

Rate-Limited Secure Function Evaluation: Definitions and Constructions

Özgür Dagdelen¹, Payman Mohassel², and Daniele Venturi³

¹ Technische Universität Darmstadt, Germany

² University of Calgary, Canada

³ Aarhus University, Denmark

Abstract. We introduce the notion of rate-limited secure function evaluation (RL-SFE). Loosely speaking, in an RL-SFE protocol participants can monitor and limit the number of distinct inputs (i.e., *rate*) used by their counterparts in multiple executions of an SFE, in a private and verifiable manner. The need for RL-SFE naturally arises in a variety of scenarios: e.g., it enables service providers to “meter” their customers’ usage without compromising their privacy, or can be used to prevent oracle attacks against SFE constructions.

We consider three variants of RL-SFE providing different levels of security. As a stepping stone, we also formalize the notion of commit-first SFE (cf-SFE) wherein parties are committed to their inputs before each SFE execution. We provide compilers for transforming any cf-SFE protocol into each of the three RL-SFE variants. Our compilers are accompanied with simulation-based proofs of security in the standard model and show a clear tradeoff between the level of security offered and the overhead required. Moreover, motivated by the fact that in many client-server applications clients do not keep state, we also describe a general approach for transforming the resulting RL-SFE protocols into *stateless* ones.

As a case study, we take a closer look at the oblivious polynomial evaluation (OPE) protocol of Hazay and Lindell, show that it is commit-first and instantiate efficient rate-limited variants of it.

Keywords: secure function evaluation, foundations, secure metering, oracle attacks, oblivious polynomial evaluation.

1 Introduction

Secure function evaluation (SFE) allows a set of mutually distrustful parties to securely compute a function f of their private inputs. Roughly speaking, SFE protocols guarantee that the function is computed correctly and that the parties will not learn any information from the interaction other than their output and what is inherently leaked from it. Seminal results in SFE show that one can securely compute any functionality [29,30,15,8,2]. There has been a large number of follow-up work improving the security, strengthening adversarial models, and studying efficiency. Recent work on practical SFE has also led to real-world deployments [7,6], and the design and implementation of several SFE frameworks [24,5,12,20,21].

In practice, most applications of SFE considered in the literature need to accommodate *multiple executions* of a protocol.¹ Consider a client that searches for multiple patterns in a large text via a secure pattern matching protocol [17,19], searches several keywords in a private database via an oblivious keyword search [27,13], or an individual who needs to run a software diagnostic program, or an intrusion detection system (IDS) to analyze data via an oblivious branching program (OBP) or an automaton evaluation (OAE) protocol [22,28].

Invoking an SFE protocol multiple times raises important practical issues that are outside the scope of standard SFE, and hence are not addressed by the existing solutions. We point out two such issues and introduce *rate-limited* SFE as a means to address them. The reason for the choice of name is that rate-limiting is commonly used in network and web applications to refer to restrictions put on clients' usage (on a per user, or a per IP address basis). In this work we consider similar restrictions on a user's inputs to services that maybe implemented using SFE.

SECURE METERING OF SFE. Service providers tend to charge their clients according to their level of usage: a location-based service may wish to charge its clients based on the number of locations they use the service from; a database owner based on the number of distinct search queries; an IDS provider based on the number of suspicious files sent for vulnerability analysis. Service providers would be more willing to adopt SFE protocols if it is possible to efficiently enforce such a metering mechanism. The challenge is to do so without compromising the client's privacy, or allowing the server or the client to cheat the metering system.

ORACLE ATTACKS. Consider multiple executions of a two-party SFE protocol (such as those mentioned above), where the first party's input stays the same in different executions but the second party's input varies. A malicious second party who "adaptively" uses different inputs in each execution, can gradually learn significant information about the first party's input, and, in the worst case, fully recover it. For instance, consider an oblivious polynomial evaluation (OPE) protocol (e.g., used in oblivious keyword search) wherein the server holds a polynomial p while the client holds a private point x and wants to learn $p(x)$, but cannot learn more than it. Evaluating the polynomial p on sufficiently many points allows a malicious client to interpolate and recover p . A similar attack can be applied to OBP and OAE protocols to learn the private branching program or automaton which may embed propriety information. Learning attacks of this sort are well-understood and have been previously identified as important threats in the context of SFE; they are sometimes referred to as *oracle attacks* since the attacker has black-box access to input/output values from multiple executions (e.g., see the discussion in [1]).

A naïve solution to the problems discussed above is to limit the total number of executions of an SFE protocol, ignoring the actual input values. However, this approach does not provide a satisfactory solution in most scenarios. For

¹ Depending on the application, a subset of the participants may use the same input in different executions.

example, in case of secure metering, fixing an a priori upper bound on the total number of executions would mean charging legitimate clients multiple times for using the service with the same input; a disadvantage for clients who may need to use the same input from multiple devices, or reproduce a result due to communication errors, device upgrades, or perhaps to prove the validity of the outcome to a third-party by re-running the protocol. Similarly, in case of oracle attacks, clients need not be disallowed to use the same input multiple times since querying the same input many times does not yield new information to an attacker.

RATE-LIMITED SFE. A more accurate (and challenging) solution is to *limit* and/or *monitor* the number of distinct inputs used by an SFE participant in multiple executions. Obviously, this should be done in a secure and efficient manner, i.e., a party should not be able to exceed an agreed-upon limit, and its counterpart should not learn any additional information about his private inputs, or impose a lower limit than the one they agreed on.² We refer to the number of distinct inputs used by a participant as his *rate*, and call a SFE protocol that monitors/limits this number, a rate-limited SFE.

Of course, achieving RL-SFE is more costly than the naïve solution discussed above. However, at a minimum we require the proposed solution to avoid storing and/or processing the complete transcripts of all previous executions. (We discuss the exact overhead of our solutions in detail below.)

We note that the complementary question of what functions are *unsafe* for use in SFE (leak too much information) has also been studied, e.g., by combining SFE and differential privacy [3,26], or belief tracking techniques [25]. These works are orthogonal to ours, and can potentially be used in conjunction with rate-limited SFE as an *enforcement mechanism*. For instance, the former works can be invoked to negotiate on a function f with a measurable “safeness” from which the rate for each user can be derived. Subsequently, the abidance of this rate can be enforced through our rate-limited SFE.

Our Contribution. Motivated by the discussion above, we initiate the study of rate-limited SFE. For simplicity, in this paper we focus on the two-party case, but point out that the definitions and some of the constructions are easily extendible to the multiparty setting. Our main contributions are as follows.

DEFINITIONS. We introduce three definitions for rate-limited secure function evaluation: (i) rate-hiding, (ii) rate-revealing and (iii) pattern-revealing. All our definitions are in the real-world/ideal-world simulation paradigm and are concerned with *multiple* sequential executions of an SFE protocol. They reduce to the standard simulation-based definition (stand alone) for SFE, when applied to a single execution.

