

Efficient Server-Aided Secure Two-Party Function Evaluation with Applications to Genomic Computation

Marina Blanton and Fattaneh Bayatbabolghani
Department of Computer Science and Engineering
University of Notre Dame
mblanton@nd.edu, fbayatba@nd.edu

Abstract

Computation based on genomic data is becoming increasingly popular today, be it for medical or other purposes such as ancestry or paternity testing. Non-medical uses of genomic data in a computation often take place in a server-mediated setting where the server offers the ability for joint genomic testing between the users. Undeniably, genomic data is highly sensitive, which in contrast to other biometry types, discloses a plethora of information not only about the data owner, but also about his or her relatives. Thus, there is an urgent need to protect genomic data, especially when it is used in computation for what we call as recreational non-health-related purposes. Towards this goal, in this work we put forward a framework for server-aided secure two-party computation with the security model motivated by genomic applications. One particular security setting that we treat in this work provides stronger security guarantees with respect to malicious users than the traditional malicious model. In particular, we incorporate certified inputs into secure computation based on garbled circuit evaluation to guarantee that a malicious user is unable to modify her inputs in order to learn unauthorized information about the other user's data.

Our solutions are general in the sense that they can be used to securely evaluate arbitrary functions and offer attractive performance compared to the state of the art. We apply the general constructions to three specific types of genomic tests: paternity, genetic compatibility, and ancestry testing and implement the constructions. The results show that all such private tests can be executed within a matter of seconds or less despite the large size of one's genomic data.

1 Introduction

The motivation for this work comes from rapidly expanding availability and use of genomic data in a variety of applications and the need to protect such highly sensitive data from potential abuse. The cost of sequencing one's genome has dramatically decreased in the recent years and is continuing to decrease, which makes such data more readily available for a number of applications. Examples of such applications include:

- personalized medicine, where genomic tests are performed prior to prescribing a drug treatment to ensure its effectiveness;
- paternity testing, which use DNA data to determine whether one individual is the father of another individual;
- genomic compatibility tests, which allow potential or current partners to determine whether their future children are likely to inherit genetic conditions;
- determining ancestry and building genealogical trees by examining DNA data of many individuals and finding relationships among specific individuals.

Genomic tests are increasingly used for medical purposes to ensure the best treatment. A number of services for what we call the “leisure” use of DNA data has flourished as well (examples are [1, 2, 3]) allowing for various forms of comparing DNA data, be it for the purposes of building ancestry trees, genomic compatibility or other.

It is clear that DNA is highly sensitive and needs to be protected from unintended uses. It can be viewed as being even more sensitive than other types of biometry associated with an individual, as not only it allows for unique identification of an individual, but it also allows to learn a plethora of information about the individual such as predisposition to medical conditions and relatives of the individual thus exposing information about others as well. Furthermore, our understanding of genomes is continuously growing and exposure of DNA data now can lead to consequences which we cannot even anticipate today. For that reason, a number of publications that enable genomic computation while preserving privacy of DNA data have appeared in the literature (see, e.g., [14, 12, 26, 9, 11, 10]). Such publications span several types of genomic computation including medical (such as personalized medicine and disease risk computation) and non-medical applications (such as paternity testing).

While protecting privacy of genomic data is important regardless of the reason why computation on such data takes place, in our opinion, it is more difficult for an individual to influence the way medical procedures are conducted than services in which individuals decide to voluntarily participate. That is, if genetic tests are necessary for a patient to determine the most effective treatment and the procedures in place are considered law-compliant, the patient has little possibility for influencing the way the DNA tests are conducted (besides, perhaps, declining to take the test and risking the possibility that the prescribed generic treatment is ineffective for that patient or has severe side effects). With what we consider as “leisure” uses of DNA information, the situation is different. An individual who meets a potential partner through a gene-based matchmaking online service (such as [3]) might be reluctant to share her DNA data with the service (or the partner) for the purpose of genetic compatibility tests. However, if the individual is assured that no sensitive information about her DNA will be revealed to any party throughout the computation other than the intended outcome, she might revisit the decision to participate in such services. Thus, in the rest of this work, when we refer to genomic computation, we focus on applications which are not detrimental to the well-being of an individual and rather consider tests in which individuals might choose to participate.

The first observation we make about such types of genomic computation is that they are normally facilitated by some service or third party. For example, both ancestry and gene-based matchmaking web sites allow participants to interact with each other through the service provider. Such service providers serve as a natural point for aiding the individuals with private computation on their sensitive genomic data. In some prior publications on genomic computation (e.g., [12]), it is assumed that computation such as paternity testing or genetic compatibility is run between a client and a server, while we believe that it is more natural to assume that such computation is carried out by two individuals through some third-party service provider. Thus, in this work we look at private genomic computation in the light of server-mediated setting and utilize the server to lower the cost of the computation for the participants. Throughout this work, we will refer to the participants as Alice (A), Bob (B), and the server (S).

From the security point of view, participants in a protocol that securely evaluates a function are normally assumed to be either semi-honest (also known as honest-but-curious or passive) or malicious (also known as active). In our application domain, we may want to distinguish between different security settings depending on how well Alice and Bob know each other. For example, if Alice and Bob are relatives and would like to know how closely they are related (i.e., how closely their genealogical trees overlap), it would be reasonable to assume that they will not deviate from

the prescribed computation in the attempt to cheat each other, i.e., they can be assumed to be semi-honest. On the other hand, if Alice and Bob meet each other through a matchmaking web site and do know each other well, it is reasonable for them to be cautious and engage in a protocol that ensures security (i.e., correctness and privacy) even in presence of malicious participants. The server can typically be expected not to deviate from its prescribed behavior, as it would lose its reputation and consequently revenue if any attempts at cheating become known. If, however, adding protection against server’s malicious actions is not very costly, it can also be meaningful to assume a stronger security model.

Another important consideration from a security point of view is enforcing correct inputs to be entered in the computation when, for instance, the inputs are certified by some authority. This requirement is outside the traditional security model for secure multi-party computation (even in presence of fully malicious actors), and to the best of our knowledge certified inputs were previously considered only for specific functionalities such as private set intersection [21, 28] or anonymous credentials and certification [18], but not for general secure function evaluation. We bring this up in the context of genomic computation because for certain types of genomic computation it is very easy for one participant to modify his inputs and learn sensitive information about genetic conditions of the other participant. For example, genetic compatibility tests evaluate the possibility of two potential or existing partners to determine the possibility of transmitting to their children a genetic disease. Such possibility is present when both partners are (silent) carriers of that disease (see section 3.1 for more detail). Then if the partners can each separately evaluate their DNA for a fingerprint of specific disease, the joint computation can consist of a simple AND of the bits provided by both parties (for one or more conditions). Now if a malicious participant sets all of his input bits to 1 and the outcome is positive, the participant learns that the other party is a carrier for a specific medical condition (or at least one condition from the set of specific conditions). We thus would like prevent malicious participants from modifying their inputs used in genomic computation in cases such data can be certified by certification authorities such as medical facilities.

The aspect of secure computation related to security properties that we treat in this work is fairness. In particular, it is known that full fairness cannot be achieved in the case of two-party computation in the malicious security model [23], but it becomes possible in the server-aided setting. Fairness has been considered in the server-aided literature in the past [34, 41] and achieving fairness only adds minimal overhead to the solutions in the settings we consider.

Contributions. While we draw motivation from genomic computation, several of our results are general and can be applied to any function. Thus, we categorize our contributions in two main groups: (i) results applicable to general secure function evaluation and (ii) results specific to genomic tests, both of which we consequently describe. All constructions rely on garbled circuit evaluation typically used in the two-party setting (which we consecutively describe in section 3.2), but which we adopt to the three-party computation between the server and two users.

Based on the motivation given above, we consider different adversarial settings, which we present from the simplest and enabling most efficient solutions to the most complex with added security guarantees.

1. Our most efficient solution is designed for the setting where A and B are semi-honest and S can be malicious (as in the ancestry testing scenario). In this setting, the solution consists of a single circuit garbling and single evaluation of the garbled circuit and the need for oblivious transfer is eliminated all together.
2. Our second solution works in the setting where A and B can be malicious, but S is semi-honest (applicable to the paternity test) and achieves fairness for A and B. In this solution,

the combined work for all participants is approximately the same as the combined work of two participants in a two-party protocol in presence of semi-honest participants only.

3. Our last solution strengthens the model of malicious A and B with input certification (applicable to the genomic compatibility test). In more detail, in addition to being able to behave arbitrarily, A and B may maliciously modify their true inputs. To combat this, the function f being evaluated is modified to mark any suitable subset of the inputs as requiring certification. At the time of secure function evaluation, A and B have to prove that the inputs that they enter in the protocol are identical to the values signed by a trusted authority (a medical facility that performs genomic tests in our case). Achieving this involves the use of additional tools such as a signature scheme and zero-knowledge proofs of knowledge (ZKPKs). Handling of the remaining inputs and the rest of the computation is not affected by the shift to a stronger security model.

All of our constructions offer conceptual simplicity and at the same time achieve highly attractive performance. To the best of our knowledge, the strongest of our models which enforces correctness of the inputs have not been treated in the context of general secure multi-party computation and computation based on garbled circuits in particular. Despite the drastic differences in the techniques for garbled circuit evaluation and data certification, we show how they can be integrated by using oblivious transfer as the connecting point or even when oblivious transfer is not used.

Based on the solutions described above, we build implementations of three genetic tests, namely, genetic common ancestry, paternity, and genetic compatibility tests. Each test uses a different security setting. We show through experimental results that each of the implemented tests is efficient with the worst runtime being on the order of a couple of seconds. The performance favorably compares to the state of the art (as detailed in section 7), in some cases achieving orders of magnitude performance improvement over existing solutions.

2 Related Work

Literature on secure two-party (or multi-party) computation is extensive, and its review is beyond the scope of this work. In what follows, we concentrate on two lines of related work, namely, secure server-aided two- or multi-party computation and work that develops privacy-preserving solutions for genetic tests.

Server-aided computation. The closest to our work is that of Herzberg and Shulman [34, 35] that considers two-party secure function evaluation based on garbled circuits in presence of weakly trusted servers that aid in the computation. The solution achieves security and fairness in presence of malicious A and B. The authors also informally discuss (in [35]) extensions to guarantee security in presence of malicious servers or collusion. Compared to that work, our solution in presence of malicious A and B is more efficient in that that [34, 35] require the parties to perform $O(\kappa n)$ signature verifications and engage in $O(\kappa n)$ oblivious transfers, where κ is the security parameter and n is the number of (B's) inputs. The server's work is also larger than in our solution. The use of the server, however, is more constrained in [34, 35] compared to our work.

Kamara et al. [41] assumes a different setting, where a number of parties use the help of a server to reduce computational burden for some of the parties. Using a solution based on garbled circuits, the work achieves work sublinear in the circuit size for a subset of the parties and work polynomial in the circuit size for the remaining parties and the server. Security holds when either the server and another party are malicious or when the server is semi-honest and all but one party are malicious. The model relies on non-colluding adversaries (termed non-cooperating adversaries

in [40, 41]), which even when behave maliciously do not collude with other parties. The work also addresses fairness. While not directly comparable to our result, the work of [41] uses what can be viewed as a more challenging security setting because all of our security settings assume a fixed semi-honest party and thus allow for more efficient constructions.

Carter et al. [22] also uses the aid of a server to reduce the cost of two-party computation based on garbled circuits in the setting where all participants are malicious. One party is assumed to be very weak (such as a mobile phone), while the second participant and the server are more powerful. The solution lifts most of the burden of associated with secure two-party function evaluation in presence of malicious participants from the weak party, but the overhead of the second party and the server is still comparable to the overhead of the participants in regular secure two-party function evaluation based on garbled circuits.

Kolesnikov et al. [44] considers the problem of input consistency in two-party secure function evaluation in presence of malicious players with the aid of a semi-honest server. The goal is to ensure that both A and B enter the same input during multiple interactions, which is enforced with the help of the semi-honest server at low cost. This solution is not suitable for our goal of guaranteeing input correctness as a malicious participant can consistently provide incorrect inputs and by doing so violate privacy of possibly multiple users. Furthermore, there may not be multiple interactions between the same pair of users to enforce input consistency. That work also mentions the possibility of input certification in secure two-party computation, but we are not aware of realizations of this idea.

