

Concurrent Secure Computation with Optimal Query Complexity

Ran Canetti^{1*}, Vipul Goyal², and Abhishek Jain^{3**}

¹ Boston University and Tel Aviv University
canetti@bu.edu

² Microsoft Research, India
vipul@microsoft.com

³ Johns Hopkins University
abhishek@cs.jhu.edu

Abstract. The multiple ideal query (MIQ) model [Goyal, Jain, and Ostrovsky, Crypto'10] offers a relaxed notion of security for concurrent secure computation, where the simulator is allowed to query the ideal functionality *multiple times per session* (as opposed to just once in the standard definition). The model provides a quantitative measure for the degradation in security under concurrent self-composition, where the degradation is measured by the number of ideal queries. However, to date, all known MIQ-secure protocols guarantee only an overall *average* bound on the number of queries per session throughout the execution, thus allowing the adversary to potentially fully compromise some sessions of its choice. Furthermore, [Goyal and Jain, Eurocrypt'13] rule out protocols where the simulator makes only an adversary-independent constant number of ideal queries per session.

We show the first MIQ-secure protocol with worst-case per-session guarantee. Specifically, we show a protocol for any functionality that matches the [GJ13] bound: The simulator makes only a *constant* number of ideal queries in *every* session. The constant depends on the adversary but is independent of the security parameter.

As an immediate corollary of our main result, we obtain the first password authenticated key exchange (PAKE) protocol for the fully concurrent, multiple password setting in the standard model with no set-up assumptions.

1 Introduction

General feasibility results for secure computation were established nearly three decades ago in the seminal works of [33,14]. However, these results only promise

* Supported by the Check Point Institute for Information Security, ISF grant 1523/14, and NSF Frontier CNS 1413920 and 1218461 grants.

** Work done in part while visiting Microsoft Research, India, and at Boston University and MIT, where the author was supported in part by NSF 1218461 and DARPA FA8750-11-2-0225. Presently supported in part by a DARPA Safeware grant.

security for a protocol if it is executed in isolation, “unplugged” from any network activity. In particular, these results are not suitable for the Internet setting where multiple protocol executions may occur *concurrently* under the control of a common adversary.

A brief history of concurrent security. Towards that end, an ambitious effort to understand and design concurrently secure protocols kicked into gear with early works such as [15,10], and later the study of the *concurrent zero knowledge* setting [11,32,7,23,30]. For other functionalities and in more general settings, however, far-reaching impossibility results were established [6,8,24,3,18,1,13]. These results refer to the “plain model” where the participating parties have no trusted set-up, and hold even if the parties have access to pairwise authenticated communication and a broadcast channel.

Two main lines of research have emerged in order to circumvent these impossibility results. The first concerns with the use of *trusted setup assumptions* such as a common random string, strong public key infrastructure or tamper-proof hardware tokens (see, e.g. [5,2,22]).

The second line of research is dedicated to the study of weaker security definitions that allow for positive results in the plain model, without additional trust assumptions. The most notable examples of this include security w.r.t. super-polynomial time simulation [28,31,4,9,12] and input-indistinguishable computation [26,12]. One main drawback in this line of research is that it is not always clear by “how much” is the definition of security relaxed, or in other words “how much security” is being lost due to concurrent attacks.

The multiple ideal query model and its applications. The multiple ideal query model (or, the MIQ model in short) of Goyal, Jain and Ostrovsky [21] takes a different approach to the problem of quantifying the security loss. In this model, the simulator is allowed to query the ideal functionality *multiple times per session* (as opposed to just once in the standard definition). On the technical side, allowing the simulator multiple queries indeed facilitates proofs of security in a concurrent setting. On the conceptual side, this model allows for a natural quantification of the “security loss” incurred by concurrent attack: the more ideal queries, the weaker the security guarantee. Furthermore, the effect of multiple ideal queries strongly depends on the task at hand, thus allowing for more fine-tuned notions of security for a given problem or setting.

One functionality where this approach proved very effective is that of password-based key exchange (namely the two-party function that outputs a secret random value to both parties if the inputs provided by the two parties are equal). When the number of queries made by the simulator per session is a constant, the security guarantees of the MIQ model actually imply fully concurrent password-based authenticated key exchange (see [16,17,21]). This fact was exploited by Goyal et. al [21] to get the first concurrent PAKE in the plain model — albeit with the significant restriction that the *same* password is to be used as input in every session. This restriction results from a weakness in their modeling and analysis - a weakness that we overcome in this work.

The central question: how many queries? So, how to best bound the number of ideal queries made by the simulator? Intuitively, if we allow a large number of queries, then the security guarantee may quickly degrade and become meaningless; in particular, if enough queries are allowed, then the adversary may be able to completely learn the inputs of the honest parties. On the other hand, if the number of allowed queries is very small (say only $1 + \epsilon$ per session) then the security guarantee is very close to that of the standard definition.

To exemplify this further, consider 1-out-of- m OT. Here, as long as λ , the simulator’s query complexity, is smaller than m , MIQ provides meaningful security which degrades gracefully with λ . More generally, the remaining security for any session i in concurrently secure computation of function f is proportional to the “level of unlearnability” of $f(\cdot, x_i)$ after q queries, where x_i is the secret input of the honest party in session i . Password-based key exchange is an extreme case of an unlearnable function. Ideally, we would like to bring λ as close as possible to 1.

Prior work: Average case vs. worst case guarantees. The best positive result in the MIQ model is due to Goyal, Gupta, and Jain [19] (improving upon [21]). They provide a construction where the number of ideal queries in a session are $(1 + \frac{\log^6 n}{n})$, where n is the security parameter. However, this is only an *average-case* guarantee over the sessions that provides very weak security. In particular, it does not preclude the ideal adversary from making an arbitrarily large number of queries in some chosen sessions (while keeping the number of queries low in the other sessions). In cases of interest, such as the PAKE functionality or the above oblivious polynomial evaluation functionality, this means that the security in some sessions may be *completely compromised!*

Furthermore, Goyal and Jain [20] recently proved an unconditional lower bound on the number of ideal queries per session. Specifically, they show that there exists a two-party functionality that cannot be securely realized in the MIQ model with any (adversary independent) constant number of ideal queries per session. A natural and important question is thus what is the best worst-case bound we can give on the number of ideal queries asked per session?

1.1 Our Results

In this work, we fully settle the question of worst-case number of per session ideal queries in the context of general function evaluation. Our main result is stated below.

Theorem 1.1 (Main result (informally stated)). *Under standard cryptographic assumptions, for every PPT functionality f , there exists a protocol in the MIQ model where the simulator makes only a constant number of ideal queries in every session. The aforementioned constant is dependent upon the adversary, and, in particular on the number of sessions (rather than being universal).*

If the number of concurrent sessions being executed by the adversary is n^c , then the constant in the above theorem will be derived from c .

