

Policy-Hiding Access Control in Open Environment

(Extended Abstract)

Jiangtao Li
Purdue University
250 N. University Street
West Lafayette, IN 47907
jtli@cs.purdue.edu

Ninghui Li
Purdue University
250 N. University Street
West Lafayette, IN 47907
ninghui@cs.purdue.edu

ABSTRACT

In trust management and attribute-based access control systems, access control decisions are based on the attributes (rather than the identity) of the requester: Access is granted if Alice's attributes in her certificates satisfy Bob's access control policy. In this paper, we develop a policy-hiding access control scheme that protects both sensitive attributes and sensitive policies. That is, Bob can decide whether Alice's certified attribute values satisfy Bob's policy, without Bob learning any other information about Alice's attribute values or Alice learning Bob's policy. To enable policy-hiding access control, we introduce the notion of certified input private policy evaluation. Our construction uses Yao's scrambled circuit protocol and two new techniques introduced in this paper. One novel technique is constructing circuits with uniform topology that can compute arbitrary functions in a family. The other technique is committed-integer based oblivious transfer.

Categories and Subject Descriptors

D.4.6 [Operating Systems]: Security and Protection—*Access controls*; K.6.5 [Management of Computing and Information Systems]: Security and Protection

General Terms

Security, Design

Keywords

Privacy, Access Control, Automated Trust Negotiation, Cryptographic Commitment, Cryptographic Protocol, Digital Credentials, Secure Function Evaluation

1. INTRODUCTION

In trust management and attribute-based access control systems [5, 47, 23, 18, 38, 37], access control decisions are based on attributes of clients, which are often documented by public key certificates. Each certificate associates a public key with the key

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PODC'05, July 17–20, 2005, Las Vegas, Nevada, USA.
Copyright 2005 ACM 1-59593-994-2/05/0007 ...\$5.00.

holder's identity and/or attributes such as employer, group membership, credit card information, birth-date, citizenship, and so on. Because these certificates are digitally signed, they can serve to introduce strangers to one another without online contact with the Certificate Authorities (CA's). As attribute information may be sensitive, the certificates that contain attribute data need protection just as other resources do. Often times, the policies for determining who can access the resources are sensitive also and need protection as well. Consider the following example.

EXAMPLE 1. Bob is a bank offering certain special-rate loans and Alice would like to know whether she is eligible for such a loan before she applies. Alice has a digital driver license certificate issued by the state authority; the certificate contains her birth-date, address, and other attribute data. Alice has also an income certificate issued by her employer documenting her salary and the starting date of her employment. Bob determines whether Alice is eligible for a special-rate loan based on Alice's attribute information. For example, Bob may require that one of the following two conditions holds: (1) Alice is over 30 years old, has an income of no less than \$43K, and has been in the current job for over six months; (2) Alice is over 25 years old, has an income of no less than \$45K, and has been in the current job for at least one year.

Bob is willing to reveal that his loan-approval policy uses the applicant's birth-date, current salary, and the length of the current employment; however, Bob considers the detail of his policy to be commercial secret and does not want to reveal it to others. Alice is interested in this loan and would like to go forward; however, she wants to reveal as little information about her attributes as possible. In particular, Bob shouldn't learn anything about her address (which is also in her driver license) or learn her actual birth-date. Ideally, Alice wants Bob to know whether she is eligible for the loan, but nothing else.

In the above example, the policy is a commercial secret, and knowledge of Bob's policy would compromise Bob's strategy and invite unwelcome imitators. In other examples, the motivation for hiding the policy is not necessarily protection from an evil adversary, but simply the desire to prevent legitimate users from gaming the system; e.g., changing their behavior based on their knowledge of the policy. This is particularly important for policies that are not incentive-compatible in economic terms. Furthermore, it is important to point out that a process that protects Alice's certificates from Bob is not only to Alice's advantage but also to Bob's: Bob no longer needs to worry about rogue insiders in his organization illicitly leaking (or selling) Alice's private information, and may even lower his liability insurance rates as a result of this. Privacy-preservation is a win-win proposition, one that is appealing even if Alice and Bob are honest and trustworthy entities.

Motivated by the preceding applications, we introduce and study the problem of *policy-hiding access control*. In this framework, Bob has a private policy and Alice has several sensitive certificates. In the end, Bob learns whether Alice’s attributes in her certificates satisfy his policy but nothing else about her attribute values; at the same time, Alice does not learn Bob’s policy except for what attributes are required for his policy.