In a *rate-hiding* RL-SFE, in each execution, the only information revealed to the parties is whether the agreed-upon rate limit has been exceeded or not.

² In fact the problem becomes significantly easier when the parties are assumed to be semi-honest.

In a *rate-revealing* RL-SFE, the parties additionally learn the current rate (i.e., the number of distinct inputs used by their counterpart so far). In a *pattern-revealing* RL-SFE, parties also learn the pattern of occurrences of each other’s inputs during the previous executions. These notions provide a useful spectrum of tradeoffs between security and efficiency: our constructions become more efficient as we move to the more relaxed notions, to the extent that *our pattern-revealing transformation essentially adds no overhead* to the underlying SFE protocol.

COMMIT-FIRST SFE. In order to design rate-limited SFE protocols, we formalize the auxiliary notion of commit-first SFE (cf-SFE). Roughly speaking, a protocol is commit-first if it can be naturally divided into a (i) *committing phase*, where each party becomes committed to its input for the second phase, and (ii) a *function evaluation phase*, where the function f is computed on the inputs committed to in the first phase.³

We utilize cf-SFE as a stepping stone to design rate-limited SFE. It turns out that the separation between the input commitment phase and the function evaluation phase facilitates the design of efficient rate-limited SFE. In particular, now a party only needs to provide some evidence of a particular relation between the committed inputs in the first phase. In contrast, if we had not started with a commit-first protocol, such an argument would have involved the complete history of the transcripts for all the previous executions, rendering such an approach impractical.

The related notion of “evaluating on committed inputs” is well-known (e.g. see [15,23]), but we need and put forth a formal (and general) definition for cf-SFE in order to prove our RL-SFE protocols secure. We then show that several existing SFE constructions are either commit-first or can be efficiently transformed into one. Examples include variants of Yao’s garbled circuit protocol, the oblivious polynomial evaluation of Hazay and Lindell [16], the private set intersection protocol of Hazay and Nissim [18], and oblivious automaton evaluation of Gennaro *et al.* [14]. We also show that the GMW compiler [15], outputs a commit-first protocol. This is of theoretical interest as it provides a general compiler for transforming a semi-honest SFE protocol into a malicious cf-SFE (and eventually a rate-limited SFE using the compilers in this paper). We elaborate on these cf-SFE instantiations in the full version of this paper [11].

COMPILERS & TECHNIQUES. We design three compilers for transforming a cf-SFE into each of the three variants of RL-SFE discussed above, and provide simulation-based proofs of their security. All our compilers start from a cf-SFE protocol and add a “proof of repeated-input phase” between the committing phase and the function evaluation phase. An exception is our pattern-revealing compiler, where a proof of repeated-input is implicit given that we force the commitments to be deterministic. In our first compiler (rate-hiding), whenever the j -th execution begins, party P_1 first checks whether its input is “fresh” or has

³ Note that adding input commitments to the beginning of a protocol does not automatically yield a cf-SFE, since parties are not necessarily bound to using the committed inputs in their evaluation.

already been used in a previous run. In the former case, P_1 encrypts the value “1” and, otherwise, the value “0” using a semantically secure public-key encryption scheme (E, D) for which it holds the secret key sk . Denote the resulting ciphertext with c^j . Party P_1 forwards to P_2 a ZK proof of the following statement:

$$\begin{aligned} &(\text{“committed to old input”} \wedge E(0)) \\ &\vee (\text{“committed to new input”} \wedge E(1) \wedge \text{“}\sum_{i \leq j} D(sk, c^i) \leq \text{rate”}). \end{aligned}$$

Intuitively, the proof above only leaks the fact that the rate is not exceeded in the current execution, but nothing else. In order to generate this proof (resp. verify the proof generated by the counterpart), P_1 needs to store all the commitments and ciphertexts sent to (resp. received from) P_2 in previous executions.

For our second compiler (rate-revealing), we can do without the encryptions. Parties can instead prove a simpler statement giving evidence that the current (committed) input corresponds to one of the commitments the other party received earlier. Clearly, this approach reveals the current rate, but as we prove nothing more.

Finally, our third compiler (pattern-revealing) exploits a PRF to generate the randomness used in the committing phase of the underlying cf-SFE protocol. In this way, the commitment becomes deterministic (given the input), allowing the other party to check whether the current input has already been used and *in which runs*. This approach discloses the pattern of inputs used by the parties; on the other hand, it is extremely efficient adding little computational overhead (merely one invocation of a PRF) to the original cf-SFE protocol.

MAKING RL-SFE STATELESS. The above compilers suffer from the limitation that the parties need to keep a state which grows linearly in the total number of executions of the underlying SFE protocol. In many applications, clients do not keep state (and outsource this task to the servers), either due to lack of resources or because they need to use the service from multiple locations/devices. We show a general approach for transforming the stateful RL-SFE protocols generated above into *stateless* ones. Here, the client keeps merely a small secret (whose size is independent of the total number of executions), but is still able to prevent cheating by a malicious server, and preserve privacy of his inputs. At a high level, the transformation requires the client to store its *authenticated* (MACed) state information on the server side and retrieve/verify/update it on-the-fly as needed. We show how to apply this transformation to our rate-revealing compiler to obtain a stateless variant and prove its security. A similar technique can be applied to our rate-hiding compiler. Our pattern-revealing compiler is already stateless (client only needs to store a PRF key) for the party who plays the role of the client.

CASE STUDY. We take a closer look at the oblivious polynomial evaluation protocol of Hazay and Lindell [16]. Their protocol is secure against malicious adversaries. We show that it is also a commit-first OPE, by observing that a homomorphic encryption of the parties’ inputs together with ZK proofs of their validity, can be interpreted as a commitment to their inputs. This immediately

yields an efficient pattern-revealing RL-SFE for the OPE problem, based on the compiler we design. We also provide an efficient rate-hiding and rate-revealing RL-OPE by instantiating the ZK proofs for membership in the necessary languages, efficiently.

Roadmap. We start introducing notations and our model for commit-first SFE in Section 2 and 3. In Section 4 we give the definition of rate-limited SFE. A fortaste of our rate-hiding, rate-revealing and pattern-revealing compilers are given in Section 5, whereas Section 6 describes the stateless version of the rate-revealing compiler. Finally, Section 7 deals with concrete instantiations for the case of OPE.

2 Preliminaries

Throughout the paper, we denote the security parameter by λ . A function $negl(\lambda)$ is negligible in λ (or just negligible) if it decreases faster than the inverse of every polynomial in λ . A machine is said to run in polynomial-time if its number of steps is polynomial in the security parameter.

Let $X = \{X_\lambda\}_{\lambda \in \mathbb{N}}$ and $Y = \{Y_\lambda\}_{\lambda \in \mathbb{N}}$ be two distribution ensembles. We say X and Y are computationally indistinguishable (and we write $X \equiv_c Y$) if for every non-uniform polynomial-time adversary \mathcal{A} there exists a negligible function $negl$ such that $|\Pr[\mathcal{A}(X) = 1] - \Pr[\mathcal{A}(Y) = 1]| \leq negl(\lambda)$. Note that all our security statements can be straightforwardly proven for uniform polynomial-time adversaries, as well.

