels, of incoercibility:

Full incoercibility captures all the above three requirements: local consistency, global consistency, and unframeability. More specifically, it considers the case where parties may have competing interests and there is no coordination among them. In particular, while some parties present fake randomness, inputs, outputs, others may present their true randomness, inputs and outputs. Moreover, here the adversary make fully *corrupt* some of the parties, while coercing others.[1]

Cooperative incoercibility captures only local and global consistency. More specifically, it considers the case where the parties have a common interest and want to protect themselves against an external coercer. For this purpose, all parties either present their real randomness, inputs, and outputs, or else they all present randomness computed via their faking algorithms - even if they choose to report their true inputs and outputs[2].

---

[1]Recall that, when corrupting a party, the adversary is guaranteed to obtain the party's true internal state. In contrast, when coercing a party, the adversary may either obtain the true internal state or a fake one, depending on the party's decision. See more details on the modeling later on.

[2]As we explain later, one of our protocols only provides security in this model.

Still, each party runs the faking algorithm locally, without any coordination of secret randomness. The inputs and outputs claimed by the parties may or may not be coordinated ahead of time. In particular, they need not be globally consistent with the evaluated function.

In both settings, the requirement is that the external coercive entity be unable to distinguish an interaction with the protocol from an interaction with an idealized system where all the communication and all the internal states of the coerced and corrupted parties are generated by a simulator that obtains only the claimed inputs and outputs of the coerced parties and the true inputs and outputs of the corrupted parties.

We show:

- A cooperatively incoercible protocol for general secure multi-party function evaluation. Our protocol works in the common reference string (CRS) model and requires 4 rounds of communication[3].

- A fully incoercible protocol for secure two-party function evaluation, for functions with poly-size input domains for at least one of the parties. For this protocol, we build an incoercible oblivious transfer (OT) from any bideniable encryption with certain properties. When instantiated with the bideniable encryption of [CPP20], twe obtain a 3-message protocol in the CRS model.

- A negative result: For $n \geq 3$, no n-party protocols with a certain communication pattern can be secure even against coercion of 2 parties, except for trivial functions.

On the significance of full incoercibility. With cooperative incoercibility only, the adversarial parties may be able to provide an unequivocal proof that other parties are lying; thus, this type of incoercibility doesn't protect the participants against each other. In contrast, full incoercibility guarantees that the coercer will not be able to use the protocol transcript to verify claims of parties — any deduction made by the coercer will be exclusively based on "taking the word" of the coerced parties.[4]

We note that full incoercibility is similar in flavor to "off the record deniability" of bideniable encryption schemes [CPP20], in the sense that both notions consider the case where the coerced parties may try "frame" each others. We note however that, while there "off the record deniability" is incomparable to plain bideniability, here full incoercibility implies cooperative incoercibility.

1.1 Technical Overview

On the definition of incoercible computation. We use the definition of incoercible computation from [CGP15], which can be seen as the "UC equivalent" of the definition of [CG96], with one critical difference.

---

[3]Our notion of incoercibility concentrates on the requirements from the protocol for evaluating the given function. It is agnostic to whether the evaluated function allows coercion in and of itself. For instance, consider the "voting functionality" that outputs the tally, together with the names of the voters and their individual votes. This functionality is clearly coercible, even if implemented in a completely ideal way. In fact, even the coercion-aware variant of this functionality is coercible. This means that even realizing this functionality in a coercion-resilient manner will not help to thwart coercion attacks.

[4]The difference between these two notions can be best manifested by considering the protection they provide against a coercer that gives parties incentives to demonstrate that the other is lying. Full incoercibility guarantees that the parties are immune to such coercive situation, since they are provably unable to prove that otherds are lying. In contrast, since cooperative incoercibility does allow protocols where parties are able to prove that other players are lying, it does not protect the players from a "prisoner's dilemma" situation where players are incentivized to expose their real secrets, against their better interests.

(See section 1.2 for the discussion of the difference between the two.) Specifically, this definition models coercion as the following special form of corruption: When a party is notified that it is coerced, it first contacts the its caller to ask whether to disclose true or fake randomness, and if fake, what value (input and output) to report to the coercer. The response can be either "fake" together with an input and an output, or "tell the truth". In the former case, the coerced party runs the faking algorithm with the prescribed value; in the latter case it reveals its actual internal state.

In the ideal process, when the simulator asks to coerce a party, the ideal functionality obtains from the environment either the value v to be presented, or the "tell the truth" directive. If the response was a value v, then the functionality forwards v to the simulator. If the response was "tell the truth", then the ideal functionality provides the actual input and output values of the coerced party to the simulator. *Crucially, the simulator isn't told if this value is true or fake.* Intuitively, the fact that the simulator can simulate the protocol without learning whether the inputs were real or fake, means that in the real world the adversary doesn't learn this information either.

Observe that this definition captures full incoercibility. Indeed, it is up to the environment to decide whether to instruct the parties to fake to a set of inputs and outputs that is either globally consistent or globally inconsistent. Unframeability is captured by allowing the adversary to corrupt some of the parties and obtain their real local state. *Cooperative incoercibility* is obtained from the above definition by restricting the environment in two ways: first, it must give "tell the truth" instructions to either *all* parties, or else to *no* party; second, we prohibit party corruption.

As noted in [CG96], this definition of incoercibility immediately implies semi-honest adaptive security.

Obstacles to incoercibility: inversion and coordination.   We start by giving some intuition for why it is hard to build incoercible protocols. For instance, consider a two-party computation protocol based on Yao garbled circuits [Yao86], where the sender sends a garbled circuit, together with labels for the inputs of the sender and the receiver; the latter is sent via oblivious transfer (OT). If both the sender and the receiver become coerced and decide to lie about their inputs and outputs, then:

- The receiver needs to demonstrate how it obtained a (potentially fake) output of the OT corresponding to a different receiver bit that the actual on. At the same time, the receiver shouldn't be able to obtain the *true* OT output for that bit, as this would violate sender privacy.

- The sender needs to explain how the garbled circuit was generated, i.e. it needs to provide the generation randomness. However, the sender has already sent its own labels over the public channel, and now it has to come up with different generation randomness such that those same labels, which correspond to the true input, would now correspond to another, fake input.

  Furthermore, this generation randomness also has to be consistent with the labels of the receiver and fake input of the receiver, which the sender doesn't even know.

This example already demonstrates two difficulties with designing incoercible protocols. One is the problem of *inversion*, where some or all parties have to invert some randomized function $f(x; r)$ with respect to a different x' (like the generation of a garbled circuit)[5]. The other is a problem of *coordination*, where parties

---

[5]Note that in the model where not everybody is coerced, it is easy to avoid the inversion problem altogether by, e.g., secret-sharing r across all parties, thus guaranteeing that the coercer never gets to see r.

have to lie about their intermediate states *in a consistent way*, even though parties do not know fake inputs and outputs of each other.

These problem are reminiscent of the problems arising in the context of adaptive security. However, the adaptive security setting is significantly milder, since the "faking algorithm" is accessible only to the simulator, and can make use of the simulator's global view of the computation. In contrast, here the faking algorithm needs to be run only using the local view of a party. In addition, the adversary has access to the same faking algorithm held by the simulator, so the simulator cannot keep secret "trapdoors" related to the faking algorithm.

We note that the above challenges manifest themselves even in the relatively simple context of deniable encryption, or in other words the task of realizing the incoercible variant of the ideal secure message transmission functionality. Indeed, despite having a number of clean and modular constructions of protocols that realize secure message transmission from adaptively secure (i.e. non-committing) encryption schemes [BH92, CFGN96, DN00, CDMW09, HOR15, HORR16, YKT19], the only known fully deniable encryption protocol is both complex and non-modular [CPP20]: it is built from the ground up using obfuscation, and both the construction and its security proof are quite heavy.

Extending this direct approach to constructing general incoercible MPC protocols appears daunting. We thus take an alternative, more modular approach and use fully deniable encryption as a primitive in our protocols, along with other, more standard primitives.

Fully deniable encryption. A key component in all our constructions is fully deniable (interactive) encryption. Specifically, we use the scheme of [CPP20]. In a nutshell, this is a scheme which withstands coercion of both the sender and the receiver even in the off-the-record setting. The protocol requires a common reference string, takes 3 rounds, and assumes subexponential indistinguishability obfuscation and one way functions. Roughly, deniable encryption give the following security guarantee:

1. the adversary cannot distinguish whether it sees

   - real randomness s of the sender, real randomness r of the receiver, and the communication transcript for plaintext m, or

   - fake randomness of the sender s' consistent with fake m, fake randomness of the receiver r' consistent with fake m, and the communication transcript for plaintext m'.

2. (off-the-record setting) the adversary cannot distinguish whether it sees

   - real randomness s of the sender, fake randomness r' of the receiver consistent with m', and the communication transcript for plaintext m,

   - fake randomness of the sender s' consistent with m, real randomness of the receiver r consistent with m', and the communication transcript for plaintext m',

   - fake randomness s' of the sender consistent with m, fake randomness r' of the receiver consistent with m', and the communication transcript for plaintext m''.

   This should hold for any, potentially equal, m, m', m''.

1.1.1 Incoercible oblivious transfer and 2PC with short inputs

Incoercible oblivious transfer has the functionality of a standard oblivious transfer - i.e. it allows the receiver to obtain exactly one value x_b (corresponding to its own input b), out of two values x_0, x_1 held by the sender. However, it additionally provides security guarantees against a coercer: that is, even if the coercer demands to see all randomness used by both parties in the protocol, parties can successfully lie about their inputs. That is, the sender can claim that it used any, possibly different inputs x'_0, x'_1 (and provide convincing randomness supporting this claim). Similarly, the receiver can claim it used a possibly different input bit b', and received a different output x' of its choice.

This primitive can be constructed from any receiver-oblivious deniable encryption (DE) with public receiver deniability. Here "public receiver deniability" means that the faking algorithm of the receiver doesn't take true receiver coins as input (thus anyone can fake on behalf of the receiver). "Receiver-oblivious" DE means that the adversary cannot tell if the receiver messages were generated honestly (following the algorithm of DE), or instead chosen at random (in this case, we say that these messages were generated obliviously); further, this indistinguishability should hold even given fake random coins of the sender. We note that the deniable encryption of [CPP20] has public receiver deniability, and in the appendix A we show that it is also receiver-oblivious.

Theorem 1 *Any receiver-oblivious deniable encryption, which remains deniable even in the off-the-record setting and has public receiver deniability, can be converted into fully incoercible 1-out-of-m oblivious transfer, for any polynomial m, in a round-preserving way.*

The construction of incoercible OT is inspired by the construction of adaptively secure OT from non-committing (adaptively secure) encryption [CLOS02]. Namely, let x_0, x_1 be the inputs of the sender, and b be the input of the receiver. The parties should run in parallel two instances of DE: DE_0 and DE_1. The sender's input to each DE_i is x_i, for both $i = 0, 1$. The receiver should pick random r and generate messages of DE_b honestly (using r as randomness of the receiver in the protocol), while messages of DE_{1-b} should be generated by the receiver obliviously.

It is easy to see that the receiver can learn only x_b but not x_{1-b}, since the receiver knows r, which allows it to decrypt DE_b, but doesn't know randomness for DE_{1-b} and therefore cannot decrypt it. The sender, in turn, doesn't learn the receiver bit b, since it doesn't know which execution was generated obliviously by the receiver. Further, this OT is indeed incoercible: the sender can directly use deniability of DE to claim that different inputs x'_0, x'_1 were sent. The receiver can lie about its input b by claiming that DE_b was generated obliviously, and by presenting fake r' as randomness for DE_{1-b}. This fake r' can be generated by using the faking algorithm on DE_{1-b} and y', where y' is the desired fake output of the oblivious transfer. Note that the receiver doesn't know true coins for obliviously generated DE_{1-b}, but it can generate fake r' anyway due to the fact that receiver deniability is public.