We stress that due to the worst-case guarantee of our result, we are able to achieve, for the *first* time in the study of the MIQ model, meaningful security for *all sessions*, which is much closer to standard security for secure computation. Interestingly, our protocol is the same as the [19] protocol. Still, we provide a significantly better analysis of its security. We stress that prior to this work, no approach for obtaining a worst-case bound on the ideal query complexity was known.

Our upper bound tightly matches the lower bound of Goyal and Jain [20] which rule out protocols where the simulator makes a constant number of ideal queries per session for any universal constant. Taken together, this fully resolves the central problem in the study of the MIQ problem: a (adversary dependent) constant number of ideal queries per session is both necessary and sufficient for simulation. Thus, our work can be viewed as the *final step* in understanding the simulator query complexity of the MIQ model.

Fully concurrent PAKE without setup. Say that a password-based key exchange protocol is *fully concurrent* if it remains secure in a setting where unboundedly many executions of the protocol run concurrently, on potentially different passwords. An immediate corollary of our main result is the resolution of the long standing open problem of designing a fully concurrent PAKE protocol in the standard model and with no setup assumptions.

1.2 Technical Overview

Simulator Query Complexity and Precise Simulation. The question of simulator query complexity in the MIQ model is intimately connected to the notion of precise simulation introduced by Micali and Pass [25]. Recall that traditional simulator strategies allow for the simulator’s running time to be an arbitrary polynomial factor of the (worst-case) running time of the real adversary. The notion of precise simulation concerns with the study of how low this polynomial can be. This idea is, in fact, much more general and can also be used in the context of resources other than running time, such as memory, etc. Thus, in the most general sense, the goal of precise simulation is to develop simulation strategies whose resource utilization is “close” to the resource utilization of the real adversary.

As observed in [21], the study of simulator query complexity in the MIQ model can also be cast as a precise simulation problem by viewing the trusted party queries as the resource of the simulator. Therefore, advances in precise simulation strategies go hand in hand with improvements in the simulator query complexity in the MIQ model. Indeed, prior works in the MIQ model [21,19] have relied upon sophisticated precise simulation strategies in order to obtain their positive results. We note, however, that till date, all precise simulation strategies only focus on minimizing the *total cost* of the simulator across all the sessions. Indeed, this is why these works only yield an *average-case* bound on the simulator query complexity.

In this work, we are interested in minimizing the *worst-case* simulator query complexity per session. In other words, we are interested in simulation strategies that guarantee **local precision for every session**.

Our approach in a nutshell. Towards that end, our starting observation is that the problem of bounding the simulator query complexity per session can be reduced to bounding the number of times the output message of a session appears in the entire simulation transcript.⁴ In other words, we need a precise (concurrent) simulation strategy where the output message of every session appears only a constant number of times across the *entire* simulation transcript.⁵ For this purpose, we revisit existing precise simulation strategies. Concretely, we show that a slight variant of the “sparse” rewinding strategy of Goyal, Gupta and Jain [19] (that we henceforth refer to as the GGJ simulation strategy) satisfies our desired property. We prove this by a novel, purely combinatorial analysis. Our final secure computation protocol remains essentially identical to those in the prior works in the MIQ model.

We now give an overview of the steps involved in our proof. Say that we wish to analyze the number of queries in session i . Consider the specific point in the protocol execution of session i where, the simulator actually makes a query to the ideal functionality: call this point p_i (for example, this may be the 5th message of the protocol execution in session i). This means that whenever the simulator reaches the point p_i (in the overall concurrent execution), it will have to call the trusted functionality for session i to compute the next outgoing message. Thus, now the problem reduces to *simply counting* how many times the point p_i occurs in the entire rewinding schedule. Observe that in each thread of execution, point p_i only occurs once. However, there could be multiple threads of execution resulting because of rewinding. Therefore, p_i may also occur multiple times in the rewinding schedule.

While a direct (full) analysis of the GGJ rewinding strategy [19] turns out to be complex, we are able to break it down into three different steps. Each step builds upon the previous one, with the final step yielding us the desired bound on the simulator query complexity. Below, we provide an informal overview of each of the three steps and refer the reader to the later sections for details.

Step 1. Lazy-KP with static scheduling: We first consider the warm-up case when scheduling of messages by the adversary is static. This means that the ordering of the messages of different sessions is decided by the adversary ahead of time and is fixed (and does not change upon rewinding by the simulator). Further, instead of directly analyzing the GGJ simulator [19], here we will analyze the query complexity of the (simpler) “lazy-KP” simulator [29,30,23] for the case where the simulator uses a splitting factor of n for rewinding. That is,

⁴ More concretely, we wish to bound the first message in the protocol where the simulator is forced to query the trusted party in order to obtain the function output.

⁵ Note that the output message of a session may appear more than once in the simulation transcript if the simulator employs rewinding.

during simulation, each thread is divided into n equal parts, and, each resulting part is rewound individually (resulting in different threads of execution).

In this case, we are able to prove that the simulator makes at most $O(1)$ queries to the ideal functionality in any given session. This is done by relying on the following fact. Say that the point p_i does *not* occur in a given thread. Then, since the adversary only employs static scheduling, this would mean that the point p_i also cannot occur in any threads resulting from rewinding this thread. Thus, the proof reduces to a counting argument on the number of threads resulting from rewinding the part of the main thread containing p_i . If d is the depth of recursion for our recursive rewinding schedule, then we are able to show that there are at most $O(2^d)$ threads containing point p_i . However, the depth d will be a constant for lazy-KP simulation with splitting factor n .

Step 2. Lazy-KP with *dynamic* scheduling: Now we analyze a general adversary that may dynamically change the ordering of the messages across different sessions upon being rewound. Hence, different threads of execution may have different ordering of the messages. We shall continue to analyze the lazy-KP simulation strategy with splitting factor n .

In this case, we prove that the simulator makes at most $O(\log(n))$ queries to the ideal functionality in any given session. The key difficulty in this case is that even if a given thread does *not* contain the point p_i , the threads resulting from its rewinding may still have p_i . Hence, it seems hard to rule out the possibility that p_i may show up in a large number of threads throughout the simulation.

To overcome this problem, we rely on the following fact: once the point p_i is seen in the main thread of execution, it cannot occur in any thread arising out of the main thread *after* that point. We also observe that *before* this point is seen in the main thread, there seems hope to rule out its occurrence in a “large” number of look ahead threads. This relies on the symmetry of the main and the look-ahead threads, and, on the fact that this point has roughly equal probability of occurring first in the main thread vs occurring first in any given look ahead thread. This step of the proof is more involved than the first step and we refer the reader to Section 4 for details.

Step 3. *Sparsifying* the lazy-KP simulation: In the final step, we analyze the *sparse* rewinding strategy of [19]. Very roughly speaking, the sparse rewinding strategy of [19] aims to rewind the adversary in “as few places as possible” while still solving all the sessions. More specifically, there is a cost associated with creating each look ahead, and, the goal of the rewinding strategy is to solve all sessions while minimizing the cost.