If x is a string, $|x|$ denotes the length of x . Vectors are denoted boldface; given vector \mathbf{x} , we write $\mathbf{x}[j]$ for the j -th element of \mathbf{x} . If \mathcal{X} is a set, $\#\mathcal{X}$ represents the number of elements in \mathcal{X} . When x is chosen randomly in \mathcal{X} , we write $x \leftarrow \mathcal{X}$. When \mathcal{A} is an algorithm, $y \leftarrow \mathcal{A}(x)$ denotes a run of \mathcal{A} on input x and output y ; if \mathcal{A} is randomized, then y is a random variable and $\mathcal{A}(x; r)$ denotes a run of \mathcal{A} on input x and random coins r .

Our compilers make use of standard cryptographic primitives. Due to space limitations, we assume familiarity of these primitives and define them formally in the full version of this paper [11].

3 Commit-First SFE

In this section, we formally define the notion of *commit-first secure function evaluation* (cf-SFE). Our three compilers Ψ_{RH} , Ψ_{RR} and Ψ_{PR} for designing rate-limited SFE, leverage commit-first protocols as a building block. We call a protocol π commit-first if it can be naturally divided into two phases. In the first phase (committing phase), both parties P_1 and P_2 become committed to their inputs. At the end of this phase, no information about the parties' inputs is revealed (the hiding property), and neither party can use a different input than what it is committed to in the remainder of the protocol (the binding property). In the

second phase (function evaluation phase), the function f will be computed on the inputs committed to in the last phase.

We now specify the two separate phases. Consider a polynomial-time functionality $f = (f_1, f_2)$ with $f_i : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^*$. Then, a cf-SFE protocol $\pi = (\pi_1, \pi_2)$ for evaluating f on parties' inputs x_1 and x_2 proceeds as follows.

Committing Phase: Parties P_1 and P_2 execute π_1 which is defined by the functionality $((x_1, r_1), (x_2, r_2)) \mapsto (C_2(x_2, r_2), C_1(x_1, r_1))$. Note that the commitment schemes C_1, C_2 can be arbitrary schemes (often different for each cf-SFE protocol), as long as they satisfy the hiding and the binding properties required.

Function Evaluation Phase: Afterwards, P_1 and P_2 execute π_2 on the same inputs as in the committing phase; π_2 is defined by the functionality $((x_1, C_2(x_2)), (x_2, C_1(x_1))) \mapsto (f_1(x_1, x_2), f_2(x_1, x_2))$. Note that P_1 and P_2 , can use their state information from the previous phase in the function evaluation phase, too.

Next, we formalize the security definition for a cf-SFE using the real/ideal world simulation paradigm.

THE REAL WORLD. In each execution, a non-uniform adversary \mathcal{A} following an arbitrary polynomial-time strategy can send messages in place of the corrupted parties (whereas the honest parties continue to follow π). Let $i \in \{1, 2\}$ be the index of the corrupted party. A real execution of $\pi = (\pi_1, \pi_2)$ on inputs (x_1, x_2) , auxiliary input z to \mathcal{A} and the security parameter λ , denoted by $\text{REAL}_{\pi, \mathcal{A}(z), i}^{\text{cf-SFE}}(x_1, x_2, \lambda)$ is defined as the output of the honest party and the adversary upon execution of π .

THE IDEAL WORLD. Let $i \in \{1, 2\}$ be the index of the corrupted party. We define the ideal world in two steps. During the ideal execution, the honest party sends its input x_{3-i} , and a uniformly random string r_{3-i} used by the commitment scheme, to the trusted party. Party P_i which is controlled by the ideal adversary \mathcal{S} , called the simulator, may either abort (sending a special symbol \perp) or send input x'_i , and an arbitrary randomness r'_i (not necessarily uniform) chosen based on the auxiliary input z , and P_i 's original input x_i . Denote by $((x'_1, r'_1), (x'_2, r'_2))$ the values received by the trusted party. If the trusted party receives \perp , the value \perp is forwarded to both P_1 and P_2 and the ideal execution terminates; else the trusted party computes $\gamma_1 = C_1(x'_1; r'_1)$ and $\gamma_2 = C_2(x'_2; r'_2)$, respectively. The TTP sends γ_{3-i} to \mathcal{S} , which can either continue or abort by sending \perp to the TTP. In case of an abort, the TTP sends \perp to the honest party; otherwise, it sends γ_i .

In the second phase, the honest party continues the ideal execution by sending to the TTP a continue flag, or abort by sending \perp . \mathcal{S} sends either \perp or continue based on the auxiliary input z , P_i 's original input, and the value γ_{3-i} . If the trusted party receives \perp , the value \perp is forwarded to both P_1 and P_2 and the ideal execution terminates; else the trusted party computes $y_1 = f_1(x'_1, x'_2)$ (resp. $y_2 = f_2(x'_1, x'_2)$).

The TTP sends y_i to \mathcal{S} . At this point, \mathcal{S} can decide whether the trusted party should continue, and thus send the output y_{3-i} to the honest party, or halt, in which case the honest party receives \perp . The honest party outputs the received value. The simulator \mathcal{S} outputs an arbitrary polynomial-time computable function of (z, x_i, y_i) .

The ideal execution of f on inputs (x_1, x_2) , auxiliary input z to \mathcal{S} and security parameter λ , denoted by $\text{IDEAL}_{f, \mathcal{C}_1, \mathcal{C}_2, \mathcal{S}(z), i}^{\text{cf-SFE}}(x_1, x_2, \lambda)$ is defined as the output of the honest party and the simulator.

EMULATING THE IDEAL WORLD. We define a secure commit-first protocol π as follows:

Definition 1 (Commit-first Protocols). *Let π and f be as above. We say that π is a commit-first protocol for computing $f = (f_1, f_2)$ in the presence of malicious adversaries with abort if for every non-uniform probabilistic polynomial-time adversary \mathcal{A} in the real world there exists a non-uniform probabilistic polynomial-time simulator \mathcal{S} in the ideal world, such that for every $i \in \{1, 2\}$,*

$$\left\{ \text{REAL}_{\pi, \mathcal{A}(z), i}^{\text{cf-SFE}}(x_1, x_2, \lambda) \right\}_{x_1, x_2, z, \lambda} \equiv_c \left\{ \text{IDEAL}_{f, \mathcal{C}_1, \mathcal{C}_2, \mathcal{S}(z), i}^{\text{cf-SFE}}(x_1, x_2, \lambda) \right\}_{x_1, x_2, z, \lambda}$$

where $x_1, x_2, z \in \{0, 1\}^*$ and $\lambda \in \mathbb{N}$.