One may tempt to use existing general solutions to the two-party Secure Function Evaluation (2-SFE) [55, 30, 28] (e.g., Yao’s scramble circuit protocol [55]) for policy-hiding access control. That is, Alice inputs her certificates and Bob inputs his policy; and they run a 2-SFE protocol to evaluate Bob’s policy on Alice’s attributes in her certificates. Such approach does not work well because (1) the function to compute in 2-SFE is public, whereas the function (Bob’s policy) in policy-hiding access control is private; (2) as Alice needs to input her certificates into 2-SFE, certificate verification, which involves verifying digital signatures, needs to be done as a part of 2-SFE circuit evaluation. This is extremely inefficient. Observe that Alice is not allowed to input her attribute values directly (instead of her certificates), because, Alice otherwise can input arbitrary faked attribute values at her will¹.

To avoid verifying certificates within circuit evaluation, we use *Oblivious Attribute Certificates* (OACerts) proposed in [34]. In an OACert, attribute values are not directly stored in the certificate; instead, a cryptographic commitment for each of these values is stored. A certificate authority (CA) generates the commitments of Alice’s attribute values and stores them in the certificates. Alice is able to disclose her certificates to Bob without revealing her attribute values. OACerts can be integrated into current standards for public-key certificates such as X.509 public-key certificates [6, 33] and X.509 attribute certificates [24]. A prototype of OACerts has been implemented by storing the commitments in X.509v3 extension fields [34]. Using OACerts, the policy-hiding access control problem becomes that Alice inputs her committed attributes which are documented in her OACerts and Bob inputs his policy, they want to learn whether Alice’s committed attributes satisfy Bob’s policy without revealing the other party’s private input.

In this paper, we introduce the notion of *Certified Input Private Policy Evaluation (CIPPE)*, which enables policy-hiding access control using OACerts. Formal definition of CIPPE is given in Section 3. In CIPPE, Alice has private inputs x_1, x_2, \dots, x_n , Bob has a private function f drawn from a family \mathcal{F} of functions (usually f outputs ‘yes’ or ‘no’; however, we allow functions that output more than one bit of information), and Alice and Bob share c_1, c_2, \dots, c_n , where c_i is a cryptographic commitments of x_i , for $1 \leq i \leq n$. The objective of CIPPE is for both Alice and Bob to learn the result of $f(x_1, \dots, x_n)$. Bob should not learn anything about x_1, \dots, x_n ; and Alice should not learn more than the fact that $f \in \mathcal{F}$.

We develop a CIPPE protocol for certain families of functions that we believe are useful for expressing policies. Our solution uses Yao’s scrambled circuit protocol [55, 39]. When a circuit is scrambled, the operation in each gate is hidden; however, the topological structure of the circuit is not. Therefore, Alice could infer some information about Bob’s policy by looking at the scrambled circuit if Bob constructs the circuit in the naive way. To protect Bob’s private function, we develop an efficient approach to construct circuits with uniform topology that can compute certain functions families. To ensure that Alice can evaluate the scrambled circuit only with her attribute values as committed in her certificates, we develop an efficient and provably secure Committed-Integer based Oblivious

¹In SFE, there is no way to prevent a dishonest party from changing its local input before the protocol execution.

Transfer (CIOT) protocol. The computation and communication complexity of the proposed CIPPE protocol is close to the complexity of the scramble circuit protocol that computes $f(x_1, \dots, x_n)$ where f is public, and x_1, \dots, x_n are not committed. The CIPPE protocol is efficient; and we believe it can be deployed in practice (see [39] for an implement of the scramble circuit protocol by Malkhi et al.).

The rest of this paper is organized as follows. We first describe how CIPPE can be used to enable policy-hiding access control in Section 2. Then we give a formal definition of CIPPE in Section 3. In Section 4, we review two cryptographic building blocks that we use, namely, the Pedersen commitment scheme and the scrambled circuit protocol for 2-SFE. In the next two sections, we present two building blocks that we build for CIPPE, one is circuit construction of policy functions with uniform topology, the other is the CIOT protocol. In Section 7 we give an efficient construction for CIPPE. We discuss the related work in Section 8 and conclude our paper in Section 9.