Genomic computation. There are a number of publications in the literature such as [9, 10, 11] and others that treat the problem of privately computing medical tests for the purposes personalized medicine (with the goal of choosing an optimal medical treatment or drug prescription). Ayday et al. [8] also focus on privacy-preserving systems for storing genomic data by means of homomorphic encryption. Because personalized medicine is outside the scope of this work, we do not further elaborate on such solutions.

To the best of our knowledge, paternity testing in the context of privacy-preserving computation was first considered in [14]. The authors propose privacy-preserving protocols for a number of genetic tests based on Short Tandem Repeats (STRs) representation of genomic data (see section 3.1 for genomic background information). The tests include identity testing, paternity tests with one and two parents, and common ancestry testing on the Y chromosome. The proposed protocols for these tests are based on additively homomorphic public key encryption scheme and are secure in presence of semi-honest participants. Implementation results were not given in [14], but Baldi et al. [12] estimates that the paternity test in [14] is several times slower than that in [12]. We thus compare our paternity test to the performance of an equivalent test in [12].

Baldi et al. [12] concentrate on a different representation of genomic data (in the form of fully-sequenced human genome) and provide solutions for three types of genetic tests: paternity, drug testing for personalized medicine, and genetic compatibility. The solutions use private set intersection as the primary cryptographic building block in the two-party server-client setting. The solutions were implemented and shown to result in attractive runtimes and we compare the performance of our paternity and compatibility tests to the results reported in [12]. We defer the comparison to section 7.

Related to that is the work of De Cristofaro et al. [25] that evaluates the possibility of using smartphones for performing private genetic tests. The tests considered in [25] are paternity and ancestry testing as well as personalized medicine tests. The protocol for the paternity test is the same as in [12] with certain optimizations for the smartphone platform (such as performing pre-processing on a more powerful machine). The ancestry test is performed by sampling genomic data

as using inputs of large size deemed infeasible on a smartphone. The implementation also used private set intersection as the building block. Our implementation, however, can handle inputs of very large sizes at low cost.

Two very recent articles [33, 36] describe mechanisms for private testing for genetic relatives and can detect up to fifth degree cousins. The solutions rely on fuzzy extractors. They encode genomic data in a special form and testing is performed on encoded data. The approach is not comparable to the solutions we put forward in this work as [33, 36] are based on non-interactive computation, which is limited to a specific set of functions.

Although not as closely related to our work as publications that implement specific genetic tests, there are also publications that focus on applications of string matching to DNA testing. Examples include the work of De Cristofaro et al. [27] that provides a secure and efficient protocol that hides the size of the pattern to be searched and its position within the genome. Another example is the work of Katz et al. [42] that applies secure text processing techniques to DNA matching.

3 Preliminaries

3.1 Genomic testing

Genomes represent complete hereditary information of an individual. Information extracted from one’s genome can take different forms. One type is called Single Nucleotide Polymorphisms (SNPs), each of which corresponds to a well known variation in a single nucleotide.¹ Because SNP mutations are often associated with how one develops diseases and responds to treatments, they are commonly used in genetic disease and disorder testing. The same set of SNPs (i.e., nucleotides in the same positions) would be extracted for each individual, but the values associated with each SNP differ from one individual to another. Normally each SNP is referenced by a specific index and its value in a individual is represented as a bit, while representations consisting of 3 values 0, 1, 2 are used as well.

Another type of data extracted from a genome is based on Short Tandem Repeats (STRs). STRs occur when a short region consisting of two or more nucleotides is repeated and the occurrences are adjacent to each other. Unrelated individuals are likely to have a different number of repeats of a given STR sequence in certain regions in their DNA and thus STRs are often used for identity testing or testing between close relatives (such as paternity testing).

Paternity test. This test is normally done based on STRs. STR profile of a person consists of an ordered sequence of N 2-element sets $S = \langle \{x_{1,1}, x_{1,2}\}, \{x_{2,1}, x_{2,2}\}, \dots, \{x_{N,1}, x_{N,2}\} \rangle$, where each value corresponds to the number of repeats of a specific STR sequence at specific locations in the genome. For each STR i , one of $x_{i,1}$ and $x_{i,2}$ is inherited from the mother and the other from the father.

Thus in the paternity test with a single parent, there are two STR profiles $S = \langle \{x_{i,1}, x_{i,2}\} \rangle$ and $S' = \langle \{x'_{i,1}, x'_{i,2}\} \rangle$ corresponding to the child and the contested father, respectively. To determine whether S' corresponds to the father’s child, the test computes whether for each i the child’s set $\{x_{i,1}, x_{i,2}\}$ contains (at least) one element from the contested father’s set $\{x'_{i,1}, x'_{i,2}\}$. In other words, the test corresponds to the computation

$$\bigwedge_{i=1}^N [\{x_{1,i}, x_{2,i}\} \cap \{x'_{1,i}, x'_{2,i}\} \neq \emptyset] = \text{True} \tag{1}$$

¹A nucleotide can be viewed as a simple unit represented by a letter A, C, G, or T.

When testing with both parents is performed, for each STR i one of $x_{i,1}$ and $x_{i,2}$ must appear in the mother’s set and the other in the father’s set. Using both parents’ profiles in the computation increases the accuracy of the test, but even the single parent test has high accuracy for a small number N of well-chosen STRs (e.g., the US CODIS system utilizes $N = 13$, while the European SGM Plus identification method uses $N = 10$).

Genetic compatibility test. While there is a variety of genetic tests that can be used for several purposes, we concentrate on the genetic compatibility test where potential (or existing) partners would like to determine the possibility of transmitting to their children a genetic disease with Mendelian inheritance. In particular, if a specific mutation occurs in one allele² (called minor), it often has no impact on one’s quality of life, but when the mutation occurs in both alleles (called major), the disease manifests itself in severe forms. If both partners silently carry a single mutation, they have a noticeable chance of conceiving a child carrying the major variety. Thus, a genetic compatibility test for a given genetic disease would test for the presence of minor mutations in both partners.

The current practice for screening for most genetic diseases consists of testing one SNP in a specific gene. It is, however, expected that in the future tests for more complex diseases (that involve multiple genes and mutations) will become available. Thus, a genetic disease can be characterized by a set of SNP indices and the corresponding values $(i_1, b_1), \dots, (i_t, b_t)$, where i_j is the SNP index and $b_j \in \{0, 1\}$ is the value it takes. Then if the same values are found in the appropriate SNPs of an individual, we assume that the individual is tested as positive (i.e., the individual is the disease carrier). If both partners test as positive, then the outcome of the genetic compatibility test will be treated as positive and otherwise it is negative.

Ancestry test. There are a number of tests that allow for various forms of ancestry testing, for example, tests using Y-chromosome STRs (applicable to males only), mitochondrial DNA (mtDNA) test on the maternal line, and more general SNP-based tests for common ancestry or one’s genealogy. Many such tests are not standardized and in addition current ancestry and genealogy service providers often use proprietary algorithms. The advantage of STR-based tests is that normally only a relatively small number of STRs are tested, while SNP-based tests often utilize a large number of (or even all available) SNPs, but more distant ancestry can be learned from SNP-based tests. For improved accuracy it is also possible to perform one type of testing after the other. In either case, to determine the most recent common ancestor between two individuals, the markers from the two individuals are compared and their number determines how closely the individuals are related. Certain tests such as determining geographical regions of one’s ancestors normally require genetic data from many individuals.

3.2 Garbled circuit evaluation

The use of garbled circuits allows two parties P_1 and P_2 to securely evaluate a Boolean circuit of their choice. That is, given an arbitrary function $f(x_1, x_2)$ that depends on private inputs x_1 and x_2 of P_1 and P_2 , respectively, the parties first represent it as a Boolean circuit. One party, say P_1 , acts as a circuit generator and creates a garbled representation of the circuit by associating both values of each binary wire with random labels. The other party, say P_2 , acts as a circuit evaluator and evaluates the circuit in its garbled representation without knowing the meaning of the labels that it handles during the evaluation. The output labels can be mapped to their meaning and revealed to either or both parties.

²An allele is one of the alternative versions of a gene at a given location.

In more detail, the basic idea is as follows (here we present only an overview of the approach and refer the reader to, e.g., [47] for technical details and security analysis): For each wire i of the Boolean circuit corresponding to f , the circuit generator creates a pair of randomly chosen labels ℓ_i^0 and ℓ_i^1 (of sufficient length that depends on the security parameter) which map to the values of 0 and 1, respectively, of this wire. Let g be a binary gate that takes two input bits and produces a single bit; also let the input wires to g have indices i and j and let the output wire have index k . Then to create a garbled representation of the gate, the circuit generator produces a truth table containing four entries of the form

$$\text{Enc}_{\ell_i^{b_i}, \ell_j^{b_j}}(\ell_k^{g(b_i, b_j)}).$$

Here $b_i, b_j \in \{0, 1\}$ are input bits into the gate and all entries in the table are randomly permuted. Possession of two input labels $\ell_i^{b_i}$ and $\ell_j^{b_j}$ for any given values of b_i and b_j will allow for recovery of the corresponding output label $\ell_k^{g(b_i, b_j)}$ without revealing anything else. Then upon garbling all gates of the circuit, the circuit generator communicates all garbled gates, to which we collectively refer as a garbled circuit \mathcal{G}_f , to the circuit evaluator together with a *single* label $\ell_i^{b_i}$ for each input wire i according to the input bit b_i . The labels corresponding to the input wires of the circuit generator are simply transmitted to the evaluator, while the labels corresponding to the inputs of the circuit evaluator are communicated to the evaluator by the means of oblivious transfer (detailed below). The knowledge of the input labels and garbled gates allows the circuit evaluator to evaluate the entire circuit in its garbled representation and obtain a label for each output wire representing the output. Then either the circuit generator sends the label pairs (in order) for all output wires to the circuit evaluator, which allows the evaluator to interpret the meaning of the labels and learn the output, or the evaluator sends computed labels to the circuit generator, which in turn allows the circuit generator to learn the result.

The basic approach is secure in presence of semi-honest circuit generator and a malicious evaluator [32] (and the knowledge of valid labels for the output wires implicitly provides a proof that the computation was performed correctly [31]). However, extending the security to the malicious setting (when either party can be malicious) requires additional techniques which substantially degrade performance of the approach.

An important component of garbled circuit evaluation is 1-out-of-2 Oblivious Transfer (OT). It allows the circuit evaluator to obtain wire labels corresponding to its inputs. In particular, in OT the sender (i.e., circuit generator in our case) possesses two strings s_0 and s_1 and the receiver (circuit evaluator) has a bit σ . OT allows the receiver to obtain string s_σ and the sender learns nothing. An oblivious transfer extension allows any number of OTs to be realized with small additional overhead per OT after a constant number of regular more costly OT protocols (the number of which depends on the security parameter). The literature contains many realizations of OT and its extensions, including very recent proposals, but in this work we primarily are interested in OT protocols and OT extensions secure in presence of malicious participants (such as [50, 38, 51] and others).

The fastest currently available approach for circuit generation and evaluation we are aware of is by Bellare et al. [13]. It is compatible with earlier optimizations, most notably the “free XOR” gate technique [45] that allows XOR gates to be processed without cryptographic operations or communication, resulting in virtually no overhead for such gates.

3.3 Signature schemes with protocols and commitment schemes

Our solution that enforces input correctness by means of user input certification relies on additional building blocks, which are signature schemes with protocols, commitment schemes, and

zero-knowledge proofs of knowledge.

From the available signature schemes such as [15, 16] that provide the ability to prove knowledge of a signature on a message without revealing the message itself, the Camenisch-Lysyanskaya solution [15] is of interest to us. It uses public key of the form (n, a, b, c) , where n is an RSA modulus and a, b, c are randomly chosen quadratic residues in \mathbb{Z}_n^* . A signature on a message m is a tuple (e, s, v) such that $v^e \equiv a^b b^s c \pmod{n}$, where e is prime, e and s are randomly chosen according to security parameters, and v is computed to satisfy the equation. A signature can also be issued on a block of messages. To sign a block of t messages m_1, \dots, m_t , the public key needs to be of the form $(n, a_1, \dots, a_t, b, c)$ and the signature is (e, s, v) , where $v^e \equiv a_1^{m_1} \dots a_t^{m_t} b^s c \pmod{n}$ holds.