4 Rate-Limited Secure Function Evaluation

In this section, we introduce three notions for rate-limited secure function evaluation (RL-SFE). In particular, we augment the standard notion of two-party SFE by allowing each player to monitor and/or limit, the number of distinct inputs (the *rate*) the other player uses in multiple executions. The idea is that each party can abort the protocol if the number of distinct inputs used in the previous executions raises above a threshold $\nu \in \mathbb{N}$. We call this threshold the *rate limit*, i.e. the maximum number of allowable executions with distinct inputs.

Naturally, our security definitions for RL-SFE are concerned with *multiple* executions of an SFE protocol and reduce to the standard simulation-based definition (stand alone) for SFE, when applied to a single run. We call a sequence of executions of a protocol π (ν_1, ν_2) -limited if party P_1 (resp. P_2) can use at most ν_1 (resp. ν_2) distinct inputs in the executions. In this work, we assume that the executions take place *sequentially*, i.e. one execution after the other. We emphasize that the inputs used by the parties in each execution can depend on the transcripts of the previous executions, but honest parties will always use fresh randomness in their computation.

We provide three security definitions for rate-limited SFE: (i) rate-hiding, (ii) rate-revealing and (iii) pattern-revealing. In a *rate-hiding* RL-SFE, at the end of each execution, the only information revealed to the parties (besides the output from the function being computed), is whether the agreed-upon rate limit (threshold) has been exceeded or not, but nothing else. In a *rate-revealing* RL-SFE, in addition to the above, parties also learn the current rates

(i.e., the number of distinct inputs used by their counterpart so far). Finally, in a *pattern-revealing* RL-SFE, parties further learn the pattern of occurrences of each others' inputs in the previous executions. In particular, each party learns which executions were invoked by the same input and which ones used different ones, but nothing else.

HIGH LEVEL DESCRIPTION. Let $f = (f_1, f_2)$ be a pair of polynomial-time functions such that f_i is of type $f_i : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^*$. Consider an arbitrary number ℓ of sequential executions of two-party SFE protocol π for evaluating f on parties' inputs. During the j -th execution, party P_i has input x_i^j and should learn $y_i^j = f_i(x_1^j, x_2^j)$. We will define rate-limited SFE in the general case where *both* parties are allowed to change their input in each execution. The case of oracle attacks and secure metering, where one party's input is fixed and the other party's input changes, are found as a special case. (In the case of secure metering one can also think that a change in the service provider's input reflects a software update.)

In the ideal world, during the j -th execution, each party sends its input to a trusted authority. The following is then performed for both $i = 1, 2$. The trusted party checks whether value x_i^j was already sent in a previous execution; in case it was not, a new entry (x_i^j, j) is stored in an initially empty set \mathcal{X}_i . Otherwise, the index $j' < j$ corresponding to such input is recovered. Whenever $\#\mathcal{X}_i$ exceeds κ_i the trusted party aborts. Otherwise, the current outputs $y_i^j = f_i(x_1^j, x_2^j)$ are computed. Finally: (i) in the rate-hiding definition party P_i learns only y_i^j ; (ii) in the rate-revealing definition party P_i learns also $\#\mathcal{X}_{3-i}$, i.e. the (partial) total number of distinct inputs used by P_{3-i} until the j -th execution; (iii) in the pattern-revealing definition party P_i learns j' , i.e. the index corresponding to the query where x_i^j was asked for the first time. Note that if the rate is exceeded, the trusted party aborts here, but, equivalently, we could simply ignore this execution and still allow to query previous inputs in subsequent executions.

We formalize the above intuitive security notions for all three flavors using the simulation-based ideal/real world paradigm. We first review the real execution which all three notions share.

THE REAL WORLD. In each execution, a non-uniform adversary \mathcal{A} following an arbitrary polynomial-time strategy can send messages in place of the corrupted party (whereas the honest party continues to follow π). Let $i \in \{1, 2\}$ be the index of the corrupted party. The j -th *real* execution of π on inputs (x_1^j, x_2^j) , auxiliary input z^j to \mathcal{A} and security parameter λ , denoted by $\text{REAL}_{\pi, \mathcal{A}(z^j), i}^{\kappa}(x_1^j, x_2^j, \lambda)_j$ is defined as the output of the honest party and the adversary in the j -th real execution of π . We denote by $\text{REAL}_{\pi, \mathcal{A}(\mathbf{z}), i}^{\kappa}(\mathbf{x}_1, \mathbf{x}_2, \lambda, \ell)$ the accumulative distribution at the end of the ℓ -th execution, i.e.,

$$\text{REAL}_{\pi, \mathcal{A}(\mathbf{z}), i}^{\kappa}(\mathbf{x}_1, \mathbf{x}_2, \lambda, \ell) = \text{REAL}_{\pi, \mathcal{A}(z^1), i}^{\kappa}(x_1^1, x_2^1, \lambda)_1, \dots, \text{REAL}_{\pi, \mathcal{A}(z^\ell), i}^{\kappa}(x_1^\ell, x_2^\ell, \lambda)_\ell$$

where $\mathbf{x}_1 = (x_1^1, \dots, x_1^\ell)$, $\mathbf{x}_2 = (x_2^1, \dots, x_2^\ell)$ and $\mathbf{z} = (z^1, \dots, z^\ell)$.

THE IDEAL WORLD. The trusted party keeps two sets \mathcal{X}_1 , and \mathcal{X}_2 initially set to \emptyset . Let $i \in \{1, 2\}$ be the index of the corrupted party. During the j -th ideal

execution, the honest party sends its input to the trusted party. Party P_i , which is controlled by the ideal adversary \mathcal{S} , called the simulator, may either abort (sending a special symbol \perp) or send input x_i^{tj} to the trusted party chosen based on the auxiliary input z^j , P_i 's original input x_i^j , and its view in the previous $j - 1$ ideal executions. Denote with (x_1^{tj}, x_2^{tj}) the values received by the trusted party (note that if $i = 2$ then $x_1^{tj} = x_1^j$).

If the trusted party receives \perp , the value \perp is forwarded to both P_1 and P_2 and the ideal execution terminates; else when the trusted party receives x_1^{tj} as the first party's input, it checks whether an entry $(x_1^{tj}, j') \in \mathcal{X}_1$ already exists; if so, it sets $J_1 = j'$. Otherwise, it creates a new entry (x_1^{tj}, j) , adds it to \mathcal{X}_1 , and sets $J_1 = j$. An identical procedure is applied to input of the second party x_2^{tj} to determine an index J_2 . At the end of the j -th ideal execution if $\sigma_1 := \#\mathcal{X}_1 \geq \nu_1$ or $\sigma_2 := \#\mathcal{X}_2 > \nu_2$, the value \perp is forwarded to both P_1 and P_2 and the ideal execution terminates. Otherwise, the pair $(y_1^j, y_2^j) = (f_1(x_1^{tj}, x_2^{tj}), f_2(x_1^{tj}, x_2^{tj}))$ is computed.

At this point, the ideal executions will be different depending on the variant of RL-SFE being considered.