2. USING CIPPE FOR POLICY-HIDING ACCESS CONTROL

In this section, we present a high-level framework for policy-hiding access control using CIPPE. We describe how policy-hiding access control in Example 1 can be enabled. In what follows, we use *commit* to denote the commitment algorithm of a commitment scheme. Let *Params* denote the public parameters for *commit*. To be secure, a commitment scheme cannot be deterministic; thus a commitment of a value a also depends on an auxiliary input, a secret random value r . We use $c = \text{commit}_{\text{Params}}(a, r)$ to denote a commitment of a .

1. **CA Setup.** Let Bureau of Motor Vehicles (BMV) be the CA who issues digital driver licenses. BMV runs the CA setup program, i.e., BMV picks a signature scheme, a commitment scheme denoted by *commit*, a pair of public/private keys, and the public parameters for the commitment scheme, *Params*. Let Company C be Alice’s employer, the CA that issues an income certificate for Alice. Company C runs the CA setup program analogously.
2. **Alice-CA Interaction.** In this phase, Alice obtains two OACerts, one from BMV and the other from Company C. Alice applies for a digital driver license certificate from BMV as follows. BMV first verifies the correctness of her attribute values through some (possibly off-line) channels, then issues an OACert for Alice. The OACert is signed using the BMV’s key and contains Alice’s public key, BMV’s public key, and a commitment for each attribute value that is to be included in the certificate. For example, let x be Alice’s birth-date (encoded as an integer), BMV generates a random number r , computes $c = \text{commit}_{\text{Params}}(x, r)$, and stores c in the OACert. The BMV sends the signed OACert to Alice, together with all the secret random values that have been used. Similarly Alice obtains an income certificate from her employer Company C.
3. **Alice-Bob Setup.** Alice applies for a special-rate loan from Bob. Bob reveals that the loan policy takes three attributes: birth-date, current salary, and the length of current employment. Alice shows her driver license OACert and income OACert to Bob. Alice then proves the ownership of her OACerts using the usual techniques [33]. Recall that OACerts can be used as regular digital certificates (e.g., X.509 certificates) except the attribute values are stored in the committed form.

4. **Alice-Bob Interaction.** Alice and Bob run an interaction protocol, where Alice inputs her attribute values and secret random values she has stored from Phase 2 (Alice-CA Interaction) and Bob inputs his private policy function. In the end, both Alice and Bob learn whether Alice satisfies Bob’s policy without getting other information about Alice’s attributes or Bob’s policy.




Figure 1: An example of policy-hiding access control procedures between Alice and Bob.

Figure 1 depicts how CIPPE can be used in the policy-hiding access control. We observe that the two CA’s are involved only in issuing certificates to Alice. When Alice is interacting with various servers such as Bob, the CA’s are not involved and can be off-line.

3. DEFINITION OF CERTIFIED INPUT PRIVATE POLICY EVALUATION (CIPPE)

We now give a formal definition of CIPPE, which allows us to prove our protocol for CIPPE is secure.

DEFINITION 1 (CIPPE). A CIPPE scheme is parameterized by a commitment scheme commit . A CIPPE scheme involves a client C , a server S , and a trusted CA, and has the following four phases:

CA Setup CA takes a security parameter σ and another parameter ℓ (which specifies the desired range of the attribute values), and outputs public parameters Params for commit . The domain of commit contains $[0..2^\ell - 1]$ as a subset. CA sends Params to C and S .

Client-CA Interaction C chooses n values $x_1, \dots, x_n \in [0..2^\ell - 1]$ (these are C ’s attribute values) and sends them to CA. For each i , $1 \leq i \leq n$, CA generates a new random number r_i and computes the commitment $c_i = \text{commit}_{\text{Params}}(x_i, r_i)$. CA gives c_i and r_i to C , and c_i to S .

Recall that in the actual usage scenario in Section 2, CA does not directly communicate with S . Instead, CA verifies C ’s attribute values before computing the commitments and storing c_1, \dots, c_n into C ’s OACerts. The OACerts are then sent by C to S , enabling S to have the commitment values as if they were sent from CA. Here we abstract these steps away to have CA sending c_i to S . We stress that CA does *not* directly participate in the policy-hiding access control process between C and S .

Client-Server Setup S chooses a family \mathcal{F} of functions and sends the description of \mathcal{F} to C (this models the fact that \mathcal{F} is public knowledge). Each f in \mathcal{F} maps n ℓ -bit integers to a bit, i.e., $f : ([0..2^\ell - 1])^n \rightarrow \{0, 1\}$. S chooses a function $f \in \mathcal{F}$ privately.