This construction can be extended to 1-out-of-m incoercible OT in a straightforward way.

Incoercible 2PC for short inputs from incoercible OT.   Recall that, when the number m of possible inputs of some party is polynomial, standard 1-out-of-m OT immediately implies general 2PC [GMW87]: The OT sender should input to the OT m possible values of $f(x, y)$, corresponding to m possible values of the receiver's input y, and a single sender's input x. Using incoercible 1-out-of-m OT in this protocol immediately makes the resulting 2PC protocol incoercible.

Incoercible MPC from OT? Despite the fact that standard OT implies general secure multi-party computation [GMW87], it is not clear whether *incoercible* OT implies *incoercible* MPC as well. In particular, simply plugging (even ideal) incoercible OT into the protocol of [GMW87] doesn't seem to result in an incoercible protocol, even just for two parties. The problem here is the following: recall that this protocol works by letting the parties compute additive secret shares of each wire of the circuit of $f(x_1, x_2)$. On one hand, since in the normal execution two shares add up to the value of the wire of $f(x_1, x_2)$, the same should hold in the fake case: fake secret shares should add up to the value of the wire of $f(x_1', x_2')$. However, it is not clear how, upon coercion, parties can compute these fake shares *locally*, without the knowledge of the other party's input.

1.1.2   Incoercible MPC

A natural starting point for building an incoercible MPC is to have the parties run any secure MPC protocol, where each message is encrypted under a separate instance of deniable encryption. If in addition the parties are allowed to communicate outside of the view of the adversary - e.g. by meeting physically - and if they are comfortable sharing their fake inputs with each other, this method immediately gives incoercible MPC. Indeed, upon coercion parties can use their out-of-band channel to all agree on some transcript $\mathsf{tr}' = \mathsf{tr}(\{x_i', r_i'\})$ of an underlying MPC executed on their fake inputs. When coerced, each party can use deniability of encryption to lie (by presenting consistent randomness and keys of deniable encryption) that it sent and received messages of tr'. In addition, each party should claim that x_i', r_i' are the true input and randomness which it used to compute the messages of tr'.

However, this protocol fails when no out-of-band interaction is possible, since the parties do not have means to agree on tr'. To fix this problem, we combine deniability with adaptive security. That is, we use MPC which is adaptively secure and has a special property called *corruption oblivious simulation* (defined in [BCH12] in a setting of leakage tolerance). Roughly, it means that there is a "main" simulator which simulates the transcript, and in addition each party has its own, "local" simulator which simulates the coins of that party, using that party's inputs only and the state of the "main" simulator (but not the inputs of other parties). If parties had a way to agree on the same simulation randomness r_{Sim}, then upon coercion, they could do the following: First they should run the main simulator on r_{Sim} to generate (the same) simulated transcript tr' of an underlying adaptive MPC, and then each party should use its own local simulator to locally compute fake coins consistent with this simulated transcript and its own input. Finally, as before, each party can use deniability of encryption to claim that the messages of tr' were indeed sent.

It remains to determine how the parties agree on the random coins r_{Sim} of the main simulator. A natural approach to do this is to let one of the parties (say, the first) choose r_{Sim} at random and send it, encrypted under deniable encryption, to each other party at the beginning of the protocol, for case that they need to fake later. However, this introduces another difficulty: now the adversary can demand to see r_{Sim}, and revealing it would allow the adversary to check that the transcript was simulated and thus detect a lie. Therefore, instead of sending r_{Sim}, the first party should send randomly chosen seed s to all other parties. This seed is not used by parties in the execution of the protocol. However, upon coercion each party can use a pseudorandom generator to expand s into a string $r_{\mathsf{Sim}}||s'$, where r_{Sim}, as before, is used to produce the same simulated transcript of an adaptive MPC, and s' is what parties will claim as their fake seed (instead of a true seed s). Note that it is safe to reveal s' to the adversary, since $s'||r_{\mathsf{Sim}}$ is pseudorandom, and therefore s' cannot help the adversary to indicate in any way that tr' was simulated.

We underline that security of this protocol is only maintains in the cooperative setting. As a result, this protocol is useful in scenarios where parties "work together" and are interested in keeping all their inputs secret, rather than turn against each other trying to make sure others get caught cheating. We note however that the protocol remains secure even if *inputs* of some parties are real and inputs of some other parties are fake - as long as *randomness* of all parties is fake. Indeed, it might happen so that a certain party is not interested in lying about its input, but still wishes the whole group of people to succeed in deceiving; then this party may provide fake randomness for its real input, thus not ruining the joint attempt to deceive, while achieving its own goals[6]. Further, this protocol maintains the best possible security even in the case when the claimed inputs and outputs are clearly inconsistent. The following informal theorem summarizes the discussion:

Theorem 2 *Assuming deniable encryption and adaptively secure MPC protocol with a global CRS and corruption-oblivious simulation, there exist cooperatively incoercible secure function evaluation protocols.*

We spell out the protocol. We need the following ingredients:

- 2-round adaptively secure MPC aMPC with global CRS[7] and corruption-oblivious simulator, e.g. that of [CPV17].

- 3-round delayed-input[8] deniable encryption DE, e.g. that of [CPP20]. While that construction is not delayed-input, we observe that it is easy to turn any deniable encryption into its delayed-input version. This can be done by letting the sender send a randomly chosen key k using deniable encryption, and also send $m \oplus k$ in the clear at the last round.

The protocol then proceeds as follows, with 4 rounds of communication:

1. In rounds $1 - 3$ parties exchange the messages of the *first* round of aMPC, encrypted under point-to-point deniable encryption.

2. In rounds $2 - 4$ parties exchange the messages of the *second* round of aMPC, encrypted under point-to-point deniable encryption. It is important that deniable encryption requires its input only by the last round, since parties receive the messages of the first round of aMPC only after round 3.

3. In rounds $2 - 4$ party 1 sends to each party randomly chosen seed, encrypted under point-to-point deniable encryption. Note that each party receives the same value of seed.

After round 4, parties learn all messages of aMPC and therefore can compute the output. Note that our protocol is delayed input, since inputs are required only by round 3. Upon coercion, each party first computes fake transcript tr' of aMPC. tr' is computed by running the "main" simulator of aMPC on r_{Sim}, where r_{Sim} is obtained by expanding seed into $\mathsf{seed}'||r_{\mathsf{Sim}}$ using a prg. (Note that parties use the same r_{Sim} and therefore obtain the same tr' upon coercion). Next, each party can use its local simulator to produce fake coins consistent with tr' and fake input x'. Therefore, each party can claim that the transcript of the underlying protocol was tr', and this claim will be consistent with party's own fake input, and across different parties. Finally, each party should claim that the seed value sent by party 1 was in fact seed'.

---

[6]Note that this scenario highlights a subtle but important difference between the modelling of coercion in [CG96] and [CGP15]. Indeed, in [CG96], if the party is given a real input, it has to provide its true randomness.

[7]The CRS of the protocol is said to be global, if the simulator can simulate the execution, *given* the CRS (as opposed to generating the CRS on its own, possibly from a different distribution, or with underlying trapdoors).

[8]That is, only the third message of the sender depends on the plaintext.

Note that our construction crucially uses the fact that underlying adaptive MPC has *global* CRS. Indeed, this allows to put this CRS as part of the final CRS of the protocol, and lets parties simulate the transcript of underlying adaptive MPC with respect to that CRS. Had the CRS been local, parties would have to generate it during the protocol and thus eventually provide the adversary with the generation coins; yet, security of protocols with local CRS usually holds only as long as the generation randomness of this CRS remains private.

1.1.3 Impossibility of incoercible MPC with lazy parties

Impossibility of incoercible MPC with lazy parties. We show that unlike 2-party protocols, multiparty protocols with some communication structure cannot be incoercible (this holds even against coercion of only 2 parties). Concretely, let us say that a party is *lazy*, if it only sends its messages in the first and the last round of a protocol, but doesn't send anything in intermediate rounds (if any). In particular, in all 2 round protocols all parties are lazy by definition. We show that coercing a lazy party and some output-receiving party allows to learn information about inputs of other parties, therefore rendering the protocol insecure for most functions:

Theorem 3 *(Informal) Assume there exists an n-party protocol withstanding 2 corruptions and 1 coercion for computing function f with a lazy party, where $n \geq 3$. Then the function f is such that for any inputs $x_1, \ldots, x_n$ it is possible, given x_1, x_n, and $f(x_1, \ldots, x_n)$, to compute $f(x, x_2, \ldots, x_n)$ for any x.*

We consider this negative result to be especially important in light of the fact that building fully incoercible protocols may require complicated obfuscation-based constructions. For instance, consider the following natural attempt to build a 3-round fully incoercible protocol. Take deniable encryption of [CPP20] which essentially lets the sender send an encryption of a plaintext together with some auxiliary information, which the receiver can decrypt using an obfuscated decryption program. This protocol features a "ping-pong" communication pattern, with a total of 3 messages sent between a sender and a receiver. One could attempt to turn it into MPC with a similar "ping-pong" communication pattern by letting $n-1$ senders $P_1, \ldots, P_{n-1}$ send its input to a single receiver P_n in a similar manner, and let the obfuscated evaluation program of the receiver decrypt the messages and evaluate the result. While this approach sounds very plausible and appealing in a sense that it potentially requires only minor modifications of the construction of deniable encryption, our impossibility result implies that such protocol cannot be incoercible.

Finally, it is interesting to note that this impossibility result is "tight" both with respect to the number of participants n, and with respect to coercion operation (as opposed to adaptive corruption). Indeed, there exists a 3-round *two-party* incoercible protocol (e.g. our OT-based protocol), and a 3-round multi-party *adaptively secure* protocol [DKR14], which features such a "ping-pong" communication pattern.[9]

To get an idea of why impossibility holds, consider standard MPC with a super-lazy party who only sends its messages in the very last round; clearly, such a protocol is insecure, since the adversary who corrupts this party together with some output-receiving party can rerun the protocol on many inputs of the lazy party and therefore infer some information about the inputs of uncorrupted parties.

Such an attack in the standard MPC case doesn't work when a lazy party sends messages in two rounds of the protocol. However, we show that in case of incoercible protocols there is a way for a lazy party to modify

---

[9]Note that formally speaking, the protocol of [DKR14] takes 4 rounds; however, the receiver learns the output already after round 3. The 4-th round is only required to send this output back to everyone.

its last message such that the protocol now thinks that a different input is used - despite the fact that its first message still corresponds to the original input. With this technique in place we can mount the same attack as described before. This technique is based on the observation in [CPP20] that sender-deniability in any deniable encryption implies that a party can "fool" its own protocol execution into thinking that a different input is being used. We refer the reader to section 5 for details.

1.1.4 On the difference between $n = 2$ and $n \geq 3$ parties

Since some of our results only hold for the case of $n = 2$ or $n \geq 3$, we briefly comment on the difference.

Impossibility of incoercible MPC with lazy parties.   This result inherently uses the fact that the number of parties is at least 3. The reason is that our attack allows the adversary, who corrupts P_1 and P_n, learn $f(x, x_2, \ldots, x_n)$ for any x (in addition to x_1, x_n, and $f(x_1, \ldots, x_n)$). As we can see, in case $n = 2$, $f(x, x_2, \ldots, x_n)$ becomes $f(x, x_2)$, which the adversary learns anyway upon corruption of both parties; thus this attack doesn't give the adversary any power when $n = 2$. And indeed, our OT-based 2PC has the property that the sender is lazy (i.e. it sends messages only in the first and in the last round).