Rate-Hiding. The trusted party forwards to the malicious party P_i the output y_i^j . At this point, \mathcal{S} can decide whether the trusted party should continue, and thus send the pair y_{3-i} to the honest party, or halt, in which case the honest party receives \perp .

Rate-Revealing. The trusted party forwards to the malicious party P_i the pair (y_i^j, σ_{3-i}) . At this point, \mathcal{S} can decide whether the trusted party should continue, and thus send the pair (y_{3-i}^j, σ_i) to the honest party, or halt, in which case the honest party receives \perp .

Pattern-Revealing. The trusted party forwards to the malicious party P_i the pair (y_i^j, J_{3-i}) . The integer $1 \leq J_{3-i} \leq j$ represents the index of the first execution where the input x_{3-i}^j has been used. At this point, \mathcal{S} can decide whether the trusted party should continue, and thus send the pair (y_{3-i}^j, J_i) to the honest party, or halt, in which case the honest party receives \perp .

The honest party outputs the received value. The simulator \mathcal{S} outputs an arbitrary polynomial-time computable function of (z^j, x_i^j, y_i^j) .

The j -th ideal execution of f on inputs (x_1^j, x_2^j) , auxiliary input z^j to \mathcal{S} and security parameter λ , denoted by $\text{IDEAL}_{f, \mathcal{S}(z^j), i}^{\nu-X}(x_1^j, x_2^j, \lambda)_j$ is defined as the output of the honest party and the simulator in the above j -th ideal execution. Here, $X \in \{\text{RH}, \text{RR}, \text{PR}\}$ determines the flavor of rate-limited SFE. We denote by $\text{IDEAL}_{f, \mathcal{S}(\mathbf{z}), i}^{\nu-X}(\mathbf{x}_1, \mathbf{x}_2, \lambda, \ell)$ the accumulative distribution at the end of the ℓ -th execution, i.e.,

$$\text{IDEAL}_{f, \mathcal{S}(\mathbf{z}), i}^{\nu-X}(\mathbf{x}_1, \mathbf{x}_2, \lambda, \ell) = \text{IDEAL}_{f, \mathcal{S}(z^1), i}^{\nu-X}(x_1^1, x_2^1, \lambda)_1, \dots, \text{IDEAL}_{f, \mathcal{S}(z^\ell), i}^{\nu-X}(x_1^\ell, x_2^\ell, \lambda)_\ell$$

where $\mathbf{x}_1 = (x_1^1, \dots, x_1^\ell)$, $\mathbf{x}_2 = (x_2^1, \dots, x_2^\ell)$ and $\mathbf{z} = (z^1, \dots, z^\ell)$.

EMULATING THE IDEAL WORLD. Roughly speaking, ℓ sequential executions of a protocol π are secure under the rate limit $\nu = (\nu_1, \nu_2)$ if the real executions

can be simulated in the above mentioned ideal world. More formally, we define a secure (ν_1, ν_2) -limited protocol π as follows:

Definition 2 (RL-SFE). *Let π and f be as above, and consider $\ell = \text{poly}(\lambda)$ sequential executions of protocol π . For $X \in \{\text{RH}, \text{RR}, \text{PR}\}$, we say protocol π is a secure X ν -limited SFE for computing $f = (f_1, f_2)$, in presence of malicious adversaries with abort with $\nu = (\nu_1, \nu_2)$, if for every non-uniform probabilistic polynomial-time adversary \mathcal{A} there exists a non-uniform probabilistic polynomial-time simulator \mathcal{S} , such that for every $i \in \{1, 2\}$,*

$$\left\{ \text{REAL}_{\pi, \mathcal{A}(\mathbf{z}), i}^{\nu}(\mathbf{x}_1, \mathbf{x}_2, \lambda, \ell) \right\}_{\mathbf{x}_1, \mathbf{x}_2, \mathbf{z}, \lambda} \equiv_c \left\{ \text{IDEAL}_{f, \mathcal{S}(\mathbf{z}), i}^{\nu-X}(\mathbf{x}_1, \mathbf{x}_2, \lambda, \ell) \right\}_{\mathbf{x}_1, \mathbf{x}_2, \mathbf{z}, \lambda}$$

where $\mathbf{x}_1, \mathbf{x}_2, \mathbf{z} \in (\{0, 1\}^*)^\ell$, such that $|\mathbf{x}_1[j]| = |\mathbf{x}_2[j]|$ for all j , and $\lambda \in \mathbb{N}$.

It is easy to see that the rate-hiding notion is strictly stronger than the rate-revealing notion, which in turn is strictly stronger than the pattern-revealing notion. A proof to this fact can be found in the full version [11].

5 Compilers for Rate-Limited SFE

In this section, we introduce our three compilers to transform an arbitrary (two-party) cf-SFE protocol into a *rate-limited* protocol for the same functionality.

Our first compiler Ψ_{RH} achieves the notion of rate-hiding RL-SFE through the use of general ZK proofs and (additively) homomorphic public key encryption. Our second compiler Ψ_{RR} achieves the notion of rate-revealing RL-SFE and is more efficient in that it needs to prove a simpler statement and does not rely on homomorphic encryption. Our last compiler Ψ_{PR} introduces essentially no overhead and avoids the use of general ZK proofs, yielding our third notion of pattern-revealing RL-SFE.

Let π_f be a two-party (single-run) commit-first protocol for secure function evaluation of a function $f = (f_1, f_2)$ (cf. Definition 1). Our compilers get as input (a description of) π_f , together with the rate $\nu = (\nu_1, \nu_2)$, and the number of executions ℓ , and output (a description of) $\hat{\pi}_f \leftarrow \Psi(\pi_f, \nu, \ell)$. The compilers are functionality preserving, meaning that protocol $\hat{\pi}$ repeatedly computes the same functionality f .

Due to space limitations, in this section we only provide a full description and analysis for the rate-revealing compiler. The other two compilers (rate-hiding Ψ_{RH} , and pattern-revealing Ψ_{PR}) are only covered at a high level here. The complete descriptions and analyses are given in the full version of this paper [11].

5.1 A Rate-Hiding Compiler

THE OVERVIEW. We naturally divide the cf-SFE protocol into a committing phase and a function evaluation phase and introduce a new phase in between where P_1 and P_2 convince each other that they have not exceeded the rate limit.

The latter step is achieved as follows. Whenever one of the parties is going to use a “fresh” input, it transmits an encryption of “1” to the other party; otherwise, it sends an encryption of “0”. The encryptions are obtained using a CPA-secure (homomorphic) PKE scheme $(\tilde{G}, \tilde{E}, \tilde{D})$. Then, the party proves in ZK that “the last commitment transmitted hides an already used input *and* it encrypted 0, *or* the last commitment transmitted hides a fresh input *and* it encrypted 1 *and* the sum of all the plaintexts, encrypted until now, does not exceed the rate”. A successful verification of this proof convinces the other party that the rate is not exceeded, leaking nothing more than this. We instantiate such ZK proofs for the OPE problem in Section 7. Notice that to generate such a proof each party needs to store all the ciphertexts transmitted to the other player, together with all the inputs and randomness used to generate the previous commitments. On the other hand, to verify the other party’s proof, one needs to store the ciphertexts and the commitments received in all earlier executions. The remainder of the messages exchanged during each execution, however, can be discarded.