Given a public verification key (n, a, b, c) , to prove knowledge of a signature (e, s, v) on a secret message m , one forms a commitment $c = \text{Com}(m)$ to m and proves that she possesses a signature on the value committed in c (for the details of which we refer the reader to [15]). The commitment c can consecutively be used to prove additional statements about m in zero knowledge. Similarly, if one wants to prove statements about multiple messages included in a signature, multiple commitments will be formed.

The commitment scheme used in [15] is that of Damgård and Fujisaki [24]. The setup consists of a public key (n, g, h) , where n is an RSA modulus, h is a randomly chosen quadratic residue in \mathbb{Z}_n^* , and g is an element in the group generated by h . Here the modulus n can be the same as or different from the modulus used in the signature scheme. For simplicity, we will assume that the same modulus is used. To produce a commitment to x using the key (n, g, h) , one randomly chooses $r \in \mathbb{Z}_n$ and sets $\text{Com}(x, r) = g^x h^r \pmod{n}$. In the context where the value of r is not essential, we may omit it from the notation and use $\text{Com}(x)$ instead. This commitment scheme is statistically hiding and computationally binding. The values x, r are called the opening of the commitment $\text{Com}(x, r)$.

Zero-knowledge proofs of knowledge (ZKPKs) allow one to prove a particular statement about private values without revealing additional information besides the statement itself. Following [19], we sometimes use notation $PK\{(vars) : statement\}$ to denote a ZKPK of the given statement, where the values appearing in the parentheses are private to the prover and the remaining values used in the statement are known to both the prover and verifier. If the proof is successful, the verifier is convinced of the statement of the proof. For example, $PK\{(\alpha) : y = g_1^\alpha \vee y = g_2^\alpha\}$ denotes that the prover knows the discrete logarithm of y to either the base g_1 or g_2 . Lastly, because a proof of knowledge of a signature is cumbersome to write in this detailed form, we use abbreviation $\text{Sig}(x)$ and $\text{Com}(x)$ in ZKPKs to indicate the knowledge of a signature and commitment, respectively. For example, $PK\{(\alpha) : \text{Sig}(\alpha) \wedge y = \text{Com}(\alpha) \wedge (\alpha = 0 \vee \alpha = 1)\}$ denotes a proof of knowledge of a signature on a bit committed to in y . Because proving the knowledge of a signature on x in [15] requires a commitment to x (which is either computed as part of the proof or may already be available from prior computation), we explicitly include the commitment into all proofs of a signature.

4 Security Model

We formulate security using the standard ideal model/real model for secure multi-party computation, where the view of any adversary in the real protocol execution should be indistinguishable from its view in the ideal model that uses a trusted party to evaluate the function. Because the server does not contribute any input to the computation, it is meaningful to consider that either A or B is honest since the goal is to protect the honest party.

As previously mentioned, we are primarily interested in the setting where the server is semi-honest, but parties A and B may either be semi-honest or fully malicious. Thus, we target security models where S complies with the computation, with the exception of the first setting of semi-honest A and B, where we get security in presence of a malicious server for free. We similarly assume that the server will not collude with users (putting its reputation at risk) or let users affect its operation.

We obtain security settings where (1) A and B can be corrupted by a semi-honest adversary, while S can act on behalf of a fully malicious adversary and (2) A and B can be malicious, but the server is semi-honest. Because we assume that the parties (or the adversaries who corrupt them) do not collude with each other, at any given point of time there might be multiple adversaries, but they are independent of each other. This is similar to the setting considered in [40, 41]. We note that based on the security settings listed above, at most one adversary would be fully malicious. In other words, if in (2) A is malicious, the goal is to protect B who is assumed to not be malicious and S is semi-honest, while in (1) S can be malicious, while A and B are semi-honest. Kamara et al. [41], however, show that in presence of non-cooperating adversaries who corrupt only one party showing security can be reduced to showing that the protocol is secure in presence of semi-honest adversaries only followed by proving for each malicious adversary \mathcal{A}_i that the solution is secure in presence of \mathcal{A}_i when all other parties are honest. This implies that in setting (2) a solution which is secure in presence of malicious A or B will also remain secure when A and B are corrupted by two independent malicious adversaries.

To model fairness, we modify the behavior of the trusted party in the ideal model to send \perp to all parties if any party chooses to abort (note that fairness is only applicable to A and B). We assume that A and B learn the result of evaluation of a predefined function f that takes input x_1 from A and input x_2 from B, and the server learns nothing. Because our primary motivation is genomic computation, we consider single-output functions, i.e., both A and B learn $f(x_1, x_2)$ (but two of our constructions support functions where A's and B's outputs differ and the remaining construction in the present form loses only fairness).

Execution in the real model. The execution of protocol Π in the real model takes place between parties A, B, S and a subset of adversaries \mathcal{A}_A , \mathcal{A}_B , and \mathcal{A}_S who can corrupt the corresponding party. Let \mathcal{A} collectively denote the set of adversaries present in a given protocol execution. A and B receive their respective inputs x_i , a set of random coins r_i , and auxiliary input z_i , while S receives only a set of random coins r_3 and auxiliary input z_3 . All parties also receive security parameter 1^n . Each adversary receives all information that the party it corrupted has and a malicious adversary can also instruct the corresponding corrupted party to behave in a certain way. For each $\mathcal{A}_X \in \mathcal{A}$, let $\text{VIEW}_{\Pi, \mathcal{A}_X}$ denote the view of the adversary \mathcal{A}_X at the end of an execution of Π . Also let $\text{OUT}_{\Pi, \mathcal{A}}^{\text{hon}}$ denote the output of the honest parties (if any) after the same execution of the protocol. Then for each $\mathcal{A}_X \in \mathcal{A}$, we define the partial output of a real-model execution of Π between A, B, S in presence of \mathcal{A} by

$$\text{REAL}_{\Pi, \mathcal{A}_X}(n, x_1, x_2, r_1, r_2, r_3) \stackrel{\text{def}}{=} \text{VIEW}_{\Pi, \mathcal{A}_X} \cup \text{OUT}_{\Pi, \mathcal{A}}^{\text{hon}}.$$

Execution in the ideal model. In the ideal model, all parties interact with a trusted party who evaluates f . Similar to the real model, the execution begins with A and B receiving their respective inputs x_i and each party (in A, B, and S) receiving a set of random coins r_i , auxiliary input z_i , and security parameter 1^n . Each honest (semi-honest) party sends to the trusted party $x'_i = x_i$ and each malicious party can send an arbitrary value x'_i to the trusted party. If x_1 or x_2 is equal to \perp (empty) or if the trusted party receives an abort message, the trusted party returns \perp to all

participants. Otherwise, A and B receive $f(x'_1, x'_2)$. Let $\text{OUT}_{f, \mathcal{A}}^{\text{hon}}$ denote the output returned by the trusted party to the honest parties and let $\text{OUT}_{f, \mathcal{A}_X}$ denote the output that corrupted party $\mathcal{A}_X \in \mathcal{A}$ produces based on an arbitrary function of its view. For each $\mathcal{A}_X \in \mathcal{A}$, the partial output of an ideal-model execution of f between A, B, S in presence of \mathcal{A} is denoted by

$$\text{IDEAL}_{f, \mathcal{A}_X}(n, x_1, x_2, r_1, r_2, r_3) \stackrel{\text{def}}{=} \text{OUT}_{f, \mathcal{A}_X} \cup \text{OUT}_{f, \mathcal{A}}^{\text{hon}}.$$

Definition 1 (Security) *A three-party protocol Π between A, B, and S securely computes f if for all sets of probabilistic polynomial time (PPT) adversaries \mathcal{A} in the real model, for all x_i, z_i , and $n \in \mathbb{Z}$, there exists a PPT transformation \mathcal{S}_X for each $\mathcal{A}_X \in \mathcal{A}$ such that*

$$\text{REAL}_{\Pi, \mathcal{A}_X}(n, x_1, x_2, r_1, r_2, r_3) \stackrel{c}{\approx} \text{IDEAL}_{f, \mathcal{S}_X}(n, x_1, x_2, r_1, r_2, r_3)$$

where each r_i is chosen uniformly at random and $\stackrel{c}{\approx}$ denotes computational indistinguishability.

To model the setting where some of the inputs of A and/or B are certified, we augment the function f to be executed with the specification of what inputs are to be certified and two additional inputs y_1 and y_2 that provide certification for A and B's inputs, respectively. Then in the ideal model execution, the trusted party will be charged with additionally receiving y_i 's. If the trusted party does not receive all inputs or if upon receiving all inputs some inputs requiring certification did not verify, it sends \perp to all parties. In the real model execution, verification of certified inputs is built into Π and besides using two additional inputs y_1 and y_2 the specification of the execution in the real model remains unchanged.

Definition 2 (Security with certified inputs) *A three-party protocol Π between A, B, and S securely computes f if for all sets of PPT adversaries \mathcal{A} in the real model, for all x_i, y_i, z_i , and $n \in \mathbb{Z}$, there exists a PPT transformation \mathcal{S}_X for each $\mathcal{A}_X \in \mathcal{A}$ such that*

$$\text{REAL}_{\Pi, \mathcal{A}_X}(n, x_1, x_2, y_1, y_2, r_1, r_2, r_3) \stackrel{c}{\approx} \text{IDEAL}_{f, \mathcal{S}_X}(n, x_1, x_2, y_1, y_2, r_1, r_2, r_3)$$

where each r_i is chosen uniformly at random.

5 Server-Aided Two-Party Computation

In this section we detail our solutions for server-aided two party computation based on garbled circuits. The current description is general and can be applied to any function f . In the consecutive section 6 we describe how the solutions presented in this section can be applied to the computation of genomic tests to result in fast performance.

5.1 Semi-honest A and B, malicious S

Our first security setting is where A and B are semi-honest and S can be malicious. The main intuition behind the solution is that when A and B can be assumed to be semi-honest and a solution based on garbled circuit evaluation is used, we will charge S with the task of evaluating a garbled circuit. That is, security is maintained in presence of malicious server because garbled circuit evaluation techniques are secure in presence of a malicious evaluator. Next, we notice that if A and B jointly form garbled representation of the circuit for the function f they would like to evaluate, both of them can have access to the pairs of labels (ℓ_i^0, ℓ_i^1) corresponding to the input wires. Thus, they can simply send the appropriate label ℓ_i^b to S for evaluation purposes for their value of

the input bit b for each input wire. This eliminates the need for oblivious transfer and results in a solution that outperforms a two-party protocol in presence of only semi-honest participants.

Modern techniques for garbling circuits rely on the “free XOR” technique, which removes the need for cryptographic operations and communication for XOR gates resulting substantial savings. To use the technique, the label pairs for each circuit wire are required to have a certain relationship, namely that $\ell_i^1 = \ell_i^0 \oplus \Delta$, where Δ is a global random value with the last bit set to 1, which is used for all wires of the circuit. Thus, when generating labels for circuit wires, the circuit generator chooses ℓ_i^0 at random and computes ℓ_i^1 as $\ell_i^0 \oplus \Delta$. This is of importance to us because in the current solution the circuit is being garbled jointly by two parties. This means that if one party generates labels for some, but not all wires, communicating the chosen labels to the other party involves sending only one label ℓ_i^0 per wire instead of the pair (ℓ_i^0, ℓ_i^1) . This reduces the communication for that portion of the parties’ interaction by half.

A more detailed description of the solution, which we denote as Protocol 1, is given below. In what follows, let m denote the total number of wires in a circuit (including input and output wires), wires $1, \dots, t_1$ correspond to A’s input, wires $t_1 + 1, \dots, t_1 + t_2$ correspond to B’s input, and the last t_3 wires $m - t_3 + 1, \dots, m$ correspond to the output wires. We also use κ to denote security parameter (for symmetric key cryptography). Notation $a \stackrel{R}{\leftarrow} U$ means that the value of a is chosen uniformly at random from the set U .