Off-the-record incoercibility.   Our OT-based 2PC remains incoercible even in the off-the-record scenario, whereas our MPC doesn't. However, we don't think that there is some inherent obstacle to obtaining off-the-record security in the case of many parties. We believe that it could be done, especially if each action of each party is protected by obfuscation. (The reason why our 2PC protocol is incoercible in the off-the record scenario is because so is the underlying deniable encryption scheme; in turn, the latter is incoercible because all actions of parties are protected by obfuscation).

1.1.5 Discussion, open problems, and future work

Our results naturally lead to the following open problems:

- *Round complexity:* is it possible to build an incoercible protocol, withstanding coercion of all parties, for general functions in 3 rounds? (Recall that 3 rounds is the minimum s [BNNO11].)

- *Full incoercibility:* Is it possible to obtain a protocol which withstands coercion of all parties and remains incoercible even in the off-the-record setting - with any number of rounds?

The protocols in this paper follow a blueprint of composing deniable encryption with non-deniable primitives, resulting in a simple and clean protocol design. However, it could be problematic to use this approach for answering the questions listed above. The reason is the following. Since incoercible MPC implies deniable encryption, any construction of incoercible MPC:

- either has to use some construction of deniable encryption,

- or has to build deniable encryption from scratch, at least implicitly.

As we explain in more detail next, improving on our results would likely require the latter. Given the complexity of the [CPP20] construction, such modification will require better tools for abstracting and mudularizing this construction and similar ones.

On obtaining full incoercibility for $n \geq 3$. A natural attempt to build a fully incoercible protocol is to combine deniable encryption (secure even in the off-the-record setting) with other, weaker-than-incoercible primitives (e.g. standard MPC). Unfortunately, this is unlikely to help. Indeed, a very simple argument made by [AOZZ15] shows that in any construction of off-the-record incoercible MPC with the help of secure channels, parties have to use these (perfectly deniable!) channels in an inherently non-deniable way: that is, if a party sends (receives) a message M via secure channel during the protocol, then its faking algorithm cannot instruct this party to lie about M^{10}. This can be informally interpreted as follows: in any incoercible protocol which uses deniable encryption, deniable encryption can be replaced with standard encryption such that the protocol still remains incoercible[11]. This in turn indicates that such a protocol would have to be incoercible to begin with.

1.2 Related Work

Receipt-free and coercion-resistante voting schemes. The study of incoercibility has started with the specific task of *voting*. Indeed, a long line of work, including [BT94, BMN+09] and many others, studies incoercibility in the context of voting schemes, including notions such as *receipt-freeness* and coercion resistance. It is important to stress that, while these works share the overall concerns and goals as the presnt work, there are a number of significant differences:

First, while this work concentrates on the case where the parties actually follow the protocol until the point of coercion, which happens only after the after the protocol execution has completed, these works consider the case where parties may deviate from the protocol instructions right from the start. In accordance, their solutions inherently include a non-computational, physical component (a "voting booth"). Second, their notion of coercion resistance imposes restrictions both on the protocol and on the evaluated function. In contrast, we provide general tools for incoercible secure evaluation of *any* function, thus allowing for separating the analysis of the evaluated function from the security analysis of the implementation.

Prior work on generic incoercible MPC. The prior work on generic incoercible MPC can be split into two parts, depending on whether it focuses on *semi-honest* or *active* coercion (in the language of [AOZZ15]). Intuitively, a coercer is semi-honest if it lets the party participate in the protocol as prescribed (by following the instructions of the protocol), but after that demands to see the entire view of that party and checks whether it matches the claimed input of that party. In contrast, an *active* coercer assumes full control over the party and in particular may instruct the party to deviate from the protocol, in order to make it harder for the party to deceive the coercer.

As already noted in [BT94] in the context of secure voting, active coercion is clearly unachievable with cryptography alone: coerced parties have no hope of lying about their inputs if the adversary watches over their shoulder during the computation. As a result, security against active coercion requires some form of physical unaccessibility assumption. Indeed, to come up with the protocol secure against active coercion, [AOZZ15] makes use of a stateful hardware token which can generate keys, distribute them to all parties, and encrypt.

---

[10]Roughly, this is because said party doesn't know whether its peer is lying or telling the truth; it could be telling the truth, thus revealing true M, and from definition of off-the-record deniability, their joint state should look valid even in the case when the party is lying and its peer is telling the truth - as long as their inputs and outputs are consistent.

[11]We underline again that this is an informal statement - indeed, such a statement is tricky to even formalize, let alone prove.

In contrast, semi-honest incoercibility is well within the reach of "digital cryptography", without the need to assume inaccessible hardware: sender-deniable encryption and encryption deniable for both parties was constructed by [SW14] and [CPP20] respectively, from indistinguishability obfuscation and one-way functions, and it was shown back in 1996 how to transform any sender-deniable encryption into incoercible MPC which withstands coercion of up to half participants [CG96]. The protocols of [CGP15], [DKR14], originally devised as adaptively secure protocols withstanding corruption of all parties, can withstand coercion of a single party.

Active incoercibility has been studied in the works of [MN06] for the specific functionality of voting, and later the works of [UM10], [AOZZ15], which develop a framework for UC incoercibility and give constructions of UC-incoercible protocols for general functionalities, using inaccessible hardware.

Note that, although from a practical standpoint active incoercibility is stronger and more desirable than semi-honest one, from theoretical perspective semi-honest and active incoercibility are two completely different and incomparable problems. Indeed, achieving semi-honest incoercibility requires solving the problem of "inverting the computation" - i.e. finding randomness which makes some computation appear to stem from a different input (note that this problem is also interesting on its own, without its connection to incoercibility). Active incoercibility, as discussed above, inherently requires inaccessible hardware to hide parts of the computation and thus avoids the inverting problem altogether; instead, the goal there is to ensure that the active coercer cannot force parties to output something committing, while making the underlying physical assumptions as realistic as possible.

Impossibility results.　[CG96] shows that semi-honest incoercible computation is not achievable against unbounded adversaries; this impossibility holds even in the presence of private channels. To the best of our knowledge, in the computational setting no impossibility results specific to incoercible MPC were known. However, the impossibility of non-interactive (i.e. 2-message) receiver-deniable encryption [BNNO11] immediately implies that 2-round incoercible MPC is impossible, even against coercion of a single party which receives the output[12] (in particular, the 2-round protocol of [CGP15] only withstands coercion of a party which doesn't receive the output); this impossibility holds for all functions which imply a bit transmission.

On the difference between the definitions of incoercible MPC in [CG96] and [CGP15].　In this work we use the definition of incoercible computation from [CGP15]. We briefly explain how it differs from the one in [CG96]. The definitions of [CG96] and [CGP15] are conceptually similar but differ in case when an environment instructs a party to fake, but sets its fake input and output to be exactly the same as its real input and output. In this case the definition in [CG96] instructs the party to output its true randomness, while the definition in [CGP15] instructs the party to run the fake algorithm anyways and output the resulting fake randomness.

This difference may appear minor - indeed, if a party is not going to lie about its inputs nor outputs, why fake the randomness? Nevertheless, there are situations when a party may want to fake its randomness anyways.

---

[12]Indeed, any incoercible protocol for a message transmission functionality can be turned into a 2-message receiver-deniable encryption, by letting the party R which receives the output be a receiver of deniable encryption, and letting the sender run the MPC protocol on behalf of all other parties. In particular, the first message (sent by the receiver) will consist of all messages sent by R in the first round of the protocol, and the second message (sent by the sender) will consist of all messages sent to R in rounds 1 and 2. Messages sent by R in round 2 of MPC protocol do not have to be sent, since S doesn't receive the output, nor does S have to deny later.

Indeed, as we discuss in the technical overview, our incoercible MPC protocol only retains its incoercibility properties as long as *all* parties disclose their fake coins to the coercer. In particular, there may be a party which has no interest in lying about its own input, but which anticipates that other participants may need to lie about theirs, and which thus decides to give out its fake randomness to make sure its true randomness doesn't compromise other parties' security.

2 Preliminaries

2.1 Incoercible computation

We use the definition of incoercible computation from [CGP15], which can be regarded as a re-formulation of the definition of [CG96] within the UC framework. (We note that the formulations of [MN06, AOZZ15] are similar to and consistent with the one we use, with the exception that they allow also Byzantine corruptions and incorporate modeling of ideally opaque hardware.) Specifically, for full incoercibility, we consider protocols that, in addition to the standard corruption instruction from the adversary, take a coercion instruction; upon receiving this instruction, the party notifies a predetermined external entity (say, its "caller" via subroutine output) that it was coerced and expects an instruction to either "tell the truth", in which case it reveals its entire local state to the adversary, or "fake to input x and output y", in which case a faking algorithm, which is provided as part of the protocol, is run on inputs x, y and the current local state. The output of the faking algorithm is presented to the adversary as the (fake) internal state of the party.

Cooperative environments. A cooperative environment is one that's guaranteed: (a) not to *corrupt* parties, and (b) to either instruct all coerced parties to "tell the truth", or else to instruct neither of the coerced/corrupted parties to "tell the truth." (In the latter case the environment gives each coerced party values x, y and the party is expected to generate fake randomness that explains the messages sent by the party as consistent with input x and output y.

Coercion-aware ideal functionalities. An ideal functionality can now enforce incoercibile imlpementations via the following mechanism: When asked by the adversary (or, simulator) to coerce a party P, the ideal functionality outputs a request to coerce P to the said external entity, in the same way as done by the protocol. If the response is "fake to input x and output y, then the pair x, y is returned to the adversary. If the response is "tell the truth" then the actual input x and output y are returned to the adversary. *Crucially, the simulator is not told whether the returned values are real or fake.*

This behavior is intended to mimic the situation where the computation is done "behind closed doors" and no information about it is ever exposed, other than the inputs and outputs of the parties. In particular, such an ideal functionality does not prevent situations where the outputs of the parties are globally inconsistent with their inputs, or where a certain set of inputs of the parties are inconsistent with auxiliary information that's known outside the protocol execution. Indeed, the only goal here is to guarantee that any determination made by an external coercer (modeled by the environment) after interacting with the protocol, could have been done in the ideal model, given only the claimed inputs and outputs.

Figures 1, 2 and 3 depict incoercible variants of the standard ideal functionalities for secure message transmission, oblivious transfer, and multiparty function evaluation, respectively.

Functionality $\mathcal{F}_{\text{IMT}}$

- Upon receiving input $(\texttt{Send}, sid, R, m)$ from party S, where R is an identity for the intended receiver, send $(sid, S, R, |m|)$ to the adversary. When receiving $\texttt{ok}$ from the adversary, output $(\texttt{Receive}, sid, S, m)$ to R.
- Upon receiving $(\texttt{Coerce}, sid, P)$ from the adversary, where $P \in \{S, R\}$, output $(\texttt{Coerce}, sid)$ to P. Upon receiving V from P do: If $V = (\texttt{tell-truth})$ then send m to the adversary. If $V = (\texttt{fake-to}, m')$ then send m' to the adversary.
- Upon receiving $(\texttt{Corrupt}, sid, P)$ from the adversary, where $P \in \{S, R\}$, output $(\texttt{Corrupt}, sid)$ to P, and send m to the adversary.

Figure 1: The Incoercible Message transmission Functionality $\mathcal{F}_{\text{IMT}}$.