Theorem 1. *Let π_f be a cf-SFE securely evaluating function f and $(\tilde{G}, \tilde{E}, \tilde{D})$ be a CPA-secure PKE scheme. Then $\hat{\pi}_f \leftarrow \Psi_{\text{RH}}(\pi_f, \mathfrak{r}_1, \mathfrak{r}_2, \ell)$ is a secure rate-hiding $(\mathfrak{r}_1, \mathfrak{r}_2)$ -limited protocol for the function f .*

5.2 A Rate-Revealing Compiler

THE OVERVIEW. Once again, we divide the cf-SFE protocol into a committing phase and a function evaluation phase and introduce a new phase in between where P_1 and P_2 convince each other that the current input has already been used in a previous execution. Note that the parties need to maintain a state variable Γ collecting the input commitments sent and received in all earlier executions. During the j -th execution, given a list of input commitments (and the corresponding inputs and randomness) for all the previous executions, party P_i can prove in ZK that the input commitment generated in the current execution is for the same value as one of the commitments collected previously. Party P_{3-i} also needs to collect the same set of commitments in order to verify the statement proven by P_i . The remainder of the messages exchanged during each execution, however, can be discarded. We note that while in general efficient ZK proofs of repeated inputs might be hard to find, for discrete-logarithm based statements, there exist efficient techniques for proving such statements. We refer the reader to the full version for more details. We also instantiate such ZK proofs for the OPE problem in Section 7. A complete description of the compiler is depicted in Figure 1. We prove the following result:

Theorem 2. *Let π_f be a cf-SFE securely evaluating function f . Then $\hat{\pi}_f \leftarrow \Psi_{\text{RR}}(\pi_f, \mathfrak{r}_1, \mathfrak{r}_2, \ell)$ is a secure rate-revealing $(\mathfrak{r}_1, \mathfrak{r}_2)$ -limited protocol for f .*

5.3 A Pattern-Revealing Compiler

In this section, we introduce a more efficient compiler Ψ_{PR} for designing rate-limited SFE. Given as input a cf-SFE protocol, our compiler Ψ_{PR} outputs a




Fig. 1. A compiler for rate-revealing rate-limited SFE

weaker form of rate-limited SFE where each party not only learns the current rate for its counterpart during each execution, but also the pattern of already used inputs. The main advantage is that this new compiler adds very little overhead to the original cf-SFE.

THE OVERVIEW. The idea is as follows. Besides their input, each party also stores a secret key for a PRF (a different key for each party). Before invoking the commit-first SFE protocol, each player generates the randomness it needs for the committing phase by applying the PRF on the *chosen input* for this execution. With this modification in place, the committing phase for each party becomes deterministic. If a party uses the same input in two executions, the two commitments its counterpart receives will be identical. As a result, to prove a repeated-input, each party can compare the commitment for the current execution with those used in the previous ones, and determine if the input is new or being repeated (hence also revealing the pattern). Note that the commitments still provide the required hiding and binding properties. The only overhead imposed by this compiler is the application of a PRF to generate the randomness for the committing phase.

Theorem 3. *Let π_f be a cf-SFE securely evaluating function f . Then $\hat{\pi}_f \leftarrow \Psi_{PR}(\pi_f, \nu_1, \nu_2, \ell)$ is a secure pattern-revealing (ν_1, ν_2) -limited SFE for f .*

6 Making the Compilers Stateless

One drawback of the compilers described in the previous section is that both P_1 and P_2 need to maintain state. To some extent, this assumption is necessary. It is not too hard to see that RL-SFE is impossible to achieve if neither party is keeping any information about the previous executions (we omit a formal

argument of this statement). However, as discussed earlier, in many natural client-server applications of SFE in the real world, it is reasonable to assume that the servers keep state, while the clients typically do not.

In this section, we show how to modify the compilers from Section 5 in such a way that only one of the parties needs to keep state. Our solution is efficient and works for all three compilers we discussed earlier. Throughout this section, we assume P_1 is the client and P_2 is the server. Server P_2 receives no output (as it is usually the case in the client-server setting) and wants to enforce the rate limit ν for the client. Although P_1 does not maintain any state, it needs to make sure that P_2 handles the rate, honestly. On the other hand, the server also needs to be convinced that the client is not cheating, by exceeding the rate limit ν .

THE OVERVIEW. Note that in the stateful versions of our compilers, P_1 needs to keep state in order to generate a ZK proof of repeated inputs, and verify the corresponding statement being proven by P_2 . Since we are only enforcing the rate for P_1 , we can eliminate the latter ZK proofs, and focus on the first one. Although our approach is general, for the sake of simplicity, we describe it in relation to our rate-revealing compiler from Section 5.2. The same idea can be applied to make the other compilers stateless. The basic idea is simple: We ask the server to store the list of all the commitments previously sent by P_1 sends the list to the client, during each run. For this simple approach to work, we need to address several important issues:

- For the client to learn the current rate and the previously queried inputs before each execution, it needs to store these values on the server side in a secure way. This can be easily addressed by having P_1 encrypt the message and randomness for each commitment (using a symmetric-key encryption) and send it along with the commitment itself. P_1 will just keep the private key for the encryption scheme.
- The client needs to verify that the list of commitments it receives from the server are the original commitments it sent in the previous executions. To do so, in each run P_1 computes a MAC ϕ of the string obtained by hashing all the commitments (i.e., the concatenation of the list it obtains from the server and the one it creates in the current execution) and sends it to the server.⁴ In each execution, it requests this MAC, the list of commitments along with the ciphertext storing the inputs and random coins from the server. Due to the unforgeability of the MAC, the server will only be able to use a correct list of commitments, previously issued and MACed by the client itself.
- It may seem that the above solution still allows the server to cheat and only send a subset of the commitment list along with a tag generated for that subset in one of the earlier executions, to the client. This would potentially make the input to the current execution look “new” and allow the server to decrease the rate. The client would not be able to detect this attack since it does not keep state and does not know the total number of commitments.

⁴ To save on computation, one could let the client obtain the previous hash value and compute the new one via *incremental* hashing [9,4].

However, a more careful inspection shows that the above does not really constitute an attack. In fact, the tag ϕ already *binds* the current rate to the current list of commitments, and prevents the server from decreasing the rate in this fashion. In particular, it is hard for the server to cook-up a state such that the verification of the tag is successful, and the client will think its rate is already exceeded when it is not. Essentially, coming up with such a state requires to find a collision in H or forging a tag for a fake list of commitments.

A detailed description of the compiler Ψ_{RR} and a proof of the following theorem are given in our full version [11].