Input: A has private input x_1 , B has private input x_2 , and S has no private input.

Output: A and B learn $f(x_1, x_2)$, S learns nothing.

Protocol 1:

1. A and B jointly choose $\delta \stackrel{R}{\leftarrow} \{0, 1\}^{\kappa-1}$ and m labels $\ell_i^0 \stackrel{R}{\leftarrow} \{0, 1\}^\kappa$, then set $\Delta = \delta || 1$ and $\ell_i^1 = \ell_i^0 \oplus \Delta$ for each $i \in [1, m]$. A and B jointly garble the gates to produce a garbled circuit \mathcal{G}_f for f and send \mathcal{G}_f to S.
2. For each $i \in [1, t_1]$, A locates the i th bit b_i of her input and sends to S the label $\ell_i^{b_i}$ of the corresponding wire i in the garbled circuit.
3. Similarly, for each bit $j \in [1, t_2]$, B locates the j th bit b_j of his input and sends to S the label $\ell_{i+t_1}^{b_j}$ of the corresponding wire $i + t_1$ in the garbled circuit.
4. S evaluates the circuit on the received inputs and returns to B the computed label ℓ_i^b for each output wire $i \in [m - t_3 + 1, m]$. B forwards all received information to A.
5. For each ℓ_i^b returned by S ($i \in [m - t_3 + 1, m]$), A and B do the following: if $\ell_i^b = \ell_i^0$, set $(i - m + t_3)$ th bit of the output to 0, if $\ell_i^b = \ell_i^1$, set $(i - m + t_3)$ th bit of the output to 1, otherwise abort.

In this solution, the combined work of A and B is linear in the size of a circuit for f . The work, however, can be distributed in an arbitrary manner between the parties. Besides equally splitting the work of circuit garbling between the parties, an alternative possibility is to let the weaker party (e.g., a mobile phone user) to do work sublinear in the circuit size. Let A be a weak client, who delegates as much work as possible to B. Then A can create t_1 label pairs for her input only and send them to B who produces the rest of the garbled circuit. Upon completion of the result, A learns the output from B (i.e., there is not even the need for A to know any labels for the output wires). Thus, the work and communication of the weaker client is only linear in the input and output sizes. We note that, as noted above, using the variant of circuit garbling that supports free

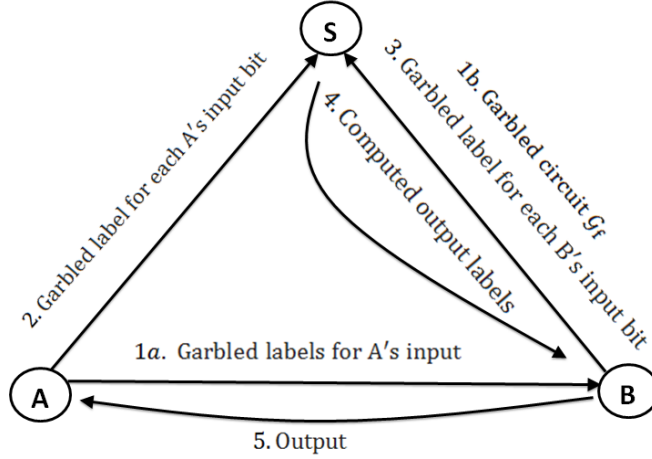


Figure 1: Illustration of Protocol 1 with weak A (who contributes only garbled labels for her input wires to the computation).

XOR gates, A needs to send only t_1 input labels to B instead of t_1 label pairs, which additionally reduces A's communication in half.

This version of the protocol (with weak A) is depicted in figure 1. Security of this solution can be stated as follows:

Theorem 1 *Protocol 1 above fairly and securely evaluates a circuit for function f in presence of semi-honest A and B and malicious S.*

The proof can be found in Appendix A.

Before we conclude this section, we would like to comment on the possibility of executing our solution when users A and B are not simultaneously or continuously online (but the server is available) and analyze the solution with respect of how the computation can proceed asynchronously.

1. One user (A or B) comes online and initiates the protocol (executes steps 1a and 2).
2. The other user comes online and participates in the computation (executes steps 1b and 3–4).

At this point the answer is ready and known to the second user. If the first user is to learn the answer as well, it will need to either come online or be notified by the second user through a different private channel (step 5). Thus, we see that interaction is minimal and the solution is suitable for asynchronous environments.

5.2 Semi-honest S, malicious A and B

In order to maintain efficiency of the previous solution by avoiding the cost of oblivious transfer, we might want to preserve the high-level structure of the computation in the first solution. Now, however, because A and B can be malicious, neither of them can rely on the other party in garbling the circuit correctly. To address this, each of A and B may garble their own circuit for f , send it to S, and S will be in charge of evaluating both of them and performing consistency check on the results (without learning the output). With this solution, A would create label pairs for her input bits/wires for both garbled circuits and communicate one set of pairs to B who uses them in constructing his circuit. What this accomplishes is that now A can directly send to S the labels corresponding to her input bits for circuit evaluation for both circuits. B performs identical operations. There is still no need to perform OT, but two security issues arise: (1) we have to

enforce that A and B provide consistent inputs into both of the circuits and (2) regardless of whether the parties learn the output (e.g., whether the computation is aborted or not), a malicious party can learn one bit of information about the other party’s input (by constructing a circuit that does not correspond to f) [37]. While the first issue can be inexpensively addressed using the solution of [44] (which works in presence of malicious users and semi-honest server), the second issue will still stand with this structure of the computation.

Instead of allowing for (1-bit) information leakage about private inputs, we change the way the computation takes place. If we now let the server garble the circuit and each of the remaining parties evaluate a copy of it, the need for OT (for both A and B’s inputs) arises. We, however, were able to eliminate the use of OT for one of A and B and construct a solution that has about the same cost as a single two-party solution in the semi-honest model. At the high-level, it proceeds as follows: A creates garbled label pairs (ℓ_i^0, ℓ_i^1) for the wires corresponding to her inputs only and sends them to S. S uses the pairs to construct a garbled circuit for f and sends it to B. S and B engage in OT, at the end of which B learns labels corresponding to his input bits. Also, A sends to B the labels corresponding to her input bits, which allows B to evaluate the circuit. We note that because A may act maliciously, she might send to B incorrect labels, which will result in the inability of B to evaluate the circuit. This, however, is equivalent to A aborting the protocol. In either case, neither A nor B learn any output and the solution achieves fairness. Similarly, if B does not follow the computation for circuit evaluation, neither party learns the output.

The next issue that needs to be addressed is that of fairly learning the output. We note that S cannot simply send the label pairs for the output wires to A and B as this would allow B to learn the output and deny A of this knowledge. Instead, upon completion of garbled circuit evaluation, B sends the computed labels to A. With the help of S, A verifies that the labels A possesses are indeed valid labels for the output wires without learning the meaning of the output. Once A is satisfied, she notifies S who sends the label pairs to A and B, both of whom can interpret and learn the result. We note that malicious A can report failure to S even if verification of the validity of the output labels received from B was successful. Once again, this is equivalent to A aborting the protocol, in which case neither party learns the output and fairness is maintained.

Our solution, denoted as Protocol 2, is given in more detail next. As with Protocol 1, we can let A generate and communicate one random label ℓ_i^0 for her input bits, while the second label ℓ_i^1 can be computed by A and S using their shared knowledge of Δ . This results in significant reduction of A’s communication.

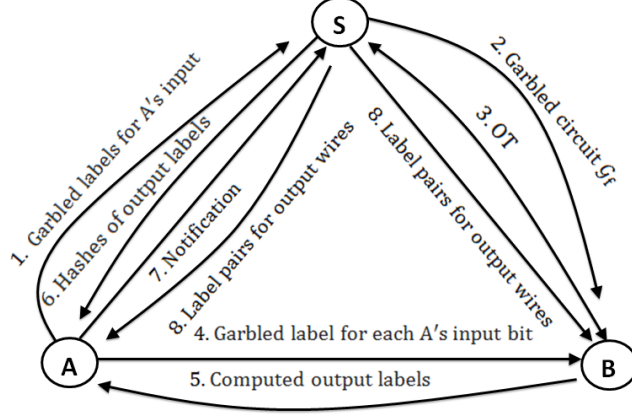


Figure 2: Illustration of Protocol 2.

Input: A has private input x_1 , B has private input x_2 , and S has no private input.

Output: A and B learn $f(x_1, x_2)$, S learns nothing.

Protocol 2:

1. A generates t_1 labels $\ell_i^0 \xleftarrow{R} \{0, 1\}^\kappa$ for each circuit wire $i \in [1, t_1]$ corresponding to A's input and sends them to S. S chooses Δ and sends it to A.
2. S checks that each label ℓ_i^0 received from A is of the correct bitlength and computes the corresponding labels $\ell_i^1 = \ell_i^0 \oplus \Delta$. S then creates random labels for the remaining wires in the same way and construct garbled gates \mathcal{G}_f . S sends \mathcal{G}_f to B.
3. S and B engage in t_2 instances of 1-out-of-2 OT, where S assumes the role of the sender and uses t_2 label pairs $(\ell_{t_1+i}^0, \ell_{t_1+i}^1)$ for $i \in [1, t_2]$ corresponding to B's input wires as its input and B assumes the role of the receiver and uses his t_2 input bits b_i as the input into the protocol. As the result of the interaction, B learns garbled labels $\ell_{t_1+i}^{b_i}$ for $i \in [1, t_2]$.
4. A sends the labels $\ell_i^{b_i}$ corresponding to her input bits b_i to B for $i \in [1, t_1]$, where $\ell_i^1 = \ell_i^0 \oplus \Delta$ for any $b_i = 1$.
5. Upon the receipt of the labels for his own and A's input, B evaluates the circuit, learns the output labels $\ell_i^{b_i}$ for $i \in [m - t_3 + 1, m]$ and sends them labels to A.
6. A asks and receives from S output verification information constructed as follows: For each output wire i , S computes $H(\ell_i^0), H(\ell_i^1)$, where $H : \{0, 1\}^* \rightarrow \{0, 1\}^\kappa$ is a collision-resistant hash function, randomly permutes the tuple, and sends it to A.
7. For each label ℓ_i received from B in step 5, A computes $H(\ell_i)$ and checks where the computed value appear among $H(\ell_i^b), H(\ell_i^{1-b})$ received from S in step 6. If the check succeeds for all output wires, A notifies S of successful completion of circuit evaluation and aborts otherwise.
8. Upon receiving confirmation of success from A, S sends the pairs (ℓ_i^0, ℓ_i^1) for all output wires i to A and B, who interpret the result and learn the output.

The protocol steps are also illustrated in figure 2.

We show security of this solution in a hybrid model where the parties are assumed to be given

access to a trusted entity computing oblivious transfer. We obtain:

Theorem 2 *Protocol 2 above fairly and securely evaluates a circuit for function f in presence of malicious A or B and semi-honest S in the hybrid model with ideal implementation of OT and where H is a collision-resistant hash function.*

The proof can be found in Appendix A.

As with the previous solution, we would like to analyze our protocol with respect to the number of asynchronous interactions needed when A and B are not continuously online, which we show below.

1. One user (A or B) comes online and initiates the protocol (executes steps 1 and 4).
2. The other user comes online and participates in the computation (executes steps 2–3 and 5–6).
3. The first user comes online and performs verification (executes steps 7–8).

At this point the answer is ready and known to the first user. If the second user is to learn the answer as well, it will need to come online the second time (step 8).

5.3 Semi-honest S , malicious A and B with input certification

We next consider the setting with enhanced security guarantees in which malicious A and B are enforced to provide correct inputs in the computation. This enforcement is performed by requiring A and B certify their inputs prior to protocol execution and prove the existence of certification on the inputs they enter in the computation.

The basic structure of our solution in this stronger security model remains the same as in Protocol 2, but we extend it with a novel mechanism for obviously verifying correctness of the inputs. The intricate part of this problem is that signature schemes use public-key operations, while garbled circuit evaluation deals with randomly generated labels and symmetric key operations. In what follows, we describe the intuition behind our solution followed by more detailed explanation.