Functionality $\mathcal{F}_{\text{IOT}}$

- Upon receiving input $(\texttt{OT-Sender}, sid, R, (m_0, m_1))$ from party S, where R is an identity for the intended receiver, send (sid, S, R) to the adversary. When receiving $\texttt{ok}$ from the adversary, output $(\texttt{OT-Receiver}, sid, S)$ to R.
- Upon receiving input $(\texttt{OT-Receiver}, sid, b)$ from R, send sid to the adversary. When receiving $\texttt{ok}$ from the adversary, output $(\texttt{OT-Receiver}, sid, m_b)$ to R.
- Upon receiving $(\texttt{Coerce}, sid, P)$ from the adversary, where $P \in \{S, R\}$, output $(\texttt{Coerce}, sid)$ to P. Upon receiving V from P do: If $V = (\texttt{tell-truth})$ then send P's input and output to the adversary. If $V = (\texttt{fake-to}, v)$ then send v to the adversary.
- Upon receiving $(\texttt{Corrupt}, sid, P)$ from the adversary, where $P \in \{S, R\}$, output $(\texttt{Corrupt}, sid)$ to P, and send P's input and output to the adversary.

Figure 2: The Incoercible Oblivious Transfer Functionality $\mathcal{F}_{\text{IOT}}$.

As discussed in the introduction, we note that the present definitional framework is only concerned about incoercible implementations of a given ideal functionality. The do not adress the question of whether the ideal functionality itself is susceptible to coercion.

We say that π is a fully incoercible message transmission protocol if π UC-realizes $\mathcal{F}_{\text{IMT}}$. If π UC-realizes $\mathcal{F}_{\text{IMT}}$ only with respect to cooperative environments then π is a cooperatively incoercible message transmission protocol.

Similarly, we say that π is a fully incoercible oblivious transfer protocol if π UC-realizes $\mathcal{F}_{\text{IOT}}$. If π UC-realizes $\mathcal{F}_{\text{IOT}}$ only with respect to cooperative environments then π is a cooperatively incoercible oblivious transfer protocol.

Similarly, say that π is a fully incoercible function evaluation protocol if π UC-realizes $\mathcal{F}_{\text{IFE}}$. If π UC-realizes $\mathcal{F}_{\text{IFE}}$ only with respect to cooperative environments then π is a cooperatively incoercible function evaluation protocol.

2.2 Deniable Encryption

Since deniable encryption is a main building block in our constructions, we review the definition, taken from [CPP20]. We note that the definition below implies the simulation-based notion of deniable communication. That is, if π satisfies Definition 1 , then π is an off-the-record incoercible secure message transmission

protocol. However, we note that we require additional properties of receiver-obliviousness and public-receiver-deniability from deniable encryption for our OT-based protocol

We denote by $\pi(\text{CRS}, s, r, m)$ the messages sent by parties when running deniable encryption protocol on input m with randomness s of the sender and r of the receiver.

Definition 1 Deniable bit encryption in the CRS model. $\pi = (\text{DE.Setup}, \text{DE.msg1}, \text{DE.msg2}, \text{DE.msg3}, \text{DE.Dec}, \text{DE.SFake}, \text{DE.RFake})$ *is a 3-message deniable interactive encryption scheme for message space* $\mathcal{M} = \{0, 1\}$, *if it satisfies the following correctness and deniability properties:*

- **Correctness:** *There exists negligible function $\nu(\lambda)$ such that for at least $(1-\nu)$-fraction of randomness* $r_{\text{Setup}} \leftarrow \{0, 1\}^{|r_{\text{Setup}}|}$ *the following holds: let* $\text{CRS} \leftarrow \text{DE.Setup}(r_{\text{Setup}})$. *Then for any $m \in \mathcal{M}$* $\Pr[m' \neq m : s \leftarrow \{0, 1\}^{|s|}, r \leftarrow \{0, 1\}^{|r|}, \text{tr} \leftarrow \pi(\text{CRS}, s, r, m), m' \leftarrow \text{DE.Dec}(\text{CRS}, r, \text{tr})] \leq \nu(\lambda)$, *where the probability is taken over the choices of s and r.*

- **Bideniability:** *No PPT adversary* Adv *wins with more than negligible advantage in the following game, for any $m_0, m_1 \in \mathcal{M}$:*

    1. *The challenger chooses random r_{Setup} and generates $\text{CRS} \leftarrow \text{DE.Setup}(r_{\text{Setup}})$. It also chooses a bit b at random.*

    2. *If $b = 0$, then the challenger generates the following variables:*

        (a) *It chooses random s^*, r^* and computes $\text{tr}^* = \pi(\text{CRS}, s^*, r^*, m_0)$.*

        (b) *It gives the adversary $(\text{CRS}, m_0, m_1, s^*, r^*, \text{tr}^*)$.*

    3. *If $b = 1$, then the challenger generates the following variables:*

        (a) *The challenger chooses random s^*, r^* and computes $\text{tr}^* \leftarrow \pi(\text{CRS}, s^*, r^*, m_1)$;*

        (b) *It sets $s' \leftarrow \text{DE.SFake}(\text{CRS}, s^*, m_1, m_0, \text{tr}^*; \rho_S)$ and $r' \leftarrow \text{DE.RFake}(\text{CRS}, r^*, m_1, m_0, \text{tr}^*; \rho_R)$, for randomly chosen ρ_S, ρ_R.*

        (c) *It gives the adversary $(\text{CRS}, m_0, m_1, s', r', \text{tr}^*)$.*

    4. Adv *outputs b' and wins if $b = b'$.*

- **Off-the-record deniability:** *No PPT adversary* Adv *wins with more than negligible advantage in the following game, for any* $m_0, m_1, m_2 \in \mathcal{M}$:

  1. *The challenger chooses random* r_{Setup} *and generates* $\mathsf{CRS} \leftarrow \mathsf{DE.Setup}(r_{\mathsf{Setup}})$. *It also chooses* $b \in \{0, 1, 2\}$ *at random.*

  2. *If* $b = 0$, *then the challenger generates the following variables:*

     (a) *The challenger chooses random* s^*, r^* *and computes* $\mathsf{tr}^* \leftarrow \pi(\mathsf{CRS}, s^*, r^*, m_0)$;

     (b) *It sets* $r' \leftarrow \mathsf{DE.RFake}(\mathsf{CRS}, r^*, m_0, m_1, \mathsf{tr}^*; \rho_R)$, *for randomly chosen* ρ_R.

     (c) *It gives the adversary* $(\mathsf{CRS}, m_0, m_1, m_2, s^*, r', \mathsf{tr}^*)$.

  3. *If* $b = 1$, *then the challenger generates the following variables:*

     (a) *The challenger chooses random* s^*, r^* *and computes* $\mathsf{tr}^* \leftarrow \pi(\mathsf{CRS}, s^*, r^*, m_1)$;

     (b) *It sets* $s' \leftarrow \mathsf{DE.SFake}(\mathsf{CRS}, s^*, m_1, m_0, \mathsf{tr}^*; \rho_S)$ *for randomly chosen* ρ_S.

     (c) *It gives the adversary* $(\mathsf{CRS}, m_0, m_1, m_2, s', r^*, \mathsf{tr}^*)$.

  4. *If* $b = 2$, *then the challenger generates the following variables:*

     (a) *The challenger chooses random* s^*, r^* *and computes* $\mathsf{tr}^* \leftarrow \pi(\mathsf{CRS}, s^*, r^*, m_2)$;

     (b) *It sets* $s' \leftarrow \mathsf{DE.SFake}(\mathsf{CRS}, s^*, m_2, m_0, \mathsf{tr}^*; \rho_S)$ *for randomly chosen* ρ_S *and* $r' \leftarrow \mathsf{DE.RFake}(\mathsf{CRS}, r^*, m_2, m_1, \mathsf{tr}^*; \rho_R)$ *for randomly chosen* ρ_R.

     (c) *It gives the adversary* $(\mathsf{CRS}, m_0, m_1, m_2, s', r', \mathsf{tr}^*)$.

  5. Adv *outputs* b' *and wins if* $b = b'$.

Receiver-obliviousness. We say that deniable encryption satisfies **receiver-obliviousness**, if no adversary can win the following game, for any m in the plaintext space:

1. The challenger chooses random r_{Setup} and generates $\mathsf{CRS} \leftarrow \mathsf{DE.Setup}(r_{\mathsf{Setup}})$. It also chooses a bit b at random.

2. If $b = 0$, then the challenger generates the following variables:

   (a) It chooses random s^*, r^* and computes $\mathsf{tr}^* = \pi(\mathsf{CRS}, s^*, r^*, m)$.

   (b) It gives the adversary $(\mathsf{CRS}, m, s^*, \mathsf{tr}^*)$.

3. If $b = 1$, then the challenger generates the following variables:

   (a) The challenger chooses random s^*, r^* and computes tr^* as follows. It sets $a1 = \mathsf{DE.msg1}(\mathsf{CRS}, s^*, m)$, chooses $a2$ at random, and sets $a3 = \mathsf{DE.msg3}(\mathsf{CRS}, s^*, m, a1, a2)$. It sets $\mathsf{tr}^* = (a1, a2, a3)$.

   (b) It gives the adversary $(\mathsf{CRS}, m, s^*, \mathsf{tr}^*)$.

4. Adv outputs b' and wins if $b = b'$.

In other words, the adversary cannot tell if messages of the receiver are honestly generated or simply chosen at random (as long as the adversary doesn't see the random coins of the receiver).

We observe that any protocol that satisfies the requirements of Definition 1 is a fully incoercible message transmission protocol as defined above.

Theorem 4 *Let π be a protocol that satisfies Definition 1. Then π UC-realizes $\mathcal{F}_{\text{IMT}}$.*

Public receiver deniability. We say that receiver deniability is public, if $\mathsf{DE.RFake}(\mathsf{CRS}, m', \mathsf{tr}; \cdot)$ only takes as input the CRS, the fake plaintext, the transcript of the execution, and its own random coins. In other words, true random coins r of the receiver are not required to run RFake.

Deniable encryption with public receiver deniability is built in [CPP20]. In appendix A we note that it also satisfies receiver-obliviousness.

2.3 Adaptively Secure Protocols with Corruption-Oblivious Simulation

Another concept that will be useful to us is adaptive security with coercion-oblivious simulation. Spcifically we consider adaptively secure protocols [CFGN96] where the simulator can be splitted in several parts, as follows:

- The "main" simulator $\mathsf{Sim}(r_{\mathsf{Sim}}) \to (\mathsf{tr}, \mathsf{state})$, which simulates the transcript tr of the protocol, and also generates a state to be used by "local" simulators;

- "Local" simulators $\mathsf{Sim}_i(\mathsf{state}, x_i, y_i; \cdot) \to r_i$, where each Sim_i takes as input the state of the main simulator, input and output of party i, and possibly its own random coins. It outputs r_i, a simulated randomness of party i.

We call such simulation coecrcion-oblivious. The security requirement is now stated as follows: the transcript of a real protocol execution, together with randomness of each party, should be indistinguishable from a simulated transcript, together with simulated randomness of each party. Note that this requirement is not trivial, since simulators have to work "locally" (without knowing other parties' inputs) and still have to jointly generate a consistently-looking picture. We refer the reader to [BCH12] for formal treatment.

A 2-round adaptively secure MPC protocol with corruption-oblivious simulator appears in [CPV17]. That protocol uses indistinguiahability obfuscation and one way functions.

3 Incoercible Oblivious Transfer

In this section we describe our construction of incoercible oblivious transfer. As noted in the introduction, such a protocol immediately implies incoercible 2PC for the case where one of the parties has polynomial input space.

Incoercible Oblivious Transfer

The CRS: $\mathsf{CRS} = \mathsf{CRS_{DE}}$, where $\mathsf{CRS_{DE}}$ is a CRS of deniable encryption with receiver-obliviousness, and public receiver deniability.

Inputs: inputs x_0, x_1 of the sender S; input bit b of the receiver R.