Theorem 4. *Let π_f be a cf-SFE securely evaluating function $f = (f_1, -)$ and (G, T, V) be a UNF-CMA MAC scheme, $(\tilde{G}, \tilde{E}, \tilde{D})$ be a CPA-secure PKE scheme, and H being picked from a family of CRHFs. Then, $\hat{\pi}_f \leftarrow \Psi_{\text{RR}}(\pi_f, \mathfrak{r}, \ell)$ is a secure rate-revealing \mathfrak{r} -limited protocol for f .*

7 Rate-Limited OPE

Hazay and Lindell [16] design an efficient two-party protocol for oblivious polynomial evaluation (OPE) with security against malicious adversaries. In an OPE protocol, the first party holds a value t while the second party holds a polynomial p of degree d . Their goal is to let the first party learn $p(t)$ without revealing anything else. The protocol takes advantage of an additively homomorphic encryption scheme (Paillier’s encryption) and efficient ZK proofs of a few statements related to the encryption scheme. While the authors (only) prove security against malicious adversaries, we observe that, with a small modification, their construction is indeed a commit-first protocol for OPE as well.

FIRST PARTY’S COMMITMENT. Consider an additively homomorphic encryption scheme (G, E, D) . The first few steps performed by the first party (the party holding the value t) are as follows: (i) it runs the key generation for the encryption scheme to generate a key pair $(\text{pk}, \text{sk}) \leftarrow G(1^\lambda)$, accompanied by a ZK proof of knowledge of the secret key; (ii) then, it encrypts powers of t , i.e. $E(\text{pk}, t), E(\text{pk}, t^2), \dots, E(\text{pk}, t^d)$, and sends the resulting ciphertexts along with a ZK proof of the validity of the ciphertexts to the other party.

We observe that sending $E(\text{pk}, t)$ and a ZK proof of its validity constitutes a commitment by the first party to its input t . This commitment scheme realizes the ideal functionality of the first phase in our definition of commit-first protocols. (Recall that this means the simulator can extract both the input and the randomness used to generate the commitment.) In particular, a careful inspection of the security proof of [16] reveals that the simulator can extract both t and the randomness used to encrypt it during the simulation. Extracting the randomness is possible since in Paillier’s encryption scheme, given the secret key sk and a ciphertext c , one can recover both the randomness and the message.

SECOND PARTY’S COMMITMENT. The commitment of the second party to its input polynomial is slightly more subtle, and requires a small modification to the

original design. In the first few steps, the second party does the following: (i) it runs the key generation to generate a key pair $(\mathbf{pk}', \mathbf{sk}') \leftarrow \mathbf{G}(1^\lambda)$, accompanied by a ZK proof of knowledge of the secret key; (ii) it computes $((\mathbf{E}(\mathbf{pk}', q_1), \mathbf{E}(\mathbf{pk}', p - q_1)), \dots, (\mathbf{E}(\mathbf{pk}', q_s), \mathbf{E}(\mathbf{pk}', p - q_s)))$ where q_i 's are random polynomials of degree d for some security parameter s ; (iii) it sends all the ciphertext pairs along with ZK proofs of the fact that the homomorphic addition of every pair encrypts the same polynomial (i.e., p), to the first party. We need to slightly modify this step to realize our ideal commitment functionality: For the first pair of ciphertexts, the second party will also include a ZK proof of validity of $(\mathbf{E}(\mathbf{pk}', q_1), \mathbf{E}(\mathbf{pk}', p - q_1))$.

The pair of ciphertexts $(\mathbf{E}(\mathbf{pk}', q_1), \mathbf{E}(\mathbf{pk}', p - q_1))$ and the accompanied ZK proof of their validity, constitute the commitment by the second party to its input polynomial p . Once again, we note that the simulator in the proof is able to extract q_1 , p , and the randomness used in the two encryptions, due to the randomness recovering property of Paillier's encryption. The proof of security provided in [16] can be easily modified to show the commit-first property of the above-mentioned variant of their OPE construction.

Claim. The modified oblivious polynomial evaluation protocol of [16] is a commit-first SFE with security against malicious adversaries.

7.1 ZK Proofs for Rate-Limited OPE

We now explain how to derive rate-limited OPE protocols from the scheme of [16], by giving concrete instantiation of our compilers from Section 5 and 6.

RATE-REVEALING OPE. Consider first our rate-revealing compiler from Figure 1. A proof of repeated-input, here, is equivalent to proving a statement for the following language:

$$\mathcal{L}^{\text{OPE}}(n) = \left\{ (\mathbf{pk}, \hat{c}, c_1, \dots, c_n) : \exists \lambda, r \text{ s.t. } (\mathbf{pk}, \mathbf{sk}) \leftarrow \mathbf{G}(1^\lambda, r) \text{ and } \left. \begin{array}{l} (\mathbf{D}(\mathbf{sk}, \hat{c}) = \mathbf{D}(\mathbf{sk}, c_1) \vee \dots \vee \mathbf{D}(\mathbf{sk}, \hat{c}) = \mathbf{D}(\mathbf{sk}, c_n)) \end{array} \right\},$$

where the ciphertexts c_1, \dots, c_n are encryptions of the inputs for n previous executions of the OPE protocol. The ciphertext \hat{c} is the encryption of the input for the current execution.

Such a proof can be obtained by exploiting ZK proofs for the languages $\mathcal{L}^{\text{zero}}$ and $\mathcal{L}^{\text{mult}}$ defined in [16]. Informally, a valid proof of a statement in the language $\mathcal{L}^{\text{zero}}$ says that a ciphertext is an encryption of 0. Language $\mathcal{L}^{\text{mult}}$ allows us to prove that given three ciphertexts, one of them decrypts to the product of the other two underlying plaintexts. Denote the plaintext for each c_i by m_i and the one for \hat{c} by \hat{m} . The high level idea is to have the prover compute $\mathbf{E}(\mathbf{pk}, (\hat{m} - m_1) \cdots (\hat{m} - m_n))$, prove correctness of this computation and show that the final ciphertext is an encryption of zero. This clearly ensures correctness when the current input equals one of the inputs used in previous executions. See the full version [11] for a complete description.

RATE-HIDING OPE. In the rate-hiding case, besides a standard proof of the statement “a ciphertext is a valid encryption of bit b ”, the prover also needs

to prove that: (i) “the current commitment corresponds to a fresh input”; (ii) “given a collection of ciphertexts, the sum of the corresponding plaintexts is below some threshold \mathfrak{n} ”.

Note that a proof for the first statement is equivalent to proving that an element is *not* in \mathcal{L}^{ope} (denoted by \mathcal{L}^{ope}). Moreover, we show that a proof for the second statement can also be reduced to a proof of membership in \mathcal{L}^{ope} by relying on the homomorphic properties of the underlying encryption scheme. It remains to show ZK proofs for \mathcal{L}^{ope} . It is possible to do so using range proofs, but we show a simple and more efficient construction.