Suppose that the party whose inputs are to be verified participates in an OT protocol on her inputs as part of garbled circuit evaluation (i.e., the party is the circuit evaluator and acts as the receiver in the OT). Then if we use the variant of OT known as committed oblivious transfer (COT) (which is also called verifiable OT in some literature), the party will submit commitments to the bits of her input as part of OT computation and these commitments can be naturally tied to the values signed by a third party authority by means of zero-knowledge proofs (i.e., without revealing any information other than equality of the signed values and the values used in the commitments). Several COT schemes that we examined (such as in [39, 43]), however, had disadvantages in their performance and/or complex setup assumptions (such as requiring the sender and receiver to hold shares of the decryption key for a homomorphic public-key encryption scheme). We thus integrate input certification with a specific instantiation of conventional OT protocol by Naor and Pinkas [50].

Before we proceed with further description, we discuss the choice of the signature scheme and the way knowledge of a signature is proved. Between the main two candidates of signature schemes with protocols [15] and [16], we chose the solution from [15] because it uses an RSA modulus. In application like ours, zero-knowledge statements are to be proved across different groups. This requires the use of statistically-hiding zero-knowledge proofs that connect two different groups through a setting in which the Strong RSA assumption (or, more generally, the difficulty of eth root extraction) holds [17, 24, 30]. Thus, the public key of the third party certification authority can be conveniently used as the common setup for other interaction between the prover and verifier.

Input: Sender S has two strings m_0 and m_1 , receiver R has a bit σ . Common input consists of prime p , generator \hat{g} of subgroup of \mathbb{Z}_p^* of prime order q , and a random element C from the group generated by \hat{g} (chosen by S).

Output: R learns m_σ and S learns nothing.

OT Protocol:

1. S chooses random $r \in \mathbb{Z}_q$ and computes C^r and \hat{g}^r .
2. R chooses $k \in \mathbb{Z}_q^*$, sets public keys $PK_\sigma = \hat{g}^k$ and $PK_{1-\sigma} = C/PK_\sigma$, and sends PK_0 to S.
3. After receiving PK_0 , S computes $(PK_0)^r$ and $(PK_1)^r = C^r/(PK_0)^r$. S sends to R \hat{g}^r and two encryptions $H((PK_0)^r, 0) \oplus m_0$ and $H((PK_1)^r, 1) \oplus m_1$, where H is a hash function (modeled as a random oracle).
4. R computes $H((\hat{g}^r)^k) = H((PK_\sigma)^r)$ and uses it to recover m_σ .

Figure 3: 1-out-of-2 Oblivious Transfer of [50].

This has important implications on the use of such solutions in practice. (If multiple signatures are issued by multiple authorities, i.e., medical facilities in our application, one of the available public keys can be used to instantiate the common setup.)

Recall that in Protocol 2 B obtains the labels corresponding to his input from S via oblivious transfer, while A knows all label pairs corresponding to her input wires and simply sends the appropriate labels to B. Now both of them will have to prove to S that the inputs they enter in the protocol have been certified by a certain third party authority. For simplicity, in what follows we will assume that all of A's and B's inputs are to be verified. (If this is not the case and only a subset of the inputs should be verified, the computation associated with input verification described below is simply omitted for some of the input bits.) Let us start with the verification mechanism for B, after which we treat the case of A.

B engages in the Naor-Pinkas OT in the role of the receiver. The details of the OT protocol are given in Figure 3. As part of OT, B forms two keys PK_0 and PK_1 , where PK_σ is the key that will be used to recover m_σ . Thus, if we want to enforce that σ corresponds to the bit for which B has a signature from a certification authority, B must prove that he knows the discrete logarithm of PK_σ where σ is the signed bit. More formally, the statement B has to prove in zero knowledge is as follows:

$$PK\{(\sigma, \beta) : \text{Sig}(\sigma) \wedge y = \text{Com}(\sigma) \wedge ((\sigma = 0 \wedge PK_0 = \hat{g}^\beta) \vee (\sigma = 1 \wedge PK_1 = \hat{g}^\beta))\}.$$

In other words, B has a signature on 0 and knows the discrete logarithm of PK_0 to the base \hat{g} (i.e., constructed PK_0 as \hat{g}^k) or B has a signature on 1 and knows the discrete logarithm of PK_1 to the same base. Using a technically more precise PK statement for showing that σ is 0 or 1 would result in the PK statement above be re-written as:

$$PK\{(\sigma, \alpha, \beta) : \text{Sig}(\sigma) \wedge y = \text{Com}(\sigma, \alpha) = g^\sigma h^\alpha \wedge ((y = h^\alpha \wedge PK_0 = \hat{g}^\beta) \vee (y/g = h^\alpha \wedge PK_1 = \hat{g}^\beta))\}.$$

We note that it is known how to realize this statement as a ZKPK as it uses only conjunction and disjunction of discrete logarithm based sub-statements (see, e.g., [20]). Executing this ZKPK would allow S to verify B's input for a particular input wire if B has a signature on a bit. In practice, however, a signature is expected to be on messages from a larger space than $\{0, 1\}$ and

thus a single signature will need to be used to provide inputs for several input wires in the circuit. This can be accomplished by, in addition to using a commitment on the signed message, creating commitments to the individual bits and showing that they correspond to the binary representation of the signed message. Then the commitments to the individual bits of the message are linked to the keys generated in each instance of the OT. More formally, the ZKPK statement used in the protocol would for a t -bit signed value become:

$$PK\{(\sigma, \sigma_1, \dots, \sigma_t, \alpha, \alpha_1, \dots, \alpha_t) : \text{Sig}(\sigma) \wedge y = g^\sigma h^\alpha \wedge$$

$$y_1 = g^{\sigma_1} h^{\alpha_1} \wedge \dots \wedge y_t = g^{\sigma_t} h^{\alpha_t} \wedge \sigma = \sum_{i=1}^t 2^{i-1} \sigma_i\} \quad (2)$$

$$PK\{(\sigma_i, \alpha_i, \beta_i) : y_i = g^{\sigma_i} h^{\alpha_i} \wedge ((y_i = h^{\alpha_i} \wedge PK_0^{(i)} = \hat{g}^{\beta_i}) \vee (y_i/g = h^{\alpha_i} \wedge PK_1^{(i)} = \hat{g}^{\beta_i}))\}. \quad (3)$$

Here notation $PK_0^{(i)}$ and $PK_1^{(i)}$ indicates the public keys used in the i th instance of Naor-Pinkas OT. Note that it was shown in [20] how to prove that discrete logarithms that satisfy a particular linear equation.

Furthermore, it is likely that signatures will contain multiple messages (e.g., the name of a genetic disease and the outcome of testing for it). In those cases, multiple messages from a single signature can be used as inputs into the garbled circuit or, depending on the function f , there may need to be other arrangements. For instance, one message can be used to provide inputs into the circuit and another be opened or partially open (i.e., some statement is proved about it in zero knowledge). We note that it is not difficult to generalize the statement in equations 2 and 3 to cover such cases.

We now can proceed with the description of the mechanism for verifying A's inputs. Recall that for each bit i of her input, A has label pairs (ℓ_i^0, ℓ_i^1) and later she sends to B the label $\ell_i^{b_i}$ corresponding to her input bit b_i . As before, consider first the case when A holds a signature on a single bit. To prove that the label $\ell_i^{b_i}$ sent to B corresponds to the bit for which she possesses a signature, we have A commit to the label $\ell_i^{b_i}$ and prove to S that either the commitment is to ℓ_i^0 and she has a signature on 0 or the commitment is to ℓ_i^1 and she has a signature on 1. Let the commitment be $c_i = \text{Com}(\ell_i^{b_i}, r_i)$. Then if verification of the ZKPKs for each input bit was successful, S forwards each c_i to B together with the garbled circuit. Now when A is to send her input label $\ell_i^{b_i}$ to B, she is also now required to open the commitment c_i by sending r_i to B. B will proceed with circuit evaluation only if $c_i = g^{\ell_i^{b_i}} h^{r_i}$ for each bit i of A's input, where $\ell_i^{b_i}$ and r_i are the values that he received from A.

More formally, the statement A proves to S in zero knowledge is:

$$PK\{(\sigma, \alpha, \beta, \gamma) : \text{Sig}(\sigma) \wedge y = g^\sigma h^\alpha \wedge z = g^\beta h^\gamma \wedge$$

$$((y = h^\alpha \wedge z/g^{\ell_i^0} = h^\gamma) \vee (y/g = h^\alpha \wedge z/g^{\ell_i^1} = h^\gamma))\}.$$

Similar to the case of B's input verification, this ZKPK can be generalized to use a single signature with multiple bits input into the circuit. More precisely, the statement in equation 2 remains unchanged, while the second statement becomes:

$$PK\{(\sigma_i, \alpha_i, \beta_i, \gamma_i) : y_i = g^{\sigma_i} h^{\alpha_i} \wedge z_i = g^{\beta_i} h^{\gamma_i} \wedge$$

$$((y_i = h^{\alpha_i} \wedge z_i/g^{\ell_i^0} = h^{\gamma_i}) \vee (y_i/g = h^{\alpha_i} \wedge z_i/g^{\ell_i^1} = h^{\gamma_i}))\}. \quad (4)$$

We summarize the overall solution below, where for simplicity of presentation we assume that all input bits of A (resp. B) are certified and signed as a single message.

Input: A has private input x_1 and signature $\text{Sig}(x_1)$, B has private input x_2 and $\text{Sig}(x_2)$, and S has no private input.

Output: A and B learn $f(x_1, x_2)$, S learns nothing.

Protocol 3:

1. S chooses Δ and sends to A. A chooses t_1 labels $\ell_i^0 \xleftarrow{R} \{0, 1\}^\kappa$. For each bit b_i of her input, A computes commitments $c_i = \text{Com}(b_i, r_i)$ and $c'_i = \text{Com}(\ell_i^{b_i}, r'_i)$ using fresh randomness r_i and r'_i , where $\ell_i^1 = \ell_i^0 \oplus \Delta$. A sends to S $\text{Sig}(x_1)$, labels ℓ_i^0 and commitments c_i, c'_i for each $i \in [1, t_1]$. A proves in ZK the statement in equation 2 using private inputs $x_1, b_1, \dots, b_{t_1}, r_1, \dots, r_{t_1}$. For each $i \in [1, t_1]$, A also proves in ZK the statement in equation 4 using private inputs $b_i, r_i, \ell_i^{b_i}, r'_i$.
2. S performs the same computation as in Protocol 2 and also sends to B A's commitments c'_i for $i \in [1, t_1]$.
3. S and B engage in t_2 instances of 1-out-of-2 OT as in Protocol 2 together with verification of B's input. Before B can learn labels $\ell_{t_1+i}^{b_i}$, B forms t_2 commitments $c''_i = \text{Com}(b_i, r''_i)$ using fresh randomness r''_i and proves in ZK the statements in equations 2 and 3 using private input $x_2, b_1, \dots, b_{t_2}, r''_1, \dots, r''_{t_2}$ and b_i, r''_i, k_i , respectively. Here k_i denotes the value chosen during step 2 of the i th instance of the OT protocol.
4. A opens commitments c'_i by sending to B pairs $(\ell_i^{b_i}, r'_i)$ for $i \in [1, t_1]$. B checks whether $\text{Com}(\ell_i, r'_i) = c'_i$ for each i and aborts if at least one check fails.
5. The remaining steps are the same as in Protocol 2.

We state the security of this solution as follows:

Theorem 3 *Protocol 3 fairly and securely evaluates a circuit for function f in presence of malicious A or B and semi-honest S in the hybrid model with ideal implementation of OT and where H is a collision-resistant hash function and inputs of A and B are verified according to definition 2.*

Because the structure of the computation in Protocol 3 is the same as in Protocol 2 and primarily only ZKPKs have been added (that have corresponding simulators for the simulation in the ideal model), we defer the proof to the full version of this work.