The sparse rewinding strategy of [19] builds upon the lazy-KP simulator with splitting factor n . Very roughly, [19] pick a subset of the total threads resulting out of the lazy-KP simulation, and choose to execute only the threads in the subset (while ignoring the remaining threads by aborting them at their start). In more detail, at each level of recursion, [19] randomly chooses $\frac{\text{polylog}(n)}{n}$ fraction of the total threads and execute them while ignoring the rest. Interestingly, Goyal et. al [19] show that, if one uses protocols with somewhat higher round

complexity, all the session will still be solved even though most of the look-ahead threads are never executed.

The key idea of our final step is to leverage this sparsification in order to reduce the number of queries from $O(\log(n))$ from the previous step to $O(1)$. Recall from above that if we were to use the full lazy-KP simulation, the point p_i would have occurred at $O(\log(n))$ places in the entire simulation. However, now, in the GGJ rewinding strategy, it will occur only $O(1)$ times because most of the threads will never be executed. More details are given in section 5.

2 Our Model

Let n denote the security parameter. We consider malicious, static adversaries that choose whom to corrupt before the start of any protocol. We work in the static input setting, i.e., we assume that the inputs of the honest parties in all sessions are fixed at the beginning. We do not require fairness.

Ideal model. In the ideal world experiment, there is a trusted party for computing the desired two-party functionality f . Let there be two parties P_1 and P_2 that are involved in multiple, say $m = m(n)$, evaluations of f . Let \mathcal{S} denote the adversary. The ideal world execution (parametrized by λ) proceeds as follows.

- I. Inputs:** P_1 and P_2 obtain a vector of m inputs, denoted \vec{x} and \vec{y} respectively. The adversary is given auxiliary input z , and chooses a party to corrupt. Without loss of generality, we assume that the adversary corrupts P_2 . The adversary receives the input vector \vec{y} of the corrupted party.
- II. Session initiation:** The adversary initiates a new session by sending a start-session message to the trusted party. The trusted party then sends (start-session, i) to P_1 , where i is the index of the session.
- III. Honest parties send inputs to trusted party:** Upon receiving the message (start-session, i) from the trusted party, P_1 sends (i, x_i) to the trusted party, where x_i denotes its input for session i .
- IV. Adversary sends input to trusted party and receives output:** At any point, the adversary may send a message $(i, \ell, y'_{i,\ell})$ to the trusted party for any $y'_{i,\ell}$ of its choice. It receives back $(i, \ell, f(x_i, y'_{i,\ell}))$ where x_i is the input value that P_1 previously sent to the trusted party for session i . For any i , the trusted party accepts at most λ tuples indexed by i from the adversary.
- V. Adversary instructs trusted party to answer honest party:** When the adversary sends a message of the type (output, i, ℓ) to the trusted party, the trusted party sends $(i, f(x_i, y'_{i,\ell}))$ to P_1 , where x_i and $y'_{i,\ell}$ denote the respective inputs sent by P_1 and adversary for session i .
- VI. Outputs:** The honest party P_1 always outputs the values $f(x_i, y'_{i,\ell})$ that it obtained from the trusted party. The adversary may output an arbitrary efficient function of its auxiliary input z , input vector \vec{y} and the outputs obtained from the trusted party.

The ideal execution of a function \mathcal{F} with security parameter n , input vectors \vec{x}, \vec{y} and auxiliary input z to \mathcal{S} , denoted $\text{Ideal}_{\mathcal{F}, \mathcal{S}}(n, \vec{x}, \vec{y}, z)$, is defined as the output pair of the honest party and \mathcal{S} from the above ideal execution.

Definition 2.1 (λ -Ideal Query Simulator). *Let \mathcal{S} be a non-uniform probabilistic (expected) PPT machine representing the ideal-model adversary. We say that \mathcal{S} is a λ -ideal query simulator if it makes at most λ output queries per session in the above ideal experiment.*

Real model. Let Π be a two-party protocol for computing \mathcal{F} . Let \mathcal{A} denote a non-uniform probabilistic polynomial-time adversary that controls either P_1 or P_2 . The parties run concurrent executions of the protocol Π , where the honest party follows the instructions of Π in each execution i using input x_i . The scheduling of all messages is controlled by the adversary. At the conclusion of the protocol, an honest party computes its output as prescribed by the protocol. Without loss of generality, we assume the adversary outputs exactly its entire view of the execution of the protocol.

The real concurrent execution of Π with security parameter n , input vectors \vec{x} , \vec{y} and auxiliary input z to \mathcal{A} , denoted $\text{Real}_{\Pi, \mathcal{A}}(n, \vec{x}, \vec{y}, z)$, is defined as the output pair of the honest party and \mathcal{A} , resulting from the above real-world process.

Definition 2.2 (λ -Secure Concurrent Computation in the MIQ Model). *A protocol Π is said to λ -securely realize a functionality \mathcal{F} under concurrent self composition in the MIQ model if for every real model non-uniform PPT adversary \mathcal{A} , there exists a non-uniform (expected) PPT λ -ideal query simulator \mathcal{S} such that for all polynomials $m = m(n)$, every pair of input vectors $\vec{x} \in X^m$, $\vec{y} \in Y^m$, every $z \in \{0, 1\}^*$,*

$$\{\text{Ideal}_{\mathcal{F}, \mathcal{S}}(n, \vec{x}, \vec{y}, z)\}_{n \in \mathbb{N}} \stackrel{c}{=} \{\text{Real}_{\Pi, \mathcal{A}}(n, \vec{x}, \vec{y}, z)\}_{n \in \mathbb{N}}$$

3 Framework for Concurrent Extraction

The Setting. Consider the following two-party computation protocol $\Pi = (P_1, P_2)$:

- **Stage 1:** First, P_1 and P_2 interact in the commit phase of an execution of an extractable commitment scheme $\langle C, R \rangle$ (described below) where P_2 acts as the committer, committing to a random string, and, P_1 acts as the receiver.
- **Stage 2:** At the end of the commitment protocol, P_1 sends a special message msg to P_2 .

Now, consider the scenario where P_1 and P_2 are interacting in multiple concurrent executions of Π . Suppose that P_2 is corrupted. Our goal is to design a simulator algorithm \mathcal{S} that satisfies the following two properties:

- **Extraction in all sessions:** \mathcal{S} must successfully extract the value committed by adversarial P_2^* in each execution of Π .
- **Minimize the query parameter:** Let λ denote the upper bound on the number of times the special message msg_s of any session s appears in the entire simulation transcript. We refer to λ as the *query parameter*. Then, the goal of \mathcal{S} is to minimize the query parameter.

In the next subsection, we describe the extractable commitment scheme $\langle C, R \rangle$ from [30]. Later, in Sections 4 and 5, we analyze the “lazy-KP” rewinding strategy [29,30,23] and the “sparse” rewinding strategy of Goyal, Gupta and Jain (GGJ) [19].