Now S has c_1, \dots, c_n , and f . C has $c_1, \dots, c_n, x_1, \dots, x_n$, and r_1, \dots, r_n .

Client-Server Interaction C and S run an interactive protocol. In the end, both C and S output $F(x_1, r_1, \dots, x_n, r_n, f)$, where F takes c_1, \dots, c_n , and \mathcal{F} as parameters, and is defined as

$$F(x_1, r_1, \dots, x_n, r_n, f) = \begin{cases} f(x_1, \dots, x_n) & \text{if } f \in \mathcal{F} \wedge \\ & c_i = \text{commit}(a_i, r_i) \text{ for each } i \\ \perp & \text{otherwise.} \end{cases}$$

When both C and S are honest, C and S will output $f(x_1, \dots, x_n)$ in this phase.

To avoid unnecessarily cluttering the exposition, in Definition 1 we assume that there is only one CA in a CIPPE scheme, and that x_1, \dots, x_n are equal-length and are committed under the same commitment parameters. The definition of CIPPE can be modified to support multiple CA’s, different input lengths, and different commitment parameters. As a matter of fact, we can easily adjust our CIPPE protocol to support the situation in which each x_i is committed under a different set of commitment parameters.

Notion of Security

The security definitions we use follow [28, 13, 14]. We consider security against three kinds of adversaries. An *adversary* is a probabilistic interactive Turing Machine [31]. A *honest-but-curious* adversary is an adversary who follows the prescribed protocol, and attempts to learn more information than allowed from the execution. A *weak-honest* adversary [14] is an adversary who may deviate arbitrarily from the protocol, as long as her behavior appears honest to parties executing the protocol. A *malicious* adversary is an adversary who may behave arbitrarily. When we consider malicious adversaries, there are certain things we cannot prevent: an adversary (1) may refuse to participate in the protocol, (2) may substitute its local input with something else, and (3) may abort the protocol prematurely. When we consider weak-honest, we cannot prevent an adversary from substituting her local input.




Figure 2: Ideal model for the CIPPE interaction protocol.

The security of a CIPPE protocol is analyzed by comparing what an adversary can do in the protocol to what she can do in the ideal model with a Trusted Third Party (TTP). In the ideal model, as depicted in Figure 2, the client sends her private input x_i and r_i , for $1 \leq i \leq n$, to the TTP, and the server sends his private input

f to the TTP. The TTP verifies that $c_i = \text{commit}(x_i, r_i)$ for each i and $f \in \mathcal{F}$, computes $f(x_1, \dots, x_n)$, sends the result back to the client and the server. If the verification fails, the TTP simply outputs a special symbol \perp .

The ideal model differs for honest-but-curious adversaries, weak-honest adversaries, and malicious adversaries. In the ideal model for honest-but-curious adversaries, an honest party outputs her output from the TTP, whereas an honest-but-curious party outputs an arbitrary function from her initial input and the output she obtained from the TTP. The ideal model for weak-honest adversaries is similar to the ideal model for honest-but-curious adversaries, but differs in that a weak-honest adversary can substitute her input before sending to the TTP. In the ideal model for malicious adversaries, a malicious adversary can terminate the protocol prematurely, even at a stage when she has received her output and the other party has not.

DEFINITION 2 (SECURITY). Let F be the function the client and server compute in the interaction phase of a CIPPE scheme. Let Π be the CIPPE protocol for computing F . We model the client and server as a pair of admissible probabilistic polynomial-time machines, where at least one of them is honest. Protocol Π *securely computes* F if for every pair of admissible probabilistic polynomial-time machines (C^*, S^*) in real model, there exists a pair of admissible probabilistic polynomial-time machines (C, S) in the ideal model, such that the joint execution of F under (C^*, S^*) in the real model is computationally indistinguishable from the joint execution of F under (C, S) in the ideal model.

Our construction for CIPPE is provably secure in the honest-but-curious model and the weak-honest model. The server’s privacy is guaranteed against any malicious client. A malicious server may learn additional information about a client’s attributes; however, this additional information is limited to at most one bit and such malicious behavior will be detected by the client.

4. CRYPTOGRAPHIC ASSUMPTIONS AND TOOLS

In this section we first give the cryptographic assumptions and then briefly review two cryptographic tools that we use for our CIPPE construction: the Pedersen commitment scheme [45] and the scrambled circuit protocol for 2-SFE [55].

Basic Cryptographic Assumptions

The security of our CIPPE protocol is based on the following two standard assumptions in cryptography.