The protocol:

The sender chooses random coins s_0, s_1 for two executions of deniable encryption, where S acts as a sender. The receiver chooses randomness r for a single execution of deniable encryption where it acts as a receiver. The sender and the receiver run two instances of deniable encryption, DE_0 and DE_1, in parallel. Here:

- In each execution i, for $i = 0, 1$, the sender computes its messages by honestly running the code of deniable encryption on its input x_i, randomness s_i, and the transcript so far;
- In the execution b the receiver computes its messages by honestly running the code of deniable encryption on its randomness r and the transcript so far. In the execution $1 - b$ the receiver instead generates all its messages at random, using randomly chosen $\tilde{r}$.

At the end of both executions, the receiver sets its output in the protocol to be $\mathsf{DE.Dec}(r; \mathsf{DE}_b)$.

Faking procedure of the sender S

Inputs: fake inputs x_0', x_1' of the sender, true inputs and randomness x_0, s_0, x_1, s_1 of the sender, the protocol transcript $(\mathsf{DE}_0, \mathsf{DE}_1)$, and the CRS.

In order to fake, the sender runs the faking algorithm of deniable encryption for each execution, i.e. computes $s_i' \leftarrow \mathsf{DE.SFake}(s_i, x_i, x_i', \mathsf{DE}_i; \cdot)$ for both $i = 0, 1$. It gives s_0', s_1' to the adversary.

Faking procedure of the receiver R

Inputs: fake input b' and fake output x' of the receiver, true inputs and randomness $b, r, \tilde{r}$ of the receiver, the protocol transcript $(\mathsf{DE}_0, \mathsf{DE}_1)$, and the CRS.

In order to fake, the receiver claims that messages of the receiver in execution $1 - b'$ were generated at random, and sets fake $\tilde{r}'$ to be the concatenation of these receiver messages. Next, it uses public deniability of the receiver to compute $r' \leftarrow \mathsf{DE.RFake}(x', \mathsf{DE}_{b'}; \cdot)$. It gives $r', \tilde{r}'$ to the adversary.

Figure 4: Incoercible Oblivious Transfer.

3.1 Protocol Description

For simplicity, we consider 1-out-of-2 OT (the construction can be generalized to 1-out-of-n OT in a straightforward way), and we also assume that all inputs are bits. Our protocol is described on fig. 4. It requires a special deniable encryption (DE) scheme, where deniability of the receiver is public (i.e. the faking algorithm of the receiver doesn't take receiver's true coins as input), and which satisfies receiver-obliviousness, i.e. the real transcript is indistinguishable from a transcript where receiver simply generated all its messages at random. As noted in [CPP20], their DE protocol satisfies public receiver deniability. In appendix A we note that this protocol is also receiver-oblivious.

Before stating the theorem, we remind that we consider the model of semi-honest coercions of potentially all parties, and we assume that all coercions happen after the protocol finishes. We refer the reader to section 2 for a description of our coercion model.

Theorem 5 *Assume* DE *is an interactive deniable encryption scheme which satisfies public receiver deniability and receiver obliviousness, and remains deniable even in the off-the-record scenario. Then the protocol*

on fig. 2 is a semi-honest, fully incoercible oblivious transfer protocol.

3.2 Proof of the Theorem

Correctness. Correctness immediately follows from correctness of deniable encryption.

Incoercibility. Consider the simulator depicted on fig. 5, which essentially generates two transcripts of deniable encryption, each encrypting plaintext $m = 0$, and then uses faking algorithm of deniable encryption to simulate the coins. Note that the simulator generates the simulated coins in the same way (by using faking algorithm), no matter whether the party is corrupted or coerced.

We need to show that for every pattern of corruptions and coercions, and every set of real and fake inputs and outputs, the real execution is indistingusihable from a simulated one. This boils down to showing indistinguishability in the following cases:

1. If claimed inputs and outputs are consistent, we should prove indistinguishability between the case where both the sender and the receiver show their true coins, the case where both the sender and the receiver show their fake coins, the case where the sender shows true coins and the receiver shows fake coins, and the case where the sender shows fake coins and the receiver shows true coins.

2. If claimed inputs and outputs are inconsistent, we should prove indistinguishability between the case where the sender shows true coins and the receiver shows fake coins, the case where the sender shows fake coins and the receiver shows true coins, and the case where they both show fake coins.

The proof is very straighforward and uses two main steps - (a) switching between normally and obliviously generated execution of DE, using obliviousness and public receiver deniability of DE, and (b) switching randomness of DE of the sender between real and fake, using sender-deniability of DE.

Below we formally prove indistinguishability between the simulated execution ($\mathsf{Hyb}_{\mathsf{Sim}}$) and the real execution with consistent inputs inputs x'_0, x'_1, b' and output $x'_{b'}$, where both parties tell the truth (i.e. disclose their true coins) ($\mathsf{Hyb}_{\mathsf{Real}}$). Indistinguishability between other distributions can be shown in a very similar manner.

- $\mathsf{Hyb}_{\mathsf{Sim}}$. This is the execution from fig. 5, where both the sender and the receiver are either corrupted or coerced, and the values reported to the simulator are the following: inputs x'_0, x'_1 of the sender, input b' of the receiver, output $x' = x'_{b'}$ of the receiver. The simulator gives the adversary $(\mathsf{DE}_0, \mathsf{DE}_1, s'_0, s'_1, r'_{b'}, \tilde{r}')$.

- Hyb_1. In this hybrid the receiver generates messages in $\mathsf{DE}_{1-b'}$ obliviously (instead of generating them honestly, using $r_{1-b'}$). Indistinguishability between this and the previous hybrid follows from obliviousness of the receiver of deniable encryption. Note that it is important for the reduction that the receiver deniability is public, since the reduction needs to compute fake randomness of execution $1 - b'$, $r'_{1-b'}$, for which it doesn't know the true coins $r_{1-b'}$.

- Hyb_2. In this hybrid the sender encrypts x'_0 (instead of 0) in the execution $i = 0$. It also gives the adversary its true randomness s_0 instead of fake s'_0. Indistinguishability follows from bideniability of the encryption scheme DE_0.

Inputs given to simulate the communication: CRS.

The simulator chooses random s_0, s_1, r_0, r_1, and computes $\mathsf{DE}_i \leftarrow \mathsf{DE}(s_i, r_i, 0)$ for both $i = 0, 1$, i.e. sets DE_i to be the transcript of the protocol for deniable encryption, computed with the sender input 0, sender randomness s_i, and receiver randomness r_i. $(\mathsf{DE}_0, \mathsf{DE}_1)$ is a simulated transcript of the protocol.

Simulation of corruption and coercion of the sender S

Inputs additionally given to simulate the coercion of S: claimed inputs x'_0, x'_1 of S.

The simulator computes $s'_i \leftarrow \mathsf{DE.SFake}(s_i, 0, x'_i, \mathsf{DE}_i; \cdot)$ for both $i = 0, 1$. It gives s'_0, s'_1 to the adversary.

Simulation of corruption and coercion of the receiver R

Inputs additionally given to simulate the coercion of R: claimed input b', claimed output x'.

The simulator claims that messages of the receiver in execution $1 - b'$ were generated at random, and sets fake $\tilde{r}'$ to be the concatenation of these receiver messages. Next, it computes $r'_{b'} \leftarrow \mathsf{DE.RFake}(x', \mathsf{DE}_{b'}; \cdot)$. It gives $r'_{b'}, \tilde{r}'$ to the adversary.

Figure 5: Simulation

- $\mathsf{Hyb}_{\mathsf{Real}}$. In this hybrid the sender encrypts x'_1 (instead of 0) in the execution $i = 1$. It also gives its true randomness s_1 instead of fake s'_1. Indistinguishability follows from sender deniability of the encryption scheme DE_1.

  Note that this distribution corresponds to the real world where parties use x'_0, x'_1, b' as inputs.

4   4-Round Incoercible MPC

4.1   Description of the protocol

In this section we describe our protocol achieving incoercibility even when all parties are coerced, but only in cooperative scenario. That is, as discussed in the introduction, the deception remain undetectable only as long as *all* parties lie about their randomness (however, then can still tell the truth about their inputs, if they choose so). We remind that in this work we only focus on coercions and corruptions which happen after the protocol execution.

Our protocol is presented on fig. 6. As discussed more in detail in the introduction, the protocol essentially instructs parties to run the underlying adaptively secure protocol, where each message is encrypted under a separate instance of deniable encryption. In addition, party P_1 sends to everyone the same seed seed of the prg, to be used in the faking procedure. Parties' faking algorithm instructs parties to use seed to derive (the same for all parties) coins r_{Sim}, which are used to generate (the same for all parties) simulated transcript σ' of the underlying MPC. Next each party uses the local simulator of that MPC (recall that we need that MPC to have corruption-oblivious simulator) to simulate its own fake coins of the underlying MPC. Finally parties claim that they indeed exchanged messages of σ', using deniability of encryption.

Faking the inputs vs faking the inputs and the outputs.   We note that it is enough for parties to be able to fake their inputs (as opposed to inputs and outputs), due to the standard transformation allowing parties to

mask their output with a one time pad k: $f'((x_1, k_1), (x_2, k_2)) = f(x_1, x_2) \oplus k_1 || f(x_1, x_2) \oplus k_2$. Indeed, here faking the output can be achieved by faking inputs k_i instead. Thus, in the protocol, we only describe an input-faking mechanism.

Theorem 6 *Assume the existence of the following primitives:*

- aMPC = (aMPC.msg1, aMPC.msg2, aMPC.Eval, aMPC.Sim, aMPC.Sim$_i$) *is a 2-round adaptively secure MPC with corruption-oblivious simulation, in a global CRS model;*

- DE = (DE.msg1, DE.msg2, DE.msg3, DE.Dec, DE.SFake, DE.RFake) *is a 3-message, delayed-input deniable encryption protocol, in a CRS model;*

- prg *is a pseudorandom generator.*

Then the protocol iMPC *on fig. 6, 7 is a 4-round semi-honest MPC protocol in a CRS model*[13], *which is cooperatively incoercible.*

We note that all required primitives can be built using subexponentially-secure indistinguishability obfuscation and one-way functions ([CPP20, CPV17]). Therefore we obtain the following corollary:

Corollary 1 *Assume the existence of subexponentially secure indistinguishability obfuscation and subexponentially secure one-way functions. Then in a CRS model there exists a 4-round semi-honest MPC, which is cooperatively incoercible.*

Notation and indexing.   Subscript i, j on the message of the protocol means that the message is sent from P_i to P_j. Subscript i, j of the randomness means that this randomness is used as sender or receiver randomness in the protocol where i is the sender and j is the receiver.

For example, $M1_{i,j}$ is the first message of aMPC, sent from P_i to P_j. Our protocol transmits this message inside deniable encryption, which in turn consists of messages $a1_{i,j}$, $a2_{j,i}$, and $a3_{i,j}$. To compute these messages, party P_i uses its sender randomness $s_{i,j,1}$, and party P_j uses its receiver randomness $r_{i,j,1}$.

4.2   Proof of the Theorem

Correctness.   Correctness of the protocol immediately follows from correctness of the underlying aMPC protocol and correctness of deniable encryption DE.

Incoercibility.   We define a simulator which can simulate communication and internal states of all parties, given inputs and outputs only, but without knowing whether these inputs are real or fake.

We can assume that the simulator knows the output y before the protocol starts, due to the following standard transformation, where parties additionally choose OTP keys k_i and use it to mask the output: $f'((x_1, k_1), x_2, k_2) = f(x_1, x_2) \oplus k_1 || f(x_1, x_2) \oplus k_2$. Due to this transformation, the simulator can always choose output z of parties uniformly at random, and once the first coercion occurs and the true output y

---

[13]Note that our CRS is global (recall that the notion of deniability or incoercibility only makes sense in the global CRS model).