3.1 Extractable Commitment Protocol $\langle C, R \rangle$

Let $\text{COM}(\cdot)$ denote the commitment function of a non-interactive perfectly binding string commitment scheme. Let $\ell = \omega(\log n)$. Let $N = N(n)$ which will be determined later depending on the extraction strategy. The commitment scheme $\langle C, R \rangle$ between the committer C and the receiver R is described as follows.

Commit Phase: This consists of two stages, namely, the Init stage and the Challenge-Response stage, described below:

INIT: To commit to a n -bit string σ , C chooses $(\ell \cdot N)$ independent random pairs of n -bit strings $\{\alpha_{i,j}^0, \alpha_{i,j}^1\}_{i,j=1}^{\ell, N}$ such that $\alpha_{i,j}^0 \oplus \alpha_{i,j}^1 = \sigma$ for all $i \in [\ell], j \in [N]$. C commits to all these strings using COM , with fresh randomness each time. Let $B \leftarrow \text{COM}(\sigma)$, and $A_{i,j}^0 \leftarrow \text{COM}(\alpha_{i,j}^0)$, $A_{i,j}^1 \leftarrow \text{COM}(\alpha_{i,j}^1)$ for every $i \in [\ell], j \in [N]$.

CHALLENGE-RESPONSE: For every $j \in [N]$, do the following:

- **Challenge :** R sends a random ℓ -bit challenge string $v_j = v_{1,j}, \dots, v_{\ell,j}$.
- **Response :** $\forall i \in [\ell]$, if $v_{i,j} = 0$, C opens $A_{i,j}^0$, else it opens $A_{i,j}^1$ by sending the decommitment information.

Open Phase: C opens all the commitments by sending the decommitment information for each one of them. R verifies the consistency of the revealed values. This completes the description of $\langle C, R \rangle$.

Notation. We introduce some terminology that will be used in the remainder of this paper. We refer to the committed value σ as the *preamble secret*. A slot_i of the commitment scheme consists of the i 'th Challenge message from R and the corresponding Response message from C . Thus, in the above protocol, there are N slots.

4 Lazy-KP Extraction Strategy

In this section, we discuss the “lazy-KP” rewinding strategy [29,30,23] with a “splitting factor” of n . We note that the idea of using a large splitting factor was first used in [27].

For this strategy, we will first prove that $\lambda = \mathcal{O}(1)$ for *static* adversarial schedules. Next, we will prove that for *dynamic* schedules, $\lambda = \mathcal{O}(\log n)$. In both of these results, the constants in \mathcal{O} depend on number of sessions started by the concurrent adversary.

Lazy-KP Simulator. The rewinding strategy of the lazy-KP simulator is specified by the Lazy-KP-SIMULATE procedure. Very roughly, the simulator divides the current thread (given as input) into n equal parts and then rewinds each

part individually and recursively. The input to the Lazy-KP-SIMULATE procedure consists of a triplet $(\ell, \text{hist}, \mathcal{T})$. The parameter ℓ denotes the adversary's messages to be explored, the string hist is a transcript of the *current* thread of execution, and \mathcal{T} is a table containing the contents of all the adversary's messages explored so far (to extract the preamble secrets and for sending the Stage 2 special message in protocol Π in any session).

The simulation is performed by invoking the procedure Lazy-KP-SIMULATE with appropriate parameters. Let $m = \text{poly}(n)$ denote the number of concurrent sessions in the adversarial schedule. Then, the Lazy-KP-SIMULATE procedure is invoked with input $(m(N+1), \emptyset, \emptyset)$, where $m(N+1)$ is the total number of adversary's messages in a schedule of m sessions. The Lazy-KP-SIMULATE procedure is described in Figure 1. Note that here (similar to [27]) we divide each thread into n parts. In other words, we consider a splitting factor of n . For

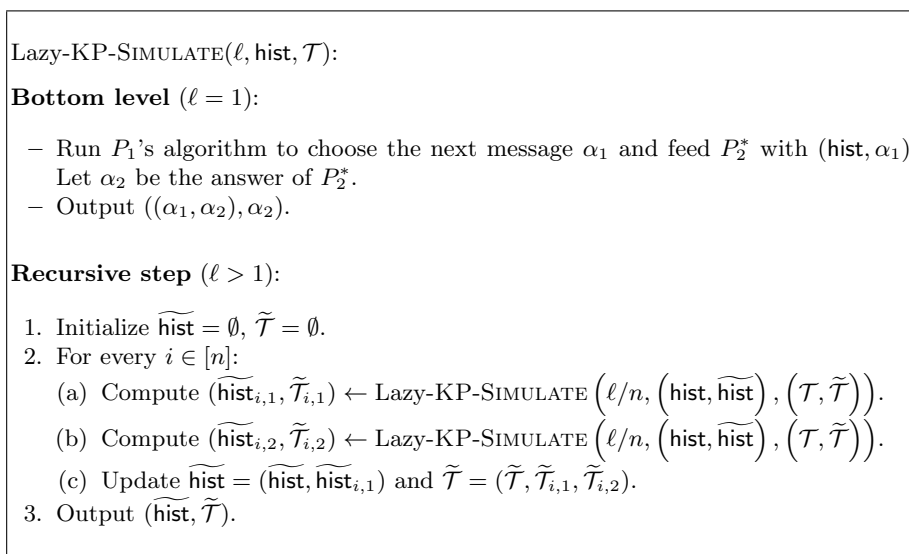


Fig. 1. Lazy-KP Simulator with splitting factor n . Even though the messages in $\{\widetilde{\text{hist}}_{i,2}\}$ do not appear in the output, some of them do appear in $\widetilde{\mathcal{T}}$.

every session s consisting of an execution of Π , the goal of the simulator is to find two instances of any slot $i \in [N]$ of the commitment protocol $\langle C, R \rangle$ where the simulator's challenges are different and adversary responds with a valid response to each challenge. Note that in this case, the simulator can extract the preamble secret of $\langle C, R \rangle$ from the two responses of the adversary. On the other hand, if the simulation reaches Stage 2 in Π at any time, without having extracted the preamble secret from the adversary, then it gives up the simulation and outputs \perp . In this case, we say the simulator *gets stuck*.

It follows from [29] that the lazy-KP simulator (as described above) gets stuck with only negligible probability.

4.1 Terminology for Concurrent Simulation

We introduce some terminology and definitions regarding concurrent simulation that will be used in the rest of the paper.

Execution Thread. Consider any adversary that starts $m = \text{poly}(n)$ number of concurrent sessions of Π . In order to extract the preamble secret in every session, the simulator creates multiple execution threads, where a thread of execution is a simulation of (part of) the protocol messages in the m sessions. We differentiate between the following:

Main Thread vs Look-ahead Thread: The *main thread* is a simulation of a complete execution of the m sessions, and this is the execution thread that is output by the simulator. In addition, from any execution thread, the simulator may create other threads by rewinding the adversary to a previous state and continuing the execution from that state. Such a thread is called a *look-ahead thread*. Note that a look-ahead thread can be created from another look-ahead thread.