Before we conclude this section, we would like to comment on the possibility of using OT extensions in combination with certified inputs. The first thing to notice is that when only a subset of the input bits is to be verified, OT and OT extensions for the remaining input bits can be used as before. Second, if there is a computational benefit to using an OT extension instead of individual instances of OT, the benefit of an OT extension is not going to be as pronounced as in the regular case with no input certification. OT extensions allow the number of public key operations to be bounded by the security parameter and be independent of the number of input bits, while with certified inputs the number of public key operations is inevitably linear in the number of input bits being verified. Furthermore, applicability of each individual OT extension mechanism to the case of certified inputs will likely need to be considered on a case by case basis and we leave this problem as a direction for future research.

At the end, we are going to mention users A and B don't need to be simultaneously online to perform the computation and the interaction is the same as in protocol 2.

6 Private Genomic Computation

For all types of genomic computation that follow we assume that A has information extracted from her genome, which she privately stores. Similarly, B possesses information associated with his genome. A and B may enter part or all of their data into the computation and they may also compute a function of their individual data, which will be used as the input into the joint computation.

6.1 Ancestry test

This test would often be invoked when A and B already know to be related or have reasons to believe to be related. Under such circumstances, they are unlikely to try to cheat each other. For that reason, we use the solution with semi-honest A and B to realize this test. Because SNP-based tests are most general and can provide information about recent as well as distant ancestry, we build a circuit that takes a large number of SNPs from two individuals and counts the number of positions which take the same values for both individuals. The computed value then can be compared to a number of thresholds to determine the closest generation in which the two individuals have the same ancestor.

To compute the number of SNPs which are equal (or, equivalently, differ) in the DNA of two individuals, the circuit first proceeds by XORing two binary input vectors from A and B (recall that the value of each SNP is a bit) and then counts the number of bits that differ in a hierarchical manner. That is, in the first round of additions, every two adjacent bits are added and the result is a 2-bit integer. In the second round of additions, every two adjacent results from the first round are added resulting in 3-bit sums. This process continues in $\lceil \log_2 t \rceil$ rounds of additions, where t is the size of A's and B's input, and the last round performs only a single addition. As mentioned earlier, the result can be interpreted by performing a number of comparisons at the end, but the cost of final comparisons is insignificant compared to the remaining size of the circuit.

6.2 Paternity test

We assess that the security setting with malicious users A and B is the most suitable for running paternity tests. That is, the participants may be inclined to tamper with the computation to influence the result of the computation. It is, however, difficult to learn the other party's genetic information by modifying one's input into the function. In particular, recall from equation 1 that the output of a paternity test is a single bit, which indicates whether the exact match was found. Then if a malicious participant engages in the computation with the same victim multiple times and modifies the input in the attempt to discover the victim's genomic data, the single bit output does not help the attacker to learn how his inputs are be modified to be closer to the victim's input. The situation is different when the output of the computation reveals information about the distance between the inputs of A and B, but we do not consider such computation in this work. Thus, we do not use input certification for paternity tests.

This test would normally be run between an individual and a contested father of that individual according to the computation in equation 1. We thus implement the computation in equation 1 using a Boolean circuit. For each i , the circuit XORs the vectors $\langle x_{i,1}, x_{i,2}, x_{i,1}, x_{i,2} \rangle$ and $\langle x'_{i,1}, x'_{i,1}, x'_{i,2}, x'_{i,2} \rangle$ and compares each of the four value in the resulting vector to 0. The (in)equality to 0 testing is performed using $k - 1$ OR gates, where k is the bitlength of all $x_{i,j}$'s and $x'_{i,j}$'s. Finally, we compute the AND of the results of the 4 equality tests, OR the resulting bits across i 's, and output the complement of the computed bit.

6.3 Genetic compatibility test

When A and B want to perform a genetic compatibility test, we assume that they want to evaluate the possibility of their children inheriting at least one recessive genetic disease. Thus, we assume that A and B agree on a list of genetic diseases that they would like to be included in the test (this list can be standard, e.g., suggested by S or recommended by a medical association). Note that performing a test for a specific genetic disease is only meaningful if both parties wish to be tested for it, and thus we assume that A and B can reconcile the differences in their lists.

To maximize privacy, we construct the function f to be as conservative as possible. In particular, given a list of genetic diseases L , A and B run a compatibility test for each disease $D \in L$, and if at least one test resulted in a positive outcome, the function will output 1, and otherwise it will output 0. That is, the function can be interpreted as producing 1 if A and B's children have a chance of inheriting the major variety for at least one of the tested diseases; and producing 0 means that their children will not inherit the major variety for any of the diseases in L . Evaluating this function can be viewed as the first step in A and B's interaction. If the output was 1, they may jointly decide to run more specific computation to determine the responsible disease or diseases themselves.

The above means that for each $D \in L$, A can locally run the test to determine whether she is a carrier of D . Similarly, B performs the same test on his own data. Thus, A's and B's input into f will consist of $|L|$ bits each and the result is 1 iff $\exists i$ such that A's and B's i th input bits are both 1. This computation can be realized as a very simple circuit consisting of $|L|$ AND gates and $|L| - 1$ OR gates.

The next thing to notice is that it is easy for a malicious A or B to learn sensitive genetic information about the other party by entering certain inputs into the computation. In particular, if a malicious participant sets all of his input bits to 1, he will be able to learn whether the other party is the carrier of least one of the genetic diseases on the list. This poses substantial privacy concerns, particularly for matchmaking services that routinely run genetic compatibility tests between many individuals. Thus, we require that A and B certify the results of testing for each genetic disease on the list (which can be done by a medical facility³) and enter certified inputs into the computation. This means that the server-aided solution with certified inputs will be used to securely perform joint computation.

For each disease $D \in L$, the signature will need to include the name of the disease D as well as the test outcome σ , which we assume is a bit. Then if we want to achieve efficiency of the computation, the disease names will not be input into the circuit, but instead S will verify that A's signature used for a particular input wire includes the same disease name as B's signature used for an equivalent input wire. A simple way to achieve this is to reveal list L to S and reveal the name of the disease including in each signature (without revealing the signature itself). If we assume that each issued signature is on the tuple (D, σ) , i.e., the signature was produced as $v^e \equiv a_1^D a_2^\sigma b^s c$, all that is needed is to adjust the value used in the ZKPK of the signature by $\frac{1}{a_2^D}$ by both the sender and the verifier for each $D \in L$ (we refer the reader to [15] for details). S will need to check that all conditions appear in the same order among A's and B's inputs (i.e., the sequences of diseases are identical) before proceeding with the rest of the protocol. Revealing the set of diseases used in the compatibility test would not constitute violation of privacy if such a set of conditions is standard or suggested by S itself.

When, however, the parties compose a custom set of genetic diseases for their genetic compatibility testing and would like to keep the set private, they may be unwilling reveal the set of diseases

³Note that the medical facility that performs sequencing can also certify the test results; alternatively, the medical facility performing test certification will require genome certification from the facility that performed sequencing.

to S. We propose that the parties instead prove that they are providing results for the same conditions without revealing the conditions themselves to the server. The difficulty in doing so arises from the fact that S interacts independently with A and B (possibly at non-overlapping times) and A and B are not proving any joint statements together. Our idea of proving that inputs of A and B correspond to the same sequence of diseases consists of forming a sequence of commitments to the diseases in L , the openings of which are known to both A and B. That is, A and B jointly generate a commitment to each disease using shared randomness and used those commitments at the time of proving that their inputs have been certified. Then if A supplies commitments $\text{Com}_1, \dots, \text{Com}_t$ and proves that the committed values correspond to the diseases in her signatures, S will check that B supplies the same sequence of commitments and also proves that the committed values are equal to the diseases in the signatures he possesses. This will ensure that A and B supply input bits for the same sequence of diseases. To jointly produce the commitments, we have both A and B contribute their own randomness and the resulting commitment will be a function of A’s and B’s contribution. It can proceed as follows (recall that A and B can be malicious):

1. A chooses a random r_A and sends to B a commitment to it $c_A = \text{Com}(r_A, z)$.
2. Similarly, B chooses a random r_B and sends to A a commitment to it $c_B = \text{Com}(r_B, z')$.
3. A and B open their commitments by exchanging (r_A, z) and (r_B, z') and verify that they match c_A and c_B , respectively.
4. They form joint randomness as $r = r_A \oplus r_B$ and use it to construct commitment $\text{Com}(D, r)$.

Given that, the common (high-level) statement that A and B will need to prove about their inputs is:

$$PK\{(\alpha, \sigma) : \text{Sig}(\alpha, \sigma) \wedge y_1 = \text{Com}(\alpha) \wedge y_2 = \text{Com}(\sigma)\}$$

for the shared value of y_1 between A and B, while the remaining portion is specific to A and B as described in section 5.3.

7 Performance Evaluation

In this section, we report on the results of our implementations for ancestry, paternity and compatibility tests. The implementation was written in C/C++ using Miracl library [5] for large number arithmetic and JustGarble library [4] for garbled circuit implementation. The tests were implemented as described in section 6, on which we further elaborate below. The security parameters for symmetric key cryptography and statistical security were set to 128. The security parameter for public key cryptography (for both RSA modulus and discrete logarithm setting) was set to 1536. Additionally, the security parameter for the group size in the discrete logarithm setting was set to 192. All tests (for A, B, and S) were run on a quad-core 3.2GHz machine with Intel i5-3470 processor running Red Hat Linux 2.6.32 in a single core. Note that in practice S is expected to have more powerful hardware and the runtimes can be significantly reduced by utilizing more cores. All experiments were run 5 times, and the mean value is reported.

To provide additional insights into which component of each protocol is main performance bottleneck, we separately report computation times for different parts of each solution (such as those associated with garbled circuit evaluation, OT, etc.). Furthermore, we separately list the times for offline and online computation, where, as in other publications, offline computation refers to all operations that can be performed before the inputs become available. Lastly, because the speed of communication channels can greatly vary, we separately report the size of communication

Party	Garbled circuit		Communication
	garble (offline)	eval (online)	
A	1.8 ms	–	4.0 MB
B	19.8 ms	–	10.0 MB
S	–	13.9 ms	10.0 MB

Table 1: Performance of ancestry test.

Party	Garbled circuit		OT		Total time		Communication
	garble (offline)	eval (online)	offline	online	offline	online	
A	21.6 ms	–	195.2 ms	1983.1 ms	216.8 ms	1983.1 ms	12.1 MB
B	–	13.9 ms	2003.3 ms	218.6 ms	2003.3 ms	232.5 ms	12.1 MB

Table 2: Performance of ancestry test without server.

for each party and communication time is not included in the runtimes. In several settings, however, overlaying computation with communication is possible (e.g., S can perform computation for OT and simultaneously transmit the garbled circuit) and the overall runtime does not need to be the sum of computation and communication time.

7.1 Ancestry test

Let us first consider the ancestry test. Recall that it is implemented in the setting where A and B are semi-honest, but S can be malicious. We ran this test using 2^{17} SNPs as the input for A and B. The resulting circuit used 655,304 XOR gates and 131,072 non-XOR gates. The computation time and communication size are given in Table 1.

The implementation assumes that A only creates labels for her input wires and communicates 2^{17} labels to B. B performs the rest of the garbling work and interacts with S. As expected, the time for circuit garbling and evaluation is small, but the size of communication is fairly large because of the large input size and consecutively circuit size. Nevertheless, we consider the runtimes very small for the computation of this size.

To provide insights into performance gains of our solution compared to the regular two-party computation in the semi-honest setting, we additionally implement the garbled circuit-based solution in presence of semi-honest A and B only. In addition to circuit garbling and evaluation, this also requires the use of oblivious transfer, which we implement using a recent optimized OT extension construction from [7] (including optimizations specific to Yao’s garbled circuit evaluation). As in [7], we use Naor-Pinkas OT for 128 base OTs [50]. The results are given in Table 2. Compared to the server-aided setting, computation is higher by at least two orders of magnitude for each party and communication is noticeably increased as well.

7.2 Paternity test

Next, we look at the paternity test. The test was implemented as described in section 6 in presence of malicious A and B and semi-honest S. The inputs for both A and B consisted of 13 2-element sets, where each element is 9 bits long. We use OT extension from [7] with 128 Naor-Pinkas based OTs. The garbled circuit consisted of 468 XOR gates and 467 non-XOR gates. The results of this experiment are reported in Table 3. The computation for input verification is reported only as part of total time.