<div style="border:1px solid black; padding:10px;">

4-round incoercible MPC protocol iMPC**:**

The CRS: $\mathsf{CRS} = (\mathsf{CRS}_{\mathsf{DE}}, \mathsf{CRS}_{\mathsf{aMPC}})$, where $\mathsf{CRS}_{\mathsf{DE}}$ is a CRS of deniable encryption, and $\mathsf{CRS}_{\mathsf{aMPC}}$ is a CRS of adaptively secure MPC protocol.

Inputs: inputs $x_1, \ldots, x_n$ of parties $P_1, \ldots, P_n$, respectively;

Randomness: each party P_i generates the following random values:

1. $s_{i,j,1}, r_{i,j,1}, j \neq i$, which is sender and receiver randomness of DE, used to send and receive aMPC messages of round 1;
2. $s_{i,j,2}, r_{i,j,2}, j \neq i$, which is sender and receiver randomness of DE, used to send and receive aMPC messages of round 2;
3. $s_{\mathsf{aMPC},i}$, which is randomness of party P_i in the underlying aMPC protocol.

In addition, party P_1 chooses at random:

1. seed, which will be used by parties to generate coins of the simulator r_{Sim} and fake seed$'$;
2. $s_{1,j,3}, j \neq 1$, which is sender randomness of DE used to send seed;

Finally, parties $P_i, i \neq 1$ generate $r_{1,i,3}$, which is receiver randomness of DE, used to receive seed.

We denote all randomness generated by each party P_i by s_i.

The protocol:

1. **Round 1:** Each party P_i sends to each other party P_j, $j \neq i$, the following:
   $$a1_{i,j} = \mathsf{DE.msg1}(\mathsf{CRS}_{\mathsf{DE}}; s_{i,j,1}).$$

2. **Round 2:** Each party P_i sends to each other party P_j, $j \neq i$, the following:
   - $a2_{i,j} = \mathsf{DE.msg2}(\mathsf{CRS}_{\mathsf{DE}}; r_{i,j,1}, a1_{j,i}).$
   - $b1_{i,j} = \mathsf{DE.msg1}(\mathsf{CRS}_{\mathsf{DE}}; s_{i,j,2}).$

   In addition, P_1 sends to each other party P_j, $j \neq 1$, the following:
   - $c1_{1,j} = \mathsf{DE.msg1}(\mathsf{CRS}_{\mathsf{DE}}; s_{1,j,3}).$

3. **Round 3:** Each party P_i for each $j \neq i$ computes $\{M1_{i,1}, \ldots, M1_{i,n}\} \leftarrow \mathsf{aMPC.msg1}(\mathsf{CRS}_{\mathsf{aMPC}}; x_i; s_{\mathsf{aMPC},i})$, and sends the following:
   - $a3_{i,j} = \mathsf{DE.msg3}(\mathsf{CRS}_{\mathsf{DE}}; s_{i,j,1}, M1_{i,j}, a1_{i,j}, a2_{j,i}).$
   - $b2_{i,j} = \mathsf{DE.msg2}(\mathsf{CRS}_{\mathsf{DE}}; r_{i,j,2}, b1_{j,i}).$

   In addition, each party P_i except P_1 sends to P_1 the following:
   - $c2_{i,1} = \mathsf{DE.msg2}(\mathsf{CRS}_{\mathsf{DE}}; r_{1,i,3}, c1_{1,i}).$

4. **Round 4:** Each party P_i, for each $j \neq i$, computes $M1_{j,i} \leftarrow \mathsf{DE.Dec}(\mathsf{CRS}_{\mathsf{DE}}; r_{j,i,1}, a1_{j,i}, a2_{i,j}, a3_{j,i})$. Next for each $j \neq i$ it computes $\{M2_{i,1}, \ldots, M2_{i,n}\} \leftarrow \mathsf{aMPC.msg2}(\mathsf{CRS}_{\mathsf{aMPC}}; x_i, M1_{1,i}, \ldots, M1_{n,i}; s_{\mathsf{aMPC},i})$, and sends the following:
   - $b3_{i,j} = \mathsf{DE.msg3}(\mathsf{CRS}_{\mathsf{DE}}; s_{i,j,2}, M2_{i,j}, b1_{i,j}, b2_{j,i}).$

   In addition, P_1 sends to each other party P_j, $j \neq 1$, the following:
   - $c3_{1,j} = \mathsf{DE.msg3}(\mathsf{CRS}_{\mathsf{DE}}; s_{1,j,3}, \mathsf{seed}, c1_{1,j}, c2_{j,1}).$

5. **Evaluation:** Each party P_i, for each $j \neq i$, computes $M2_{j,i} \leftarrow \mathsf{DE.Dec}(\mathsf{CRS}_{\mathsf{DE}}; r_{j,i,2}, b1_{j,i}, b2_{i,j}, b3_{j,i})$. Next for each $j \neq i$ it computes $y \leftarrow \mathsf{aMPC.Eval}(\mathsf{CRS}_{\mathsf{aMPC}}; x_i, M1_{1,i}, \ldots, M1_{n,i}, M2_{1,i}, \ldots, M2_{n,i}; s_{\mathsf{aMPC},i})$. It sets y to be its output in the protocol.

By $\pi = \mathsf{iMPC}(\mathsf{CRS}, (x_1, s_1), \ldots, (x_n, s_n)) = (\{a1_{i,j}, a2_{i,j}, a3_{i,j}\}_{i \neq j}, \{b1_{i,j}, b2_{i,j}, b3_{i,j}\}_{i \neq j}, \{c1_{1,j}, c2_{j,1}, c3_{1,j}\}_{j \neq 1})$ we denote the transcript of our protocol.

By $\sigma = \mathsf{aMPC}(\mathsf{CRS}_{\mathsf{aMPC}}, (x_1, s_{\mathsf{aMPC},1}), \ldots, (x_n, s_{\mathsf{aMPC},n})) = (\{M1_{i,j}, M2_{i,j}\}_{i \neq j})$ we denote the transcript of underlying adaptive MPC protocol aMPC.

</div>

Figure 6: 4-round incoercible MPC protocol.

Faking procedure of party P_i, $i = 1, \ldots, n$

Inputs: P_i's true input x_i, fake input x_i', true output y, real random coins s_i, and the protocol transcript π.

1. **learning the seed:** P_1 knows the seed seed (which it generated). For $i \neq 1$, P_i computes
   seed $\leftarrow$ DE.Dec($\mathsf{CRS_{DE}}; r_{1,i,3}, c1_{1,i}, c2_{i,1}, c3_{1,i}$).

2. **expanding the seed:** P_i computes prg(seed) and parses the result as $r_{\mathsf{Sim}}\|$seed$'$, where $|$seed$| = |$seed$'|$.

3. **computing fake transcript:** P_i computes the fake transcript and state $(\sigma', \mathsf{state}) \leftarrow$ aMPC.Sim($\mathsf{CRS_{aMPC}}, y, r_{\mathsf{Sim}}$) of the underlying 2-round MPC protocol. Let $\sigma' = \left(\left\{ M1'_{i,j}, M2'_{i,j} \right\}_{i \neq j} \right)$.

4. **computing fake coins of the underlying MPC:** P_i computes the fake coins $s'_{\mathsf{aMPC},i} \leftarrow$ aMPC.Sim$_i$($\mathsf{CRS_{aMPC}}, \mathsf{state}, x_i', y$) of the underlying MPC protocol, using the local simulator.

5. **computing fake coins of deniable encryption:** P_i computes the fake coins for each instance of deniable encryption as follows:
   $s'_{i,j,1} \leftarrow$ DE.SFake($\mathsf{CRS_{DE}}, s_{i,j,1}, M1_{i,j}, M1'_{i,j}, a1_{i,j}, a2_{j,i}, a3_{i,j}; \cdot$), to claim that it sent $M1'_{i,j}$ instead of $M1_{i,j}$;
   $s'_{i,j,2} \leftarrow$ DE.SFake($\mathsf{CRS_{DE}}, s_{i,j,2}, M2_{i,j}, M2'_{i,j}, b1_{i,j}, b2_{j,i}, b3_{i,j}; \cdot$), to claim that it sent $M2'_{i,j}$ instead of $M2_{i,j}$;
   $r'_{i,j,1} \leftarrow$ DE.RFake($\mathsf{CRS_{DE}}, r_{i,j,1}, M1_{j,i}, M1'_{j,i}, a1_{j,i}, a2_{i,j}, a3_{j,i}; \cdot$), to claim that it received $M1'_{j,i}$ instead of $M1_{j,i}$;
   $r'_{i,j,2} \leftarrow$ DE.RFake($\mathsf{CRS_{DE}}, r_{i,j,2}, M2_{j,i}, M2'_{j,i}, b1_{j,i}, b2_{i,j}, b3_{j,i}; \cdot$), to claim that it received $M2'_{j,i}$ instead of $M2_{j,i}$.
   Further, if $i = 1$, then for each $j \neq 1$ the party computes:
   $s'_{1,j,3} \leftarrow$ DE.SFake($\mathsf{CRS_{DE}}, s_{1,j,3}, \mathsf{seed}, \mathsf{seed}', c1_{1,j}, c2_{j,1}, c3_{1,j}; \cdot$), to claim that it sent seed$'$ instead of seed.
   If $i \neq 1$, then P_i computes
   $r'_{1,i,3} \leftarrow$ DE.RFake($\mathsf{CRS_{DE}}, r_{1,i,3}, \mathsf{seed}, \mathsf{seed}', c1_{1,i}, c2_{i,1}, c3_{1,i}; \cdot$), to claim that it received seed$'$ instead of seed.

The output of the faking procedure: Finally, P_i gives the adversary its fake internal state s_i', where:

- If $i \neq 1$, $s_i' = \left\{ s'_{i,j,1} \right\}_{j \neq i}, \left\{ r'_{i,j,1} \right\}_{j \neq i}, \left\{ s'_{i,j,2} \right\}_{j \neq i}, \left\{ r'_{i,j,2} \right\}_{j \neq i}, \left\{ r'_{1,i,3} \right\}, s'_{\mathsf{aMPC},i}$.

- If $i = 1$, $s_i' = \left\{ s'_{i,j,1} \right\}_{j \neq i}, \left\{ r'_{i,j,1} \right\}_{j \neq i}, \left\{ s'_{i,j,2} \right\}_{j \neq i}, \left\{ r'_{i,j,2} \right\}_{j \neq i}, \left\{ s'_{1,j,3} \right\}_{j \neq 1}, s'_{\mathsf{aMPC},i}, \mathsf{seed}'$.

(Note that all other information which P_i should know in the honest execution, e.g. seed$'$ or $M1'_{i,j}$, can be derived by the adversary using random coins s_i', input x_i', the transcript π, and the CRS.)

Figure 7: Faking procedure of party P_i, $i = 1, \ldots, n$

becomes known, set the corresponding k_i to be $z \oplus (y||y)$. From now on we assume that the simulator knows the output y ahead of time.

Simulation. The simulator is formally described on fig. 8. Informally, the simulator uses the underlying simulator of aMPC to simulate communication between parties, σ'. It then encrypts messages of σ' under deniable encryption. It encrypts randomly chosen seed' under deniable encryption as well. This concludes the description of simulation of communication.

Upon coercion of a party, given an input x'_i (without knowing whether x_i is real or fake), the simulator computes fake random coins of aMPC by running the local simulator $\mathsf{aMPC.Sim}_i$ on input x'_i. These are the only coins which are faked by the simulator; the simulator reveals true values of seed' and all randomness of DE.

Let $x_1, \ldots, x_n$ and $x'_1, \ldots, x'_n$ be some inputs to the protocol, and let y be some output. Consider the following distributions:

- $\mathsf{Hyb}_{\mathsf{Real}}$: this is the distribution corresponding to the real execution of the protocol with inputs $x'_1, \ldots, x'_n$, where parties disclose their *true* inputs and randomness.