Complete vs Partial Thread: We say that an execution thread T is a *complete* thread if it shares a prefix with the main thread: it starts where the main thread starts, and, continues until it is terminated by the simulator. Other threads that start from intermediary points of the simulation are called *partial* threads. Note that by definition, the main thread is a complete thread. In general, a complete thread may consist of various partial threads. Various complete threads may overlap with each other. For simplicity of exposition, unless necessary, we will not distinguish between complete and partial threads in the sequel.

Simulation Transcript. The simulation transcript is the set of all the messages between the simulator and the adversary during the simulation of all the concurrent sessions. In particular, this includes the messages that appear on the main thread as well as all the look-ahead threads.

Simulation Index. Consider $m = \text{poly}(n)$ concurrent executions of Π . Let $M = m(2N + 2)$, where $2N + 2$ is the round complexity of Π . Then, a simulation index i denotes the point where the i 'th message (out of a maximum of M messages) is sent on any complete execution thread in the simulation transcript.

Note that a simulation index i may appear *multiple* times over various threads in the simulation transcript. However, a simulation index i can appear at most once on any given thread (complete or partial). In particular, every simulation index $i \in [M]$ appears on the main thread (unless the main thread is aborted prematurely). Further, if a look-ahead thread T was created from a thread at simulation index i , then only simulation indices $j > i$ can appear on T .

Static vs Dynamic Scheduling. Consider the concurrent execution of $m = \text{poly}(n)$ instances of Π . Recall that the adversary controls the scheduling of the protocol messages across the m sessions. We say that a concurrent schedule is

static if the scheduling of the protocol messages is decided by the adversary ahead of time and does not change upon rewindings. Thus, in a static schedule, protocol messages appear in the *same* order on every complete thread. In particular, for every $i \in [M]$, every instance of a simulation index i in the simulation transcript corresponds to the *same* message index $j \in [2N+2]$ of the *same* session s (out of the m sessions). However note that the actual content of the j 'th message may differ on every execution thread.

We say that a concurrent schedule is *dynamic* if at any point during the execution, the adversary may decide which message to schedule next based on the protocol messages received so far. Therefore, in a dynamic schedule, the ordering of messages may be *different* on different execution threads in the simulation. In particular, each instance of a simulation index i may correspond to a *different* message j_i of a *different* session s_i .

Recursion Levels. We define recursion levels of simulation and count the number of threads at each recursion level for the lazy-KP simulator. We say that the main thread is at recursion level 0. Note that the Lazy-KP-SIMULATE divides the main thread of execution into n parts and executes each part twice. This results in $2n$ execution threads, n of which are part of the main thread, while the remaining n are look-ahead threads. All of these $2n$ threads are said to be at recursion level 1. Now, each of these threads at recursion level 1 is divided into n parts and each part is executed twice. This creates $2n$ threads at recursion level 2. Since there are $2n$ threads at recursion level 1, in total, we have $(2n)^2$ threads at recursion level 2. (Again, out of these $(2n)^2$ threads, $2n^2$ threads actually lie on the $2n$ threads at level 1.) This process is continued recursively. At recursion level ℓ , there are $(2n)^\ell$ threads. Since there are $m(2N+2)$ messages across the m sessions, the depth of recursion is a constant c' , where $c' = c + \log(2N+2)$ when $m = n^c$. Then, at recursion level c' , there are $(2n)^{c'}$ threads.

Sibling Threads. Consider Figure 2 where a thread T at some recursion level ℓ is divided into $n = 4$ parts, which leads to the creation of 8 threads at recursion level $\ell + 1$. Each pair of threads (T_i, T'_i) that are started from the same point are referred to as *sibling* threads.




Fig. 2. One recursion step for splitting factor 4. Every T_i and T'_i are sibling threads.

4.2 Analysis of λ for Static Schedules

We start by analyzing the lazy-KP extraction strategy for static schedules. Let $\lambda_{\text{lazy-KP}}$ denote the query parameter for the lazy-KP simulator.

Theorem 4.1. *For any constant c and any concurrent execution of $m = n^c$ instances of Π where the scheduling of messages is static, $\lambda_{\text{lazy-KP}} = 2^{c'}$, where $c' = c + \log(2N + 2)$.*

In order to prove Theorem 4.1, we use the following lemma that follows by a simple counting argument (the proof is deferred to the full version).

Lemma 4.2. *For any constant c and any concurrent execution of $m = n^c$ instances of Π , the simulation transcript generated by the lazy-KP simulator is such that every simulation index $i \in [M]$ appears $2^{c'}$ times, where $c' = c + \log(2N + 2)$.*

Consider any session s . From the definition of static scheduling, we have that for every $j \in [2N+2]$, if the j 'th message of session s appears at simulation index i on any thread, then every instance of simulation index i in the simulation transcript corresponds to the j 'th message of session s . Now, from Lemma 4.2, since each simulation index appears $2^{c'}$ times in the simulation transcript, we have that the special message of every session s appears $2^{c'}$ times in the simulation. Thus, we have that $\lambda_{\text{lazy-KP}} = 2^{c'}$ for static schedules.

4.3 Analysis of λ for Dynamic Schedules

Theorem 4.3. *For any polynomial $m = \text{poly}(n)$, for any concurrent execution of m instances of Π (with possibly dynamic scheduling of messages), $\lambda_{\text{lazy-KP}} = \mathcal{O}(\log n)$ except with negligible probability.*

Proof of Theorem 4.3. Fix any session s out of the $m = n^c$ sessions. Note that the special message msg_s of session s appears exactly once on the main thread. Let i_{main} denote the simulation index where msg_s appears on the main thread. Now, we will count:

1. The number of times msg_s appears in the simulation transcript before i_{main} . Let δ_1 denote this number.
2. The number of times msg_s appears in the simulation transcript at i_{main} or after i_{main} . Let δ_2 denote this number.

Thus, the total number of times msg_s appears in the simulation transcript is $\delta_1 + \delta_2$. It suffices to prove that $\delta_1 + \delta_2 = \mathcal{O}(\log n)$.

Let i_1, \dots, i_k be the *distinct* simulation indices where msg_s appears in the simulation transcript. Let i_1, \dots, i_k be ordered, i.e., for every $\ell \in [k-1]$, $i_\ell < i_{\ell+1}$. Let $k_1 \leq k$ be such that $i_{k_1} < i_{\text{main}}$ and $i_{k_1+1} \geq i_{\text{main}}$.

Lemma 4.4. *For any $\ell \in [k]$, the probability that msg_s does not appear on the main thread at simulation index i_ℓ is at most $(1 - \frac{1}{c'})$.*

Proof. Consider the simulation index i_1 . From Lemma 4.2, we have that i_1 appears on $2^{c'}$ threads in the simulation transcript. Let $T[i_1] = T_1, \dots, T_{2^{c'}}$ denote these threads. Now, let q be such that the special message msg_s appears at simulation index i_1 on q of these $2^{c'}$ threads. Let $T^*[i_1] = T_1^*, \dots, T_q^*$ denote these q threads. Let T_{main} denote the main thread. Then, we have that:

$$\Pr[T_{\text{main}} \in T^*[i_1]] = \frac{q}{2^{c'}} \quad (1)$$

To see this, recall that the Lazy-KP-SIMULATE procedure uses uniformly random coins on each execution thread, and follows the same strategy. Thus, the view of the adversary is indistinguishable on each thread. In particular, if p is the probability that a message α appears on a thread T and m' appears on its sibling thread T' with, then with probability $p - \text{negl}(n)$, m' appears on T and m appears on T' . (This is the “symmetry” property for threads in the lazy-KP simulation.) Therefore, Equation 1 follows.