Not surprisingly, the cost of OT dominates the overall runtime, but for A the overhead is negligible (the cost of generating labels for input wires and verifying the output label returned by

Party	Garbled circuit		OT		Total time		Communication
	garble (offline)	eval (online)	offline	online	offline	online	
A	$3.5 \cdot 10^{-3}$ ms	–	–	–	$3.5 \cdot 10^{-3}$ ms	3.4 ms	7.4 KB
B	–	$1.1 \cdot 10^{-2}$ ms	515.5 ms	201.7 ms	515.5 ms	201.7 ms	84.9 KB
S	$2.4 \cdot 10^{-2}$ ms	–	196.1 ms	260.9 ms	196.1 ms	260.9 ms	84.9 KB

Table 3: Performance of paternity test.

B). Thus, it is well-suited for settings when one user is very computationally limited.

Compared to two-party computation in presence of malicious participants, our solution reduces both computation and communication for the participants by at least two orders of magnitude. This is because more practical constructions rely on cut-and-choose (and other) techniques to ensure that the party who creates the garbled circuit is unable to learn unauthorized information about the other participant’s input. Recent results such as [48, 6, 49] require the circuit generator to garble on the order of 125 circuits for cheating probability of at most 2^{-40} , some of which are checked (i.e., re-generated) by the circuit evaluator, while the remaining circuits are evaluated. Thus, the work of each of A and B will have to increase by at least two orders of magnitude just for circuit garbling and evaluation, not counting other techniques that deter a number of known attacks and result in increasing the input size and introducing expensive public key operations. A notable exception to the above is the work of Lindell [46] that reduces the number of circuits to 40 for the same cheating probability. The construction, however, results in savings only for circuits of large size as it introduces a large number of additional public key operations. Thus, for paternity tests using constructions with a larger number of circuits is very likely to be faster in practice, which results in a drastic difference between our solution and regular two-party protocol with malicious participants.

Carter et al. [22] treats server-aided secure two-party computation where all parties can behave maliciously. Thus, the difference in the setting from our work is that we assume that the server is semi-honest, which is expected to result in a more efficient protocol. The goal of [22] was to decrease the amount of work for one of the participants, who is assumed to be weak (i.e., a mobile device). The work showed a drastic reduction in the overhead for one of the parties, but the work of the second party and the server is still comparable to the work during execution of regular two-party secure function evaluation in presence of malicious adversaries. We would like to point out that our solution also imposes uneven work on A and B, where A’s computation and communication overheads are minimal. This makes it possible to compare performance of our solution to that of [22]. In our solution, A only creates and sends random labels for her input bits and later performs a hash function computation. In [22], on the other hand, the weaker party has to participate in 128 base OTs, which is orders of magnitude more computation. Each bit of the party’s input is also represented using a number of bits linear in a security parameter, which also increases communication by about two orders of magnitude. The remaining party and the server still need to engage in garbling, checking, and evaluation of about 125 circuits (in addition to other work), which likewise requires about two orders of magnitude larger overhead than in our solution for B and S.

We also would like to mention the solution of [34] that, similar to us, assumes two malicious users and a semi-honest server. As mentioned earlier, the construction of [34] would require A and B to verify on the order of $O(\kappa t_2)$ signatures and engage in $O(\kappa t_2)$ OTs, which translates into tens of thousands of public key operations and significantly larger volume of communication in the paternity test. The server’s work in [34] is larger than in our solution as well.

Baldi et al. [12] also provide a private paternity test in the two-party setting (between a client

Party	Garbled circuit		OT		Sign PK		Other PK		Total time		Comm
	garble/offline	eval/online	offline	online	offline	online	offline	online	offline	online	
A	$1.0 \cdot 10^{-6}$ ms	–	–	–	1.17 s	42.1 ms	616.8 ms	20.6 ms	1.79 s	62.7 ms	46.5 KB
B	–	$7.5 \cdot 10^{-4}$ ms	15.4 ms	14.6 ms	1.17 s	42.1 ms	282.4 ms	15.7 ms	1.47 s	72.4 ms	51.8 KB
S	$2.6 \cdot 10^{-3}$ ms	–	29.3 ms	15.2 ms	0	2.06 s	0	756 ms	29.3 ms	2.83 s	91.0 KB

Table 4: Performance of compatibility test.

and a server). It uses a different computation based on Restriction Fragment Length Polymorphisms (RFLPs) and relies on private set intersection as a cryptographic building block. Both offline and online times for the client and the server are 3.4 ms and the communication size is 3KB for the client and 3.5KB for the server when the test is performed with 25 markers. All times and communication sizes double when the test is run with 50 markers. While the runtimes we report are higher, the implementation of [12] did not consider malicious participants. If protection against malicious A and B in our solution is removed, the work for all parties reduces to well below 0.1 millisecond and communication becomes a couple of KBs.

7.3 Genetic compatibility test

Lastly, we report on the implementation results of the compatibility test. Recall that the test is run in the setting where A and B are malicious and their inputs must be certified. We choose the variant of the solution that reveals the list of diseases L to the server (i.e., a standard list is used). We implement the signature scheme, OT, and ZKPKs as described earlier in this work. All ZKPKs are non-interactive using the Fiat-Shamir heuristic [29]. We used $|L| = 10$ and thus A and B provide 10 input bits into the circuit accompanied by 10 signatures. The circuit consisted of only 19 non-XOR gates. The performance of the test is illustrated in Table 4. We divide all ZKPKs into a proof of signature possession (together with a commitment), which is denoted by “Sign PK” in the table, and the remaining ZK proofs, denoted by “Other PK” in the table. As it is clear from the table, input certification contributes most of the solution’s overhead, but it is still on the order of 1–3 seconds for all parties.

As mentioned earlier, we are not aware of general results that achieve input certification for comparison. However, the comparison to general two-party computation in presence of malicious parties or server-aided two-party computation with malicious server from section 7.2 applies here as well.

Baldi et al. [12] also build a solution and report on the performance of genetic compatibility test. In [12], testing for presence of a genetic disease that client carries in the server genome consists of the client providing the disease fingerprint in the form of (*nucleotide, location*) pairs (which is equivalent to a SNP) and both parties searching whether the disease fingerprint also appears in the server’s DNA. This requires scanning over the entire genome, which our solution avoids. As a result, the solution of [12] incurs substantial offline overhead for the server (67 minutes) and large communication size (around 4GB) even for semi-honest participants. The solution utilizes authorized private set intersection, which allows inputs of one party (as opposed to both in our work) to be verified. Compared to [12], in our framework, testing for a single disease requires a fraction of a second for each party with malicious A and B, where inputs of both of them are certified. The computation is greatly simplified because the list of diseases is assumed to be known by both users. When this is the case, the cost of input certification greatly dominates the overall performance.

8 Conclusions

This work is motivated by the need to protect sensitive genomic data when such data is used in computation, especially in voluntary non-health related computation. Because computation over one's genome often happens in server-facilitated settings, we study server-aided secure two-party computation in a number of security settings. One of such security settings assumes that users A and B may act arbitrarily and, in addition to requiring that the solution remains secure in presence of malicious users, it also enforces that A and B enter their true inputs based on third party's certification. We are not aware of any prior work that combines input certification with general secure multi-party computation based on Yao's garbled circuits. We develop general solutions in our server-aided framework. Despite their generality, they lead to efficient implementations of genetic tests. In particular, we design and implement genetic paternity, compatibility, and common ancestry tests, all of which run in a matter of seconds or less and favorably compare with the state of the art.

Acknowledgments

The authors would like to thank Aaron Steele for discussing genomic tests at early stages of this work. The first author would also like to acknowledge participation in Dagstuhl seminar 13412 on Genomic Privacy, which motivated this work. This work was supported in part by grants CNS-1223699 and CNS-1319090 from the National Science Foundation and FA9550-13-1-0066 from the Air Force Office of Scientific Research. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the funding agencies.

References

- [1] 23andMe – Genetic Testing for Ancestry; DNA Test. <http://www.23andme.com>.
- [2] Genealogy, Family Trees & Family History Records at Ancestry.com. <http://www.ancestry.com>.
- [3] GenePartner.com – DNA matching: Love is no coincidence. <http://www.genepartner.com>.
- [4] The JustGarble library. <http://cseweb.ucsd.edu/groups/justgarble/>.
- [5] The Miracl library. <http://http://www.certivox.com/miracl/>.
- [6] a. shelat and C. h. Shen. Two-output secure computation with malicious adversaries. In *EUROCRYPT*, 2011.
- [7] G. Asharov, Y. Lindell, T. Schneider, and M. Zohner. More efficient oblivious transfer and extensions for faster secure computation. In *ACM Conference on Computer and Communications Security*, 2013.
- [8] E. Ayday, J. L. Raisaro, and J. Hubaux. Personal use of genomic data: Privacy vs. storage cost. In *IEEE Global Communications Conference, Exhibition and Industry Forum*, pages 2723–2729, 2013.
- [9] E. Ayday, J. L. Raisaro, and J.-P. Hubaux. Privacy-enhancing technology for medical tests using genomic data. Technical Report EPFL-REPORT-182897, EPFL, 2012.

- [10] E. Ayday, J. L. Raisaro, J.-P. Hubaux, and J. Rougemont. Protecting and evaluating genomic privacy in medical tests and personalized medicine. In *ACM Workshop on Privacy in the Electronic Society (WPES)*, pages 95–106, 2013.
- [11] E. Ayday, J. L. Raisaro, P. McLaren, J. Fellay, and J.-P. Hubaux. Privacy-preserving computation of disease risk by using genomic, clinical, and environmental data. In *USENIX Workshop on Health Information Technologies (HealthTech)*, 2013.
- [12] P. Baldi, R. Baronio, E. De Cristofaro, P. Gasti, and G. Tsudik. Countering GATTACA: Efficient and secure testing of fully-sequenced human genomes. In *ACM Conference on Computer and Communications Security (CCS)*, pages 691–702, 2011.
- [13] M. Bellare, V. Hoang, S. Keelveedhi, and P. Rogaway. Efficient garbling from a fixed-key blockcipher. In *IEEE Symposium of Security and Privacy*, pages 478–492, 2013.
- [14] F. Bruekers, S. Katzenbeisser, K. Kursawe, and P. Tuyls. Privacy-preserving matching of DNA profiles. IACR Cryptology ePrint Archive Report 2008/203, 2008.
- [15] J. Camenisch and A. Lysyanskaya. A signature scheme with efficient protocols. In *Conference on Security and Cryptography for Networks (SCN)*, pages 268–289, 2002.
- [16] J. Camenisch and A. Lysyanskaya. Signature schemes and anonymous credentials from bilinear maps. In *CRYPTO*, pages 56–72, 2004.
- [17] J. Camenisch and M. Michels. Separability and efficiency for generic group signature schemes. In *CRYPTO*, pages 413–430, 1999.
- [18] J. Camenisch, D. Sommer, and R. Zimmermann. A general certification framework with applications to privacy-enhancing certificate infrastructures. In *Security and Privacy in Dynamic Environments*, pages 25–37, 2006.
- [19] J. Camenisch and M. Stadler. Efficient group signature schemes for large groups. In *CRYPTO*, pages 410–424, 1997.
- [20] J. Camenisch and M. Stadler. Proof systems for general statements about discrete logarithms. Technical report, Institute for Theoretical Computer Science, ETH Zurich, 1997.
- [21] J. Camenisch and G. Zaverucha. Private intersection of certified sets. In *Financial Cryptography and Data Security (FC)*, pages 108–127, 2009.
- [22] H. Carter, B. Mood, P. Traynor, and K. Butler. Secure outsourced garbled circuit evaluation for mobile devices. In *USENIX Security Symposium*, 2013.
- [23] R. Cleve. Limits on the security of coin flips when half the processors are faulty. In *ACM Symposium on Theory of Computing (STOC)*, pages 573–588, 1986.
- [24] I. Damgard and E. Fujisaki. A statistically-hiding integer commitment scheme based on groups with hidden order. In *ASIACRYPT*, pages 125–142, 2002.
- [25] E. De Cristofaro, S. Faber, P. Gasti, and G. Tsudik. GenoDroid: Are privacy-preserving genomic tests ready for prime time? In *ACM Workshop on Privacy in the Electronic Society (WPES)*, pages 97–107, 2012.