Using techniques of [10], the proofs discussed above can be combined (via conjunctive/disjunctive formulas) to generate a ZK proof of membership for the language used in our rate-hiding compiler. A detailed description of the compiler and the proof of the theorem below is given in this paper’s full version [11].

Acknowledgements. Daniele Venturi acknowledges support from the Danish National Research Foundation and The National Science Foundation of China (under the grant 61061130540) for the Sino-Danish Center for the Theory of Interactive Computation, and also from the CFEM research center (supported by the Danish Strategic Research Council) within which part of this work was performed. Özgür Dagdelen was supported by CASED (<http://www.cased.de>).

References

1. Barni, M., Failla, P., Kolesnikov, V., Lазzeretti, R., Sadeghi, A.-R., Schneider, T.: Secure Evaluation of Private Linear Branching Programs with Medical Applications. In: Backes, M., Ning, P. (eds.) ESORICS 2009. LNCS, vol. 5789, pp. 424–439. Springer, Heidelberg (2009)
2. Beaver, D., Goldwasser, S.: Multiparty Computation with Faulty Majority. In: Brassard, G. (ed.) CRYPTO 1989. LNCS, vol. 435, pp. 589–590. Springer, Heidelberg (1990)
3. Beimel, A., Nissim, K., Omri, E.: Distributed private data analysis: On simultaneously solving how and what. CoRR, abs/1103.2626 (2011)
4. Bellare, M., Goldreich, O., Goldwasser, S.: Incremental Cryptography: The Case of Hashing and Signing. In: Desmedt, Y.G. (ed.) CRYPTO 1994. LNCS, vol. 839, pp. 216–233. Springer, Heidelberg (1994)
5. Bogdanov, D., Laur, S., Willemson, J.: Sharemind: A Framework for Fast Privacy-Preserving Computations. In: Jajodia, S., Lopez, J. (eds.) ESORICS 2008. LNCS, vol. 5283, pp. 192–206. Springer, Heidelberg (2008)
6. Bogdanov, D., Talviste, R., Willemson, J.: Deploying secure multiparty computation for financial data analysis. Technical report, Cryptology ePrint, 2011/662
7. Bogetoft, P., Christensen, D.L., Damgård, I., Geisler, M., Jakobsen, T., Kroigaard, M., Nielsen, J.D., Nielsen, J.B., Nielsen, K., Pagter, J., Schwartzbach, M., Toft, T.: Secure Multiparty Computation Goes Live. In: Dingledine, R., Golle, P. (eds.) FC 2009. LNCS, vol. 5628, pp. 325–343. Springer, Heidelberg (2009)
8. Chaum, D., Crépeau, C., Damgård, I.: Multiparty unconditionally secure protocols. In: ACM Symposium on Theory of Computing, pp. 11–19. ACM (1988)
9. Chaum, D., van Heijst, E., Pfitzmann, B.: Cryptographically Strong Unconditionally Secure Signatures, Unconditionally Secure for the Signer. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 470–484. Springer, Heidelberg (1992)

10. Cramer, R., Damgård, I., Schoenmakers, B.: Proof of Partial Knowledge and Simplified Design of Witness Hiding Protocols. In: Desmedt, Y.G. (ed.) CRYPTO 1994. LNCS, vol. 839, pp. 174–187. Springer, Heidelberg (1994)
11. Dagdelen, Ö., Mohassel, P., Venturi, D.: Rate-Limited Secure Function Evaluation: Definitions and Constructions. Cryptology ePrint, Report 201X/XXX
12. Damgård, I., Geisler, M., Krøigaard, M., Nielsen, J.B.: Asynchronous Multiparty Computation: Theory and Implementation. In: Jarecki, S., Tsudik, G. (eds.) PKC 2009. LNCS, vol. 5443, pp. 160–179. Springer, Heidelberg (2009)
13. Freedman, M.J., Ishai, Y., Pinkas, B., Reingold, O.: Keyword Search and Oblivious Pseudorandom Functions. In: Kilian, J. (ed.) TCC 2005. LNCS, vol. 3378, pp. 303–324. Springer, Heidelberg (2005)
14. Gennaro, R., Hazay, C., Sorensen, F.: Automata evaluation and text search protocols with simulation based security. Cryptology ePrint, Report 2010/484
15. Goldreich, O., Micali, S., Wigderson, A.: How to play any mental game. In: ACM Symposium on Theory of Computing, pp. 218–229. ACM (1987)
16. Hazay, C., Lindell, Y.: Efficient oblivious polynomial evaluation with simulation-based security. Technical report, Cryptology ePrint, Report 2009/459
17. Hazay, C., Lindell, Y.: Efficient Protocols for Set Intersection and Pattern Matching with Security Against Malicious and Covert Adversaries. In: Canetti, R. (ed.) TCC 2008. LNCS, vol. 4948, pp. 155–175. Springer, Heidelberg (2008)
18. Hazay, C., Nissim, K.: Efficient Set Operations in the Presence of Malicious Adversaries. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 312–331. Springer, Heidelberg (2010)
19. Hazay, C., Toft, T.: Computationally Secure Pattern Matching in the Presence of Malicious Adversaries. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 195–212. Springer, Heidelberg (2010)
20. Henecka, W., Kogl, S., Sadeghi, A.-R., Schneider, T., Wehrenberg, I.: TASTY: tool for automating secure two-party computations. In: ACM CCS 2007 (2010)
21. Huang, Y., Evans, D., Katz, J., Malka, L.: Faster secure two-party computation using garbled circuits. In: USENIX Security (2011)
22. Ishai, Y., Paskin, A.: Evaluating Branching Programs on Encrypted Data. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 575–594. Springer, Heidelberg (2007)
23. Jarecki, S., Shmatikov, V.: Efficient Two-Party Secure Computation on Committed Inputs. In: Naor, M. (ed.) EUROCRYPT 2007. LNCS, vol. 4515, pp. 97–114. Springer, Heidelberg (2007)
24. Malkhi, D., Nisan, N., Pinkas, B., Sella, Y.: Fairplay—a secure two-party computation system. In: USENIX Security (2004)
25. Mardziel, P., Hicks, M., Katz, J., Srivatsa, M.: Knowledge-oriented secure multiparty computation. In: ACM SIGPLAN – PLAS (June 2012)
26. McGregor, A., Mironov, I., Pitassi, T., Reingold, O., Talwar, K., Vadhan, S.: The limits of two-party differential privacy. ECCV 18:106 (2011)
27. Ogata, W., Kurosawa, K.: Oblivious keyword search. Technical report, Cryptology ePrint, Report 2002/182
28. Troncoso-Pastoriza, J., Katzenbeisser, S., Celik, M.: Privacy preserving error resilient DNA searching through oblivious automata. In: ACM CCS, pp. 519–528 (2007)
29. Yao, A.C.: Protocols for secure computations (extended abstract). In: FOCS, pp. 160–164 (1982)
30. Yao, A.C.: How to generate and exchange secrets (extended abstract). In: FOCS, pp. 162–167 (1986)