From Equation 1, we have that:

$$\Pr[T_{\text{main}} \notin T^*[i_1]] = 1 - \frac{q}{2^{c'}}$$

Note that the above probability is maximum when $q = 1$. Hence, we have that:

$$\Pr[\text{msg}_s \text{ does not occur on main thread at } i_1] \leq 1 - \frac{1}{2^{c'}}. \quad (2)$$

Now, consider simulation index i_2 . Again, from Lemma 4.2, we have that i_2 appears on $2^{c'}$ threads. Let $T[i_2]$ denote the set of these threads. Now, note that msg_s cannot appear on the look-ahead threads $T \in T^*[i_1] \cap T[i_2]$. Thus, following Equation 2, we have that:

$$\Pr[\text{msg}_s \text{ does not occur on main thread at } i_2] \leq 1 - \frac{1}{2^{c''}}.$$

where $c'' \leq c'$. Continuing the same argument, we have that for every $\ell \in [k-1]$,

$$\Pr[\text{msg}_s \text{ is not on main thread at } i_{\ell+1}] \leq \Pr[\text{msg}_s \text{ is not on main thread at } i_\ell]$$

Thus, for every i_ℓ , we have that the probability that msg_s does not occur on main thread at i_ℓ is at most $1 - \frac{1}{c'}$.

Computing δ_1 . Now, note that $(1 - \frac{1}{c'})^t = \text{negl}(n)$ for $t = \omega(\log n)$. Therefore, we have that $k_1 = \mathcal{O}(\log n)$. Now, since each of the simulation indices i_1, \dots, i_{k_1} appears $2^{c'}$ times in the simulation transcript, we have that:

$$\delta_1 \leq 2^{c'} \mathcal{O}(\log n) \quad (3)$$

Computing δ_2 . We now compute the value of γ_2 . Towards this, let us suppose that for every simulation index $i \in [\ell]$, the Lazy-KP-SIMULATE procedure runs all threads starting from simulation index i in *parallel*. That is, Lazy-KP-SIMULATE performs one step of execution on each of these threads. It then performs the

next execution step on each of these threads, and so on. Note that this is without loss of generality since the Lazy-KP-SIMULATE procedure runs all such threads *independently*.

Now, we first observe that msg_s cannot appear on a look-ahead thread that starts at a simulation index $i > i_{\text{main}}$. Thus, to compute δ_2 , we only need to consider the look-ahead threads that started at simulation indices $i < i_{\text{main}}$ and did not finish before reaching i_{main} . Let T_{good} denote the set of such threads.

By using Lemma 4.2, we can claim that $|T_{\text{good}}| \leq 2^{c'}$. Then, assuming the worst case where msg_s appears on each thread $T \in T_{\text{good}}$, we have that $\delta_2 \leq 2^{c'}$.

5 GGJ Extraction Strategy

In this section, we discuss the GGJ extraction strategy [19] and analyze the query complexity parameter for the same. Unlike [19] that used a splitting factor of 2, we will work with n as the splitting factor. For this strategy, we will prove that for every concurrent schedule of polynomial number of sessions, the query parameter $\lambda = \mathcal{O}(1)$. Here, the constant in \mathcal{O} depends on the number of concurrent sessions.

Overview. Roughly speaking, the GGJ rewinding strategy can be viewed as a “stripped down” version of the lazy-KP simulation strategy. In particular, unlike lazy-KP that executes *every* thread at every recursion level, here we only execute a small fraction of them. The actual threads that are to be executed are chosen uniformly at random, at every level. It is shown in GGJ that by slightly increasing the round complexity – (roughly) $N = n^2$ from $N = n$, executing a $\frac{\text{polylog} n}{N}$ fraction of threads at every level is sufficient to extract the preamble secret in every session.

We describe the GGJ rewinding strategy in two main steps:

1. We first describe an algorithm **Sparsify** that essentially selects which threads to execute in the lazy-KP recursion tree (Section 5.1).
2. Next, we describe the actual GGJ simulation procedure **GGJ-SIMULATE** that is essentially the same as the Lazy-KP-SIMULATE strategy, except that it only executes the threads selected by **Sparsify** (Section 5.2).

5.1 The Sparsification Procedure

We first describe the lazy-KP simulation tree and give a coloring scheme for the same. Next, we describe the **Sparsify** algorithm that takes the lazy-KP simulation tree as input and outputs a “trimmed” version of it that will correspond to the GGJ simulation tree.

Lazy-KP Simulation Tree. Let $m = n^c$ be the total number of concurrent sessions of Π started by an adversary \mathcal{A} . Then, the Lazy-KP-SIMULATE strategy for \mathcal{A} can be described by a $2n$ -ary tree $\text{Tree}_{\text{lazy-KP}}$ of constant depth c' where $c' = c + \log(2N + 2)$. The nodes in $\text{Tree}_{\text{lazy-KP}}$ are colored *white* or *black* as per the following strategy:

- The root node is colored white.
- Consider the $2n$ child nodes of any parent node. The odd numbered nodes are colored white and the even numbered nodes are colored black.

Let us explain our coloring strategy. The root node (which is colored white) corresponds to the main thread of execution. Each black colored node **Node** corresponds to a look-ahead thread that was forked from the thread corresponding to node **Parent(Node)**. A white colored node **Node** (except the root node) corresponds to a thread T' that is a part of the thread T corresponding to node **Parent(Node)**.

Figure 3 denotes the lazy-KP simulation tree for splitting factor $n = 2$ with white boxes representing white nodes and grey boxes representing black nodes.

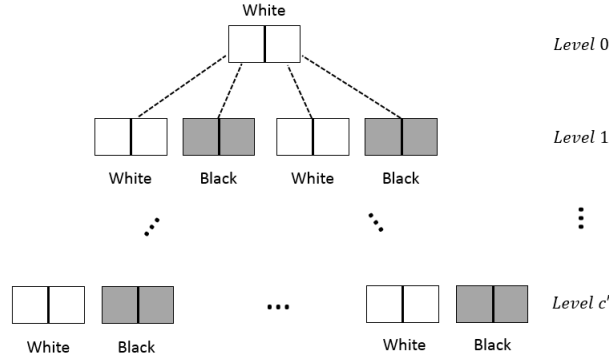


Fig. 3. The lazy-KP simulation tree for splitting factor 2.