- [26] E. De Cristofaro, S. Faber, and G. Tsudik. Secure genomic testing with size- and position-hiding private substring matching. In *ACM Workshop on Privacy in the Electronic Society (WPES)*, pages 107–118, 2012.
- [27] E. De Cristofaro, S. Faber, and G. Tsudik. Secure genomic testing with size- and position-hiding private substring matching. In *ACM Workshop on Privacy in the Electronic Society (WPES)*, pages 107–118, 2013.
- [28] E. De Cristofaro and G. Tsudik. Practical private set intersection protocols with linear complexity. In *Financial Cryptography and Data Security (FC)*, pages 143–159, 2010.
- [29] A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature scheme. In *CRYPTO*, pages 186–194, 1986.
- [30] E. Fujisaki and T. Okamoto. Statistical zero knowledge protocols to prove modular polynomial relations. In *CRYPTO*, pages 16–30, 1997.
- [31] R. Gennaro, C. Gentry, and B. Parno. Non-interactive verifiable computing: outsourcing computation to untrusted workers. In *CRYPTO*, pages 465–482, 2010.
- [32] S. Goldwasser, Y. Kalai, and G. Rothblum. One-time programs. In *CRYPTO*, pages 39–56, 2008.
- [33] D. He, N. Furlotte, F. Hormozdiari, J. Joo, A. Wadia, R. Ostrovsky, A. Sahai, and E. Eskin. Identifying genetic relatives without compromising privacy. *Genome Research*, 24:664–672, 2014.
- [34] A. Herzberg and H. Shulman. Oblivious and fair server-aided two-party computation. In *International Conference on Availability, Reliability and Security (ARES)*, pages 75–84, 2012.
- [35] A. Herzberg and H. Shulman. Oblivious and fair server-aided two-party computation. *Information Security Technical Report*, (17):210–226, 2013.
- [36] F. Hormozdiari, J. Joo, A. Wadia, F. Guan, R. Ostrovsky, A. Sahai, and E. Eskin. Privacy preserving protocol for detecting genetic relatives using rare variants. In *ISMB*, pages 204–2011, 2014.
- [37] Y. Huang, J. Katz, and D. Evans. Quid-pro-quo-tocols: Strengthening semi-honest protocols with dual execution. In *IEEE Symposium of Security and Privacy*, pages 272–284, 2012.
- [38] Y. Ishai, J. Kilian, K. Nissim, and E. Petrank. Extending oblivious transfers efficiently. In *CRYPTO*, pages 145–161, 2003.
- [39] S. Jarecki and V. Shmatikov. Efficient two-party secure computation on committed inputs. In *EUROCRYPT*, pages 97–114, 2007.
- [40] S. Kamara, P. Mohassel, and M. Raykova. Outsourcing multi-party computation. IACR Cryptology ePrint Archive Report 2011/272, 2011.
- [41] S. Kamara, P. Mohassel, and B. Riva. Salus: A system for server-aided secure function evaluation. In *ACM Conference on Computer and Communications Security (CCS)*, pages 797–808, 2012.

- [42] J. Katz and L. Malka. Secure text processing with applications to private DNA matching. In *ACM Conference on Computer and Communications Security (CCS)*, pages 485–492, 2010.
- [43] M. Kiraz, T. Schoenmakers, and J. Villegas. Efficient committed oblivious transfer of bit strings. In *Information Security Conference (ISC)*, pages 130–144, 2007.
- [44] V. Kolesnikov, R. Kumaresan, and A. Shikfa. Efficient verification of input consistency in server-assisted secure function evaluation. In *Cryptology and Network Security (CANS)*, pages 201–217, 2012.
- [45] V. Kolesnikov and T. Schneider. Improved garbled circuit: Free XOR gates and applications. In *International Colloquium on Automata, Languages and Programming (ICALP)*, pages 486–498, 2008.
- [46] Y. Lindell. Fast cut-and-choose based protocols for malicious and covert adversaries. In *CRYPTO*, 2013.
- [47] Y. Lindell and B. Pinkas. A proof of security of Yao’s protocol for two-party computation. *Journal of Cryptology*, 22(2):161–188, 2009.
- [48] Y. Lindell and B. Pinkas. Secure two-party computation via cut-and-choose oblivious transfer. In *Theory of Cryptography Conference (TCC)*, 2011.
- [49] P. Mohassel and B. Riva. Garbled circuits checking garbled circuits: More efficient and secure two-party computation. In *CRYPTO*, 2013.
- [50] M. Naor and B. Pinkas. Efficient oblivious transfer protocols. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 448–457, 2001.
- [51] J. Nielsen, P. Nordholt, C. Orlandi, and S. Burra. A new approach to practical active-secure two-party computation. In *CRYPTO*, pages 681–700, 2012.

A Security Proofs

Proof of Theorem 1 Fairness is achieved based on the fact that A and B are semi-honest. This means that if B ever learns any output, he will share (the part of) the output he receives with the other party. Thus, either both of them learn the (possibly partial) output or neither party learns the output.

To show simulation-based security, we build three independent simulators \mathcal{S}_S , \mathcal{S}_B , and \mathcal{S}_A for three independent adversaries \mathcal{A}_S , \mathcal{A}_B , and \mathcal{A}_A , respectively. Note that \mathcal{A}_S can act in an arbitrary way while \mathcal{A}_B and \mathcal{A}_A are semi-honest.

We first consider a simulator \mathcal{S}_S that communicates with malicious \mathcal{A}_S pretending to be A and B (without access to their private inputs) in such a way that \mathcal{S}_S is unable to distinguish it from an execution in the real model. \mathcal{S}_S proceeds by garbling a circuit corresponding to f and sending it to \mathcal{A}_S . For each wire $i \in [1, t_1 + t_2]$, \mathcal{S}_S sends label $\ell_i^{b_i}$ for a randomly chosen bit b_i . If \mathcal{A}_S does not abort, \mathcal{S}_S obtains from \mathcal{A}_S a set of labels corresponding to the output wires. It is easy to see that the view simulated by \mathcal{S}_S is indistinguishable from the view of \mathcal{A}_S in the real model execution (since the view of garbled circuit evaluator is independent of the input on which the circuit is evaluated). Note that the simulator can safely use random inputs because S is not entitled to receiving any information about the result of function evaluation.

We next consider a simulator \mathcal{S}_B for semi-honest \mathcal{A}_B . Note that \mathcal{A}_B 's advantage is maximized when B constructs the entire garbled circuit (or simply knows the random label pairs for all wires of the circuit), which we assume is the case. After \mathcal{A}_B constructs the circuit and forms the input labels according to x_2 , \mathcal{S}_B queries the trusted party and obtains B's output $f(x_1, x_2)$. \mathcal{S}_B then extracts from \mathcal{A}_B 's view the output labels corresponding to the received output $f(x_1, x_2)$ from the set of the label pairs for the output wires and sends them to \mathcal{A}_B . \mathcal{S}_B also receives assembled $f(x_1, x_2)$ from \mathcal{A}_B destined to A. This view is the same as the view of \mathcal{A}_B in the real model execution given the same set of random coins.

A simulator for \mathcal{A}_A is built analogously with the exception that \mathcal{A}_A receives (from \mathcal{S}_A) $f(x_1, x_2)$ directly instead of learning labels for the output wires. \square

Proof of Theorem 2 As before, we start by showing fairness and then proceed with security. The only way for A or B to learn any output is when A is satisfied with the verification of the output labels she received from B. Recall that each received label ℓ_i is checked against $H(\ell_i^b), H(\ell_i^{1-b})$ for some bit b , where H is a collision-resistant hash function. The probability that this check succeeds for some ℓ_i that is not equal to ℓ_i^0 or ℓ_i^1 is negligible. Thus, A is guaranteed to possess the result of garbled circuit evaluation, at which point both parties are given access to the output.

We next construct simulators for all of the (independent) adversaries $\mathcal{A}_A, \mathcal{A}_B$, and \mathcal{A}_S . We start with a simulator \mathcal{S}_A for malicious \mathcal{A}_A . \mathcal{S}_A runs \mathcal{A}_A and simulates the remaining parties. \mathcal{A}_A produces t_1 random labels ℓ_i^0 and sends them to \mathcal{S}_A , while \mathcal{S}_A chooses Δ and sends it to \mathcal{A}_A . If at least one label is of an incorrect bitlength, \mathcal{S}_A aborts. If \mathcal{S}_A did not abort, \mathcal{A}_A sends t_1 labels to \mathcal{S}_A . If the i th label sent by \mathcal{A}_A does not correspond to one of the labels in the i th pair of labels $(\ell_i^0, \ell_i^0 \oplus \Delta)$ corresponding to \mathcal{A}_A 's inputs, \mathcal{S}_A aborts. If \mathcal{S}_A did not abort, it interprets the meaning of the input labels received from \mathcal{A}_A and stores the input as x'_1 . At some point \mathcal{S}_A creates a random label ℓ_i for each bit of the output and sends them to \mathcal{A}_A . Upon \mathcal{A}_A 's request, \mathcal{S}_A also chooses another random label ℓ'_i for each bit of the output. For each bit i of the output, \mathcal{S}_A sends to \mathcal{A}_A the pair $H(\ell_i), H(\ell'_i)$ in a randomly permuted order. If \mathcal{A}_A notifies \mathcal{S}_A of successful verification of the output labels, \mathcal{S}_A queries the trusted party for the output $f(x'_1, x_2)$. For each i th bit b_i of the output, if $b_i = 0$, \mathcal{S}_A sends to \mathcal{A}_A the pair (ℓ_i, ℓ'_i) , otherwise, \mathcal{S}_A sends the pair (ℓ'_i, ℓ_i) .

Now we examine the view of \mathcal{A}_A in the real and ideal model executions and correctness of the output. After receiving the label pairs from \mathcal{A}_A , \mathcal{S}_A performs the same checks on them as S would and thus both would abort in the same circumstances. Similarly, if \mathcal{A}_A provides malformed labels for circuit evaluation, \mathcal{S}_A will immediately detect this in the ideal model and abort, while B in the real world will be unable to evaluate the circuit and also abort. Otherwise, in both cases the function will be correctly evaluated on the input provided by \mathcal{A}_A and B's input. In the remaining interaction, \mathcal{A}_A sees only random values, which in the ideal world are constructed consistently with \mathcal{A}_A 's view in the real model execution. Thus, \mathcal{A}_A 's view is indistinguishable in the two executions.

Let us now consider malicious \mathcal{A}_B , for which we construct simulator \mathcal{S}_B in the ideal model execution who simulates correct behavior of A and S. First, \mathcal{S}_B simulates the OT. It records the input bits used by \mathcal{A}_B during the simulation, which it stores as x'_2 and returns t_2 random labels to \mathcal{A}_B . \mathcal{S}_B also sends another set of t_1 random labels to \mathcal{S}_B . \mathcal{S}_B queries the trusted party for \mathcal{A}_B 's output $f(x_1, x'_2)$ and chooses a pair of random labels (ℓ_i^0, ℓ_i^1) for each bit i of the output. \mathcal{S}_B gives to \mathcal{A}_B a simulated garbled circuit (as described in [47]) so that the i th computed output label corresponds to the i th bit of $f(x_1, x'_2)$. If after circuit evaluation, \mathcal{A}_B does not send the correct output labels to \mathcal{S}_B , \mathcal{S}_B aborts the execution. Otherwise, \mathcal{S}_B sends the pairs (ℓ_i^0, ℓ_i^1) to \mathcal{A}_B .

The only difference between the view of \mathcal{A}_B during real model execution and the view simulated by \mathcal{S}_B in the ideal model is that \mathcal{A}_B evaluates a simulated circuit in the ideal model. Computational indistinguishability of the simulated circuit follows from the security proofs of Yao's garbled circuit

construction [47]. Thus, \mathcal{A}_B is unable to tell the two worlds apart.

It is also straightforward to simulate the view of semi-honest \mathcal{A}_S because it contributes no input and receives no output. \square