- $\mathsf{Hyb}_{\mathsf{Fake}}$: this is the distribution corresponding to the real execution of the protocol with inputs

$x_1, \ldots, x_n$, where parties disclose *fake* inputs $x'_1, \ldots, x'_n$, output y, and fake randomness.

- $\mathsf{Hyb}_{\mathsf{Sim}}$: this is the distribution corresponding to the simulation from figure 8, where the simulator is given output y and claimed inputs $x'_1, \ldots, x'_n$.

We need to show the following:

1. If $x'_1, \ldots, x'_n$ and y are consistent (i.e. $f(x'_1, \ldots, x'_n) = y$), then we need to show that $\mathsf{Hyb}_{\mathsf{Sim}} \approx \mathsf{Hyb}_{\mathsf{Real}}$ and $\mathsf{Hyb}_{\mathsf{Sim}} \approx \mathsf{Hyb}_{\mathsf{Fake}}$.

2. If $x'_1, \ldots, x'_n$ and y are not consistent, then we need to show that $\mathsf{Hyb}_{\mathsf{Sim}} \approx \mathsf{Hyb}_{\mathsf{Fake}}$.

We show this below. First, we show indistinguishability between $\mathsf{Hyb}_{\mathsf{Sim}} \approx \mathsf{Hyb}_{\mathsf{Fake}}$, for any values $x_1, \ldots, x_n, x'_1, \ldots, x'_n$, and y:

- $\mathsf{Hyb}_{\mathsf{Fake}}$. We start with the distribution corresponding to the real-world execution of the protocol, where parties fake their random coins upon coercion. In other words, the adversary sees CRS, π, and x'_i, s'_i for each i, generated as in fig. 6, 7. In particular, the truly sent transcript σ of the underlying MPC is a transcript on inputs x_i; however, parties claim that they instead sent (simulated) transcript σ', which appears consistent with fake inputs x'_i.

- Hyb_1. In this hybrid P_1 sends seed$'$ instead of seed inside $\{c1_{1,j}, c2_{j,1}, c3_{1,j}\}_{j \neq 1}$, and parties (both senders and receivers) give the adversary true randomness for this deniable encryption (instead of faking it to seed$'$). Indistinguishability between this and the previous distribution holds by $n - 1$ invocations of bideniability of encryption for plaintexts seed and seed$'$.

- Hyb_2. In this hybrid we switch $r_{\mathsf{Sim}} \| \mathsf{seed}'$ from $\mathsf{prg}(\mathsf{seed})$ to uniformly random. Indistinguishability holds by security of a prg. Note that seed is not used anywhere else in the distribution, thus the reduction is possible.

- $\mathsf{Hyb}_{\mathsf{Sim}}$. In this hybrid we set $\{a1_{i,j}, a2_{i,j}, a3_{i,j}\}_{i \neq j}$ to encrypt 1-round messages of simulated σ' (consistent with fake x'_i), instead of encrypting 1-round messages of real transcript σ (consistent with x_i). Also, all parties give true randomness $\{s_{i,j,1}\}_{j \neq i}, \{r_{i,j,1}\}_{j \neq i}$, instead of giving fake randomness consistent with σ'.

  Similarly, we change $\{b1_{i,j}, b2_{i,j}, b3_{i,j}\}_{i \neq j}$ to encrypt 1-round messages of simulated σ' (consistent with fake x'_i), instead of encrypting 1-round messages of real transcript σ (consistent with x_i). Also, all parties give true randomness $\{s_{i,j,2}\}_{j \neq i}, \{r_{i,j,2}\}_{j \neq i}$, instead of giving fake randomness consistent with σ'.

  Indistinguishability between this and the previous distribution holds by $2n(n - 1)$ invocations of bideniability of encryption, where plaintexts are messages of σ and σ'.

  Note that this is the simulated distribution.

Further, for the case when $f(x'_1, \ldots, x'_n) = y$, in one last step we show that $\mathsf{Hyb}_{\mathsf{Sim}} \approx \mathsf{Hyb}_{\mathsf{Real}}$:

- $\mathsf{Hyb}_{\mathsf{Real}}$. Compared to $\mathsf{Hyb}_{\mathsf{Sim}}$, we switch the messages of aMPC, encrypted inside deniable encryption, from simulated σ' to real σ, which is the true transcript of aMPC on inputs x'_i. In addition, parties reveal their true randomness $s_{\mathsf{aMPC},i}$ instead of computing simulated $s'_{\mathsf{aMPC},i}$ consistent with x'_i using the local simulator $\mathsf{aMPC.Sim}_i$.

Indistinguishability between this and the simulation follows from adaptive security of aMPC. Note that indeed r_{Sim}, randomness of the simulator, is not used anywhere else in the distribution.

This distribution corresponds to the real execution of the protocol on inputs x_i', where parties disclose their true randomness upon being coerced.

This concludes the security proof.

5 Incoercible MPC with Lazy Parties is Impossible

In this section we describe our impossibility result for incoercible MPC protocols with a certain communication pattern. We consider the synchronous model of communication, where parties send their messages in rounds. We call a party *lazy*, if it sends its messages only in the first and in the last round of the protocol, but not in any other round[14]. We show that a protocol for 3 or more parties cannot be incoercible, as long as there is at least one lazy party Z, and there is another party (different from Z) which receives the output.

In particular, this impossibility rules out protocols with the following communication structure, which is a natural extention of a "ping-pong" communication of 3-message 2PC to a multiparty setting: assume just one party receives the output; we call this party the receiver, and call all other parties the senders. Then the communication proceeds as follows:

- In round 1 the senders send out their messages to everybody;

- In round 2 the receiver sends its messages to the senders;

- In round 3 the senders send out their messages to everybody[15].

Our impossibility is based on the fact that in an incoercible protocol with lazy party Z it is possible to do a variation of a residual function attack, similar to impossibility of standard (non-incoercible) non-interactive MPC. Concretely, we show that security against coercion of a lazy party Z implies that Z can always pick an input x' different from its real input x and generate a different last message of the protocol corresponding to new input x', such that the resulting transcript will be a valid transcript for this new input x', as if Z used x' even in the first message (despite the fact that in reality its first message was generated using x). As a result, the adversary may corrupt Z together with some output-receiving party and evaluate the function on any possible input of Z, thus compromising security of other parties.

Theorem 7 *Let $n \geq 3$, and assume there exists an n-party protocol for evaluating function $f(x_1, \ldots, x_n)$, such that P_1 is lazy and P_n receives the output. Further, assume it is secure against coercion of P_1, and against corruption of P_1 and P_n. Then the function f is such that for any inputs $x_1, \ldots, x_n$ it is possible, given x_1, x_n, and $f(x_1, \ldots, x_n)$, to compute $f(x, x_2, \ldots, x_n)$ in polynomial time for any x of the same length as x_1.*

Note that, while the theorem statement also holds for the case of 2 parties, it doesn't imply any impossibility since for any 2-input function f it is always possible to compute $f(\cdot, x_2)$ given x_1, x_2, and thus the theorem doesn't impose any restrictions on functions f which can be computed incoercibly using 2-party protocols.

---

[14]In particular, when the protocol requires only 2 rounds, each party is lazy by definition.

[15]Note that in standard, non-deniable MPC the last message doesn't need to be sent to parties who don't receive the output. However, in deniable MPC parties who don't get the output may still need the last message in order to fake.

Figure 9: Algorithm NewMessage to generate the last message consistent with a different x_1'.

Proof of theorem 7. Without loss of generality we assume that the lazy party is P_1, and party which receives the output is P_n. Further, we assume that P_1 is the first to send its messages in round 1, and the last to send its messages in round N.

Let us denote the randomness of P_1 by r_1, the concatenated randomness of all other parties by $R = r_2 || \ldots || r_n$, the input of P_1 by x_1, the concatenated input of all other parties by $X = x_2 || \ldots || x_n$. In addition, let X^0 denote some fixed set of inputs such that $|X| = |X^0|$, e.g. all-zero inputs $0^{|X|}$. Let NMF_i denote the next message function of the protocol for party 1 in round i. Let $\mathsf{Eval}(x_n; transcript; r_n)$ denote the output evaluation function of party P_n which takes as input its randomness r_n, input x_n, and all communication in the protocol. Let $\alpha = \mathsf{NMF}_1(x_1; r_1)$ denote the concatenated messages sent by P_1 to all other parties in round 1, $\mathsf{com} = \mathsf{com}(\alpha; X; R)$ denote the concatenated messages sent by parties $P_2, \ldots, P_n$ in all rounds, com_{N-1} denote com except for messages of the last round, and $\beta = \mathsf{NMF}_N(x_1, \mathsf{com}_{N-1}; r_1)$ denote the concatenated messages sent by P_1 to all other parties in round N. Finally, let $\mathsf{Fake}_1(r_1, x_1, x_1', \mathsf{com}; \rho)$ denote the faking algorithm of party P_1, which takes as input its true coins and input r_1, x_1, desired fake input x_1', and com, all communication sent to P_1. Fake_1 could be deterministic or randomized; without loss of generality we assume that it is randomized using its own random coins ρ.

Consider the following algorithm NewMessage (fig. 9) which for any x_1' allows P_1 to generate a different β' such that $(\alpha, \mathsf{com}, \beta')$ is a valid transcript resulting in the output $f(x_1', X)$. The intuition behind this procedure is as follows: First, P_1 computes a transcript which starts with the same α but continues with a different $\widetilde{\mathsf{com}}$ (computed under freshly chosen randomness of other parties and fixed inputs X^0). Next, it runs its faking algorithm to generate fake coins r_1' which make this transcript look consistent with x_1' (in particular, this makes r_1', x_1' look like valid coins and input for α, even though α was generated under x_1). Finally, it uses fake r_1' to generate its last message β' using the original communication com and new input x_1'. In the following lemma 1 we claim that β', together with the original communication (α, com), forms a valid transcript for inputs x_1, X which will be evaluated correctly by the output-receiving party:

Lemma 1 *Let $\alpha, \mathsf{com}, \beta'$ be generated as described above, and let the protocol be secure against the coercion of P_1. Then for any f, x_1, X, x_1', with overwhelming probability over the choice of r_1, R, ρ it holds that $\mathsf{Eval}(x_n; \alpha, \mathsf{com}, \beta'; r_n) = f(x_1', X)$.*

Proof of lemma 1. This statement follows from the correctness of the protocol and its security against a coercion of P_1. Indeed, consider the adversary who tries to distinguish whether it sees the transcript corresponding to inputs x_1, X^0, and fake randomness of P_1 for input x_1', or the transcript of the protocol on inputs x_1', X^0 with the true randomness of P_1. More formally:

- In the first case, the adversary receives $(\alpha, \widetilde{\mathsf{com}}, \widetilde{\beta})$ and r'_1, where r_1, ρ_1, ρ_2 are uniformly chosen, $\alpha \leftarrow \mathsf{NMF}_1(x_1; r_1)$, $\widetilde{\mathsf{com}} \leftarrow \mathsf{com}(\alpha, X^0; \rho_1)$, $\widetilde{\beta} = \mathsf{NMF}_N(x_1, \widetilde{\mathsf{com}}_{N-1}; r_1)$, and $r'_1 \leftarrow \mathsf{Fake}_1(r_1, x_1, x'_1, \widetilde{\mathsf{com}}; \rho_2)$. In other words, $(\alpha, \widetilde{\mathsf{com}}, \widetilde{\beta})$ is a transcript of the protocol on inputs x_1, X^0 with randomness r_1, ρ_1, and r'_1 is fake randomness for input x'_1.