Node Labeling. To facilitate the description of the GGJ simulation strategy, we first describe a simple tree node labeling strategy for $\text{Tree}_{\text{lazy-KP}}$. The root node is labeled 1. The i 'th child (out of $2n$ children) of the root node is labeled $(1, i)$. More generally, consider a node **Node** at level $\ell \in [c']$. Let **path** be its label. Then the i 'th child of **Node** is labeled (path, i) .

Below, whenever necessary, we shall refer to the nodes by their associated labels.

The Sparsify Procedure. Let p be such that $\frac{1}{p} = \frac{\text{polylog}(n)}{N}$. The Sparsify function transforms the lazy-KP simulation tree $\text{Tree}_{\text{lazy-KP}}$ into a “sparse” tree Tree_{sp} in the following manner.

Let the root node correspond to level 0 and the leaf nodes correspond to level c' . The Sparsify procedure starts at level 0 and traverses down $\text{Tree}_{\text{lazy-KP}}$, stopping at level c' . It performs the following steps at every level $\ell \in [c']$:

1. Choose $\frac{1}{p}$ fraction of the total black nodes at level ℓ , uniformly at random. Let B_ℓ denote the set of these nodes.
2. Delete from $\text{Tree}_{\text{lazy-KP}}$, every black node Node at level ℓ that is not present in set B_ℓ . Further, delete the entire subtree of Node from $\text{Tree}_{\text{lazy-KP}}$.

The resultant tree is denoted as Tree_{sp} . Looking ahead, we will describe the GGJ rewinding strategy as essentially a modification of Lazy-KP-SIMULATE in that it only executes the threads corresponding to the nodes in Tree_{sp} .

5.2 The GGJ-Simulate Procedure

The rewinding strategy of the GGJ simulator is specified by the GGJ-SIMULATE procedure. The input to the GGJ-SIMULATE procedure is a tuple $(\text{path}, \ell, \text{hist}, \mathcal{T})$. The parameter path denotes the label of the node in Tree_{sp} that is to be explored, ℓ denotes the number of adversary's messages to be explored (on the thread corresponding to the node labeled with path), the string hist is a transcript of the *current* thread of execution, \mathcal{T} is a table containing the contents of all the adversary's messages explored so far (to extract the preamble secrets and for sending the Stage 2 special message in Π in any session).

The simulation is performed by invoking the procedure GGJ-SIMULATE with appropriate parameters. Let $m = \text{poly}(n)$ denote the number of concurrent sessions in the adversarial schedule. Then, the GGJ-SIMULATE procedure is invoked with input $(1, m(N+1), \emptyset, \emptyset)$, where $m(N+1)$ is the total number of adversary's messages in a schedule of m sessions. The GGJ-SIMULATE procedure is described in Figure 4. Note that unlike [19], where each thread is recursively divided into two parts, here we divide each thread into n parts. In other words, we consider a splitting factor of n . For every session s consisting of an execution of Π , the goal of the simulator is to find two instances of any slot $i \in [N]$ of the commitment protocol $\langle C, R \rangle$ where the simulator's challenges are different and adversary responds with a valid response to each challenge. Note that in this case, the simulator can extract the preamble secret of $\langle C, R \rangle$ from the two responses of the adversary. On the other hand, if the simulation reaches Stage 2 in Π at any time, without having extracted the preamble secret from the adversary, then it gives up the simulation and outputs \perp . In this case, we say the simulator *gets stuck*.

It is implicit in [19] that the GGJ simulator (as described above) gets stuck with only negligible probability when $N = \mathcal{O}(n^2)$. We now analyze the query parameter λ_{GGJ} for the GGJ simulation strategy. A formal proof is deferred to the full version.

Theorem 5.1. *For every constant c , every $m = n^c$ number of concurrent executions of Π , the query parameter $\lambda_{\text{GGJ}} = \mathcal{O}(1)$, where the constant depends on c .*

Proof (Sketch). Fix any session s . We will show that the special message msg_s can appear at most $\mathcal{O}(1)$ times at each recursion level RL_ℓ . Then, since there are only a constant number of recursion levels, it will follow that $\lambda_{\text{GGJ}} = \mathcal{O}(1)$.

GGJ-SIMULATE($\text{path}, \ell, \text{hist}, \mathcal{T}$):

Bottom level ($\ell = 1$):

- Run P_1 's algorithm to choose the next message α_1 and feed P_2^* with (hist, α_1) . Let α_2 be the answer of P_2^* .
- Output $((\alpha_1, \alpha_2), \alpha_2)$.

Recursive step ($\ell > 1$):

1. Initialize $\widetilde{\text{hist}} = \emptyset, \widetilde{\mathcal{T}} = \emptyset$.
2. For every $i \in [n]$:
 - If node $(\text{path}, 2i - 1) \notin \text{Tree}_{\text{sp}}$, set $\widetilde{\text{hist}}_{i,1} = \emptyset, \widetilde{\mathcal{T}}_{i,1} = \emptyset$.
Else, compute:
 $(\widetilde{\text{hist}}_{i,1}, \widetilde{\mathcal{T}}_{i,1}) \leftarrow \text{GGJ-SIMULATE} \left((\text{path}, 2i - 1), \ell/n, (\text{hist}, \widetilde{\text{hist}}), (\mathcal{T}, \widetilde{\mathcal{T}}) \right)$.
 - If node $(\text{path}, 2i) \notin \text{Tree}_{\text{sp}}$, set $\widetilde{\text{hist}}_{i,2} = \emptyset, \widetilde{\mathcal{T}}_{i,2} = \emptyset$.
Else, compute:
 $(\widetilde{\text{hist}}_{i,2}, \widetilde{\mathcal{T}}_{i,2}) \leftarrow \text{GGJ-SIMULATE} \left((\text{path}, 2i), \ell/n, (\text{hist}, \widetilde{\text{hist}}), (\mathcal{T}, \widetilde{\mathcal{T}}) \right)$.
 - Update $\widetilde{\text{hist}} = (\widetilde{\text{hist}}, \widetilde{\text{hist}}_{i,1})$ and $\widetilde{\mathcal{T}} = (\widetilde{\mathcal{T}}, \widetilde{\mathcal{T}}_{i,1}, \widetilde{\mathcal{T}}_{i,2})$.
3. Output $(\widetilde{\text{hist}}, \widetilde{\mathcal{T}})$.

Fig. 4. GGJ Simulator with splitting factor n . Even though the messages in $\{\widetilde{\text{hist}}_{i,2}\}$ do not appear in the output, some of them do appear in $\widetilde{\mathcal{T}}$.

Towards that end, let's fix a recursion level ℓ . First recall from Theorem 4.3 that for the lazy-KP simulation strategy, $\lambda_{\text{lazy-KP}} = \mathcal{O}(\log n)$. In particular, this implies that at every recursion level ℓ in the lazy-KP simulation, msg_s for a session s appears on at most $\mathcal{O}(\log n)$ threads. Using the tree terminology as introduced earlier, we have that msg_s appears on (the threads corresponding to) at most $\mathcal{O}(\log n)$ black nodes at level ℓ in $\text{Tree}_{\text{lazy-KP}}$. Now, recall that at every level ℓ , the Sparsify procedure selects only $\frac{1}{p} = \frac{\text{polylog} n}{N}$ fraction of black nodes, uniformly at random, and deletes the rest of the black nodes. Using Chernoff bound, we can then show that the probability that Sparsify selects $\omega(1)$ black nodes containing msg_s is negligible.

6 From Concurrent Extraction to Concurrent Secure Computation

Theorem 6.1. *Assuming 1-out-of-2 oblivious transfer, for any efficiently computable functionality f there exists a protocol Π that $\mathcal{O}(1)$ -securely realizes f in the MIQ model.*

We construct such a protocol by following the exact recipe of [21,19]. We note that the works of [21,19] show how to compile a semi-honest secure computation protocol Π_{sh} for any functionality f into a new protocol Π that securely realizes f in the MIQ model. The core ingredient of their compiler is a concurrently extractable commitment $\langle C, R \rangle$: if there exists a concurrent simulator for $\langle C, R \rangle$ with query parameter λ , then the resultant (compiled) protocol Π λ -securely realizes f .

In order to prove Theorem 6.1, we construct such a protocol Π by simply plugging in our $O(n^2)$ -round extractable commitment scheme in the construction of [21,19]. Then, it follows from Theorem 5.1 that protocol Π $\mathcal{O}(1)$ -securely realizes f in the MIQ model, where the constant in \mathcal{O} depends on c , where n^c is the number of sessions opened by the concurrent adversary.

Fully concurrent PAKE in the plain model. Consider the PAKE functionality: it takes a password as input from each party, and, if they match, outputs a randomly generated key to both of them. The above protocol, when executed for the PAKE functionality gives a PAKE construction in the MIQ model where the simulator makes a constant number of queries per session in the ideal world. We then plug in Lemma 7 in [21] which shows that a PAKE construction in the MIQ model for a constant number of queries implies a concurrent PAKE as per the definition of Goldreich and Lindell [16] (with the modification that the constant in big O is adversary dependent). Put together, this gives us a construction of concurrent password-authenticated key exchange in the plain model.

References

1. Agrawal, S., Goyal, V., Jain, A., Prabhakaran, M., Sahai, A.: New impossibility results on concurrently secure computation and a non-interactive completeness theorem for secure computation. In: CRYPTO (2012)
2. Barak, B., Canetti, R., Nielsen, J., Pass, R.: Universally composable protocols with relaxed set-up assumptions. In: FOCS (2004)
3. Barak, B., Prabhakaran, M., Sahai, A.: Concurrent non-malleable zero knowledge. In: FOCS (2006)
4. Barak, B., Sahai, A.: How to play almost any mental game over the net - concurrent composition using super-polynomial simulation. In: Proc. 46th FOCS (2005)
5. Canetti, R., Lindell, Y., Ostrovsky, R., Sahai, A.: Universally composable two-party and multi-party secure computation. In: STOC (2002)
6. Canetti, R., Fischlin, M.: Universally composable commitments. In: CRYPTO (2001)
7. Canetti, R., Kilian, J., Petrank, E., Rosen, A.: Black-box concurrent zero-knowledge requires $\tilde{\Omega}(\log n)$ rounds. In: STOC. pp. 570–579 (2001)
8. Canetti, R., Kushilevitz, E., Lindell, Y.: On the limitations of universally composable two-party computation without set-up assumptions. In: Eurocrypt (2003)
9. Canetti, R., Lin, H., Pass, R.: Adaptive hardness and composable security in the plain model from standard assumptions. In: FOCS (2010)
10. Dolev, D., Dwork, C., Naor, M.: Nonmalleable cryptography. SIAM J. Comput. 30(2), 391–437 (electronic) (2000), preliminary version in STOC 1991.

11. Dwork, C., Naor, M., Sahai, A.: Concurrent zero-knowledge. In: STOC. pp. 409–418 (1998)
12. Garg, S., Goyal, V., Jain, A., Sahai, A.: Concurrently secure computation in constant rounds. In: Eurocrypt (2012)
13. Garg, S., Kumarasubramanian, A., Ostrovsky, R., Visconti, I.: Impossibility results for static input secure computation. In: CRYPTO (2012)
14. Goldreich, O., Micali, S., Wigderson, A.: How to play any mental game. In: STOC (1987)
15. Goldreich, O., Krawczyk, H.: On the composition of zero-knowledge proof systems. *SIAM J. Comput.* 25(1), 169–192 (Feb 1996), <http://epubs.siam.org/sam-bin/dbq/article/22068>, preliminary version appeared in ICALP’ 90.
16. Goldreich, O., Lindell, Y.: Session-key generation using human passwords only. In: CRYPTO. pp. 408–432 (2001)
17. Goldreich, O., Lindell, Y.: Session-key generation using human passwords only. *J. Cryptology* 19(3), 241–340 (2006)
18. Goyal, V.: Positive results for concurrently secure computation in the plain model. In: FOCS (2012)
19. Goyal, V., Gupta, D., Jain, A.: What information is leaked under concurrent composition. In: CRYPTO (2013)
20. Goyal, V., Jain, A.: On concurrently secure computation in the multiple ideal query model. In: Eurocrypt (2013)
21. Goyal, V., Jain, A., Ostrovsky, R.: Password-authenticated session-key generation on the internet in the plain model. CRYPTO (2010), full version available online
22. Katz, J.: Universally composable multi-party computation using tamper-proof hardware. In: Eurocrypt (2007)
23. Kilian, J., Petrank, E.: Concurrent and resettable zero-knowledge in poly-logalgorithm rounds. In: STOC (2001)
24. Lindell, Y.: Bounded-concurrent secure two-party computation without setup assumptions. In: STOC. pp. 683–692. ACM (2003)
25. Micali, S., Pass, R.: Local zero knowledge. In: STOC (2006)
26. Micali, S., Pass, R., Rosen, A.: Input-indistinguishable computation. In: FOCS (2006)
27. Pandey, O., Pass, R., Sahai, A., Tseng, W.L.D., Venkatasubramanian, M.: Precise concurrent zero knowledge. In: Eurocrypt (2008)
28. Pass, R.: Simulation in quasi-polynomial time, and its application to protocol composition. In: Eurocrypt (2003)
29. Pass, R., Tseng, W.L.D., Venkatasubramanian, M.: Concurrent zero knowledge, revisited. *Journal of Cryptology* 27(1), 45–66 (2014)
30. Prabhakaran, M., Rosen, A., Sahai, A.: Concurrent zero knowledge with logarithmic round-complexity. In: FOCS (2002)
31. Prabhakaran, M., Sahai, A.: New notions of security: achieving universal composability without trusted setup. In: STOC (2004)
32. Richardson, R., Kilian, J.: On the concurrent composition of zero-knowledge proofs. In: Eurocrypt (1999)
33. Yao, A.C.C.: How to generate and exchange secrets. In: FOCS (1986)