- In the second case, the adversary receives $(\alpha, \widetilde{\mathsf{com}}, \widetilde{\beta})$ and r'_1, where r'_1, ρ_1, ρ_2 are uniformly chosen, $\alpha \leftarrow \mathsf{NMF}_1(x'_1; r'_1)$, $\widetilde{\mathsf{com}} \leftarrow \mathsf{com}(\alpha, X^0; \rho_1)$, $\widetilde{\beta} = \mathsf{NMF}_N(x'_1, \widetilde{\mathsf{com}}_{N-1}; r'_1)$. In other words, $(\alpha, \widetilde{\mathsf{com}}, \widetilde{\beta})$ is a transcript of the protocol on inputs x'_1, X^0 with randomness r'_1, ρ_1.

Given one of the two distributions, the adversary proceeds as follows. It chooses randomness $R = r_2 || \ldots || r_n$ (which is concatenated randomness for parties $P_2, \ldots, P_n$) and computes $\mathsf{com} = \mathsf{com}(\alpha; X; R)$. Next it computes $\beta' = \mathsf{NMF}_N(x'_1, \mathsf{com}_{N-1}; r'_1)$. It then runs $\mathsf{Eval}(x_n; \alpha, \mathsf{com}, \beta'; r_n)$.

Note that in the second case $(\alpha, \mathsf{com}, \beta')$ is a transcript of the protocol on inputs x'_1, X with uniformly chosen random coins r'_1, R, and therefore, by correctness, it holds that with overwhelming probability $\mathsf{Eval}(x_n; \alpha, \mathsf{com}, \beta'; r_n) = f(x'_1, X)$. By indistinguishability of the first and the second case, it should also be true that in the first case $\mathsf{Eval}(x_n; \alpha, \mathsf{com}, \beta'; r_n) = f(x'_1, X)$. Finally, note that in the first case the resulting β' is obtained exactly as in algorithm $\mathsf{NewMessage}$, which concludes the proof of the lemma.

Now we finish the proof of the theorem 7. We claim that the adversary who corrupts P_1 and P_n in the real world can compute $f(x, x_2, \ldots, x_n)$ for any input x (of the same length as x_1), where $x_1, \ldots, x_n$ are inputs of the parties in the protocol. Indeed, the adversary can do so in two steps: first it corrupts P_1 to learn r_1 and x_1 and runs $\beta' \leftarrow \mathsf{NewMessage}(x_1, r_1, x, \alpha, \mathsf{com}; \rho)$ for any desired input x and random ρ (as before, α, com is the communication of P_1 in round 1 and of all other parties). Next it corrupts P_n to learn r_n and computes $\mathsf{Eval}(x_n; \alpha, \mathsf{com}, \beta'; r_n)$, which is with overwhelming probability equal to $f(x, x_2, \ldots, x_n)$, as shown in the lemma 1. Note that in the ideal world the adversary who only corrupts P_1 and P_n and learns x_1, x_n, and $f(x_1, \ldots, x_n)$ cannot compute residual function $f(\cdot, x_2, \ldots, x_n)$ (except for some functions f), and therefore the adversary in the real world has an advantage. This finishes the proof of the theorem 7.

References

[AOZZ15] Joël Alwen, Rafail Ostrovsky, Hong-Sheng Zhou, and Vassilis Zikas. Incoercible multi-party computation and universally composable receipt-free voting. In *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part II*, pages 763–780, 2015. 2, 11, 12, 13

[BCH12] Nir Bitansky, Ran Canetti, and Shai Halevi. Leakage-tolerant interactive protocols. In *Theory of Cryptography - 9th Theory of Cryptography Conference, TCC 2012, Taormina, Sicily, Italy, March 19-21, 2012. Proceedings*, pages 266–284, 2012. 7, 17

[BH92] Donald Beaver and Stuart Haber. Cryptographic protocols provably secure against dynamic adversaries. In *Advances in Cryptology - EUROCRYPT '92, Workshop on the Theory and Application of of Cryptographic Techniques, Balatonfüred, Hungary, May 24-28, 1992, Proceedings*, pages 307–323, 1992. 5

[BMN+09] Josh Benaloh, Tal Moran, Lee Naish, Kim Ramchen, and Vanessa Teague. Shuffle-sum: coercion-resistant verifiable tallying for STV voting. *IEEE Trans. Inf. Forensics Secur.*, 4(4):685–698, 2009. 11

[BNNO11] Rikke Bendlin, Jesper Buus Nielsen, Peter Sebastian Nordholt, and Claudio Orlandi. Lower and upper bounds for deniable public-key encryption. In *Advances in Cryptology - ASIACRYPT 2011 - 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings*, pages 125–142, 2011. 10, 12

[BT94] Josh Cohen Benaloh and Dwight Tuinstra. Receipt-free secret-ballot elections (extended abstract). In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, 23-25 May 1994, Montréal, Québec, Canada*, pages 544–553, 1994. 11

[CDMW09] Seung Geol Choi, Dana Dachman-Soled, Tal Malkin, and Hoeteck Wee. Improved non-committing encryption with applications to adaptively secure protocols. In *Advances in Cryptology - ASIACRYPT 2009, 15th International Conference on the Theory and Application of Cryptology and Information Security, Tokyo, Japan, December 6-10, 2009. Proceedings*, pages 287–302, 2009. 5

[CDNO96] Ran Canetti, Cynthia Dwork, Moni Naor, and Rafail Ostrovsky. Deniable encryption. *IACR Cryptology ePrint Archive*, 1996:2, 1996. 2

[CFGN96] Ran Canetti, Uriel Feige, Oded Goldreich, and Moni Naor. Adaptively secure multi-party computation. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 639–648, 1996. 5, 17

[CG96] Ran Canetti and Rosario Gennaro. Incoercible multiparty computation (extended abstract). In *37th Annual Symposium on Foundations of Computer Science, FOCS '96, Burlington, Vermont, USA, 14-16 October, 1996*, pages 504–513, 1996. 1, 2, 3, 4, 8, 12, 13

[CGP15] Ran Canetti, Shafi Goldwasser, and Oxana Poburinnaya. Adaptively secure two-party computation from indistinguishability obfuscation. In *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part II*, pages 557–585, 2015. 2, 3, 8, 12, 13

[CLOS02] Ran Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. Universally composable two-party and multi-party secure computation. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 494–503, 2002. 6

[CPP20] Ran Canetti, Sunoo Park, and Oxana Poburinnaya. Fully deniable interactive encryption. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part I*, volume 12170 of *Lecture Notes in Computer Science*, pages 807–835. Springer, 2020. 3, 5, 6, 8, 9, 10, 12, 14, 17, 18, 21, 30

[CPV17] Ran Canetti, Oxana Poburinnaya, and Muthuramakrishnan Venkitasubramaniam. Better two-round adaptive multi-party computation. In *Public-Key Cryptography - PKC 2017 - 20th IACR International Conference on Practice and Theory in Public-Key Cryptography, Amsterdam, The Netherlands, March 28-31, 2017, Proceedings, Part II*, pages 396–427, 2017. 8, 17, 21

[DKR14]   Dana Dachman-Soled, Jonathan Katz, and Vanishree Rao. Adaptively secure, universally composable, multi-party computation in constant rounds. *IACR Cryptology ePrint Archive*, 2014:858, 2014. 2, 9, 12

[DN00]   Ivan Damgård and Jesper Buus Nielsen. Improved non-committing encryption schemes based on a general complexity assumption. In *Advances in Cryptology - CRYPTO 2000, 20th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 2000, Proceedings*, pages 432–450, 2000. 5

[GMW87]   Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, New York, New York, USA*, pages 218–229, 1987. 6, 7

[HOR15]   Brett Hemenway, Rafail Ostrovsky, and Alon Rosen. Non-committing encryption from Φ-hiding. In *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part I*, pages 591–608, 2015. 5

[HORR16]   Brett Hemenway, Rafail Ostrovsky, Silas Richelson, and Alon Rosen. Adaptive security with quasi-optimal rate. In *Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part I*, pages 525–541, 2016. 5

[MN06]   Tal Moran and Moni Naor. Receipt-free universally-verifiable voting with everlasting privacy. In *Advances in Cryptology - CRYPTO 2006, 26th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 2006, Proceedings*, pages 373–392, 2006. 2, 12, 13

[SW14]   Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 475–484, 2014. 2, 12

[UM10]   Dominique Unruh and Jörn Müller-Quade. Universally composable incoercibility. In *Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings*, pages 411–428, 2010. 12

[Yao86]   A. C. Yao. How to generate and exchange secrets. In *27th Annual Symposium on Foundations of Computer Science (sfcs 1986)*, pages 162–167, 1986. 4

[YKT19]   Yusuke Yoshida, Fuyuki Kitagawa, and Keisuke Tanaka. Non-committing encryption with quasi-optimal ciphertext-rate based on the ddh problem. *IACR Cryptol. ePrint Arch.*, 2019:1151, 2019. 5

A   Receiver-oblivious DE

In this appendix we briefly explain why the construction of [CPP20] is receiver-oblivious. In their construction, the sender and the receiver exchange messages μ_1, μ_2, μ_2, where μ_1 and μ_3 are sent by the sender and only μ_2 is sent by the receiver. In their construction, parties never do any computation themselves; instead they only

choose random coins and then run special obfuscated programs $P1, P2, P3$ to compute messages μ_1, μ_2, μ_2, respectively.

In particular, the message of the receiver μ_2 is a result of applying an extracting PRF RG on randomness r of the receiver and the first message μ_1 in the protocol (again, computed using the obfuscated program); using the fact that this PRF is a strong extractor, it is possible to change the output of the PRF to a uniformly random value indistinguishably, even in a presence of a PRF key. More concretely, it can be shown that (s, μ_1, μ_2, μ_3) is indistinguishable from (s, μ_1, u, μ_3), where s is the true randomness of the sender, and u is uniformly random string of size μ_2; thus, even the sender cannot tell whether the second message of the protocol is generated by the receiver honestly or obliviously.

Let us reduce this statement to security of strong extractor. Assume someone can distinguish between the transcript with $\mu_2 = \mathsf{RG}_{k_R}(r, \mu_1)$ or random u, then we can win the strong extractor game as follows. We choose the sender randomness s at random and try to guess what the plaintext m will be (since m is simply a bit, the guess will succeed with one half probability). We compute $\mu_1 = \mathsf{P1}(s, m)$, and give μ_1 as an input to the challenger of the strong computational extractor game. The challenger chooses the key k_R for the extracting PRF RG, chooses random r and either gives us $\mu_2 = \mathsf{RG}_{k_R}(r, \mu_1)$ or uniformly random u, together with the key k_R. We can use these values to reconstruct the rest of the distribution, including the obfuscated programs (note that r is not required anywhere else). Thus, if someone breaks receiver-obliviousness, then our reduction can break the strong extractor property of the PRF.

B   Input-Delayed Deniable Encryption

In this appendix we show how to transform any deniable encryption into its input-delayed version - i.e. an encryption where only the last message depends on the plaintext.

Theorem 8 *Any deniable encryption can be transformed to a deniable encryption with the same number of rounds, where the plaintext should be determined only by the last round.*

Proof.   The idea is to transmit a one-time pad key k (instead of the plaintext m) using the deniable encryption protocol, and in the last round also send $k \oplus m$ in the clear. More formally, the sender should pick random k, such that $|k| = |m|$, before the protocol starts, and use deniable encryption to transmit k; in addition, during the last round the sender should also send $c = k \oplus m$ in the clear. The receiver should execute decryption procedure of deniable encryption to learn k, and then set its message to be $k \oplus c$. Note that in such a protocol m has to be determined only by the last round.

To fake the communication to a different plaintext m', both the sender and the receiver should compute fake $k' = c \oplus m'$, and then use the corresponding faking algorithm of deniable encryption in order to create fake random coins of the sender/receiver, making the transcript consistent with k'.