1. *Discrete Logarithm (DL) Assumption.* Given a finite cyclic group G , a generator $g \in G$, and a group element y , the DL assumption states that there exists no polynomial-time algorithm that can compute $\log_g y$ with non-negligible probability.
2. *Decision Diffie-Hellman (DDH) Assumption.* Given a finite cyclic group G and a generator $g \in G$, the DDH assumption states that there exists no polynomial-time algorithm that can distinguish between two distributions $\langle g^a, g^b, g^{ab} \rangle$ and $\langle g^a, g^b, g^c \rangle$, where a, b, c are chosen at random.

A function μ is called *negligible* in the security parameter σ if, for every polynomial p , $|\mu(\sigma)|$ is smaller than $1/p(\sigma)$ for large enough σ ; otherwise, it is *non-negligible*.

The Pedersen Commitment Scheme

Setup A trusted third party T chooses two large prime numbers p and q such that q divides $p - 1$. It is typical to have p be 1024 bits and q be 160 bits. T picks g to be a generator of G_q , the unique order- q subgroup of \mathbb{Z}_p^* . We use $s \in_R \mathbb{Z}_q$ to denote that s is uniformly randomly chosen from \mathbb{Z}_q . T picks $s \in_R \mathbb{Z}_q$ and computes $h = (g^s \bmod p)$. T keeps the value s secret and makes the values p, q, g, h public.

Commit The domain of the committed values is \mathbb{Z}_q . For a party A to commit an integer $x \in \mathbb{Z}_q$, A chooses $r \in_R \mathbb{Z}_q$ and computes the commitment $c = (g^x h^r \bmod p)$.

Open To open a commitment c , A reveals x and r , and a verifier verifies whether $c = (g^x h^r \bmod p)$.

We use a trusted third party T to generate the parameters of the Pedersen commitment scheme because the setup algorithm is run by a trusted CA in the CIPPE setting. The Pedersen commitment scheme is *unconditionally hiding*: Even with unlimited computational power it is impossible for an adversary to learn any information about the value x from c , because the commitments of any two numbers in \mathbb{Z}_q have exactly the same distribution. This commitment scheme is *computationally binding*: Under the DL assumption, it is computationally infeasible for an adversarial committer to open a value x' other than x in the open phase of the commitment scheme.

The scrambled circuit protocol for 2-SFE

The scrambled circuit protocol was developed by Yao [55] (see [30, 42, 29] for detailed descriptions of the protocol). This protocol runs between two players: a (circuit) *generator* and a (circuit) *evaluator*. Let x be the evaluator’s input, y be the generator’s input, and $f : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^*$ be a function known to both parties. The goal of the scramble circuit protocol is for the two parties to securely evaluate $f(x, y)$ without the generator learning x or the evaluator learning y . In the protocol, the generator first constructs a circuit for computing f . The generator then constructs a scrambled version of the circuit and sends the scrambled circuit to the evaluator for evaluation. In a scrambled circuit, each wire is associated with two random numbers, one corresponds to 0 and the other to 1. Before the evaluation, the evaluator uses oblivious transfer [41] to obtain the random values corresponding to each bit of the evaluator’s private input x . During the evaluation, the evaluator learns exactly one random value for each internal wire, yet she doesn’t know whether it corresponds to 0 or 1. Finally the evaluator sends the outcome of the evaluation to the generator, who then recovers the final result.

5. BUILDING CIRCUITS THAT HAVE UNIFORM TOPOLOGICAL STRUCTURE

When a circuit is scrambled, the operation in each gate is hidden; however, the topological structure of the circuit is not. Therefore, the client could infer some information about the server’s function by looking at the scrambled circuit if the server constructs the circuit in a naive way. To protect the server’s private function, we present an approach to construct circuits that can compute a family of functions and have the same topological structure.

Function definition

We propose a family \mathcal{F} of functions that can express many policy functions in real applications. We define \mathcal{F} as follows. \mathcal{F} has four parameters ℓ, n, m , and λ . Each function f in $\mathcal{F}(\ell, n, m, \lambda)$ takes

m parameters $y_1, \dots, y_m \in [0..2^\ell - 1]$ and n inputs $x_1, \dots, x_n \in [0..2^\ell - 1]$, and maps them to $\{0, 1\}$. Let $f(x_1, \dots, x_n) = p(x_{i_1} \text{op}_1 y_1, x_{i_2} \text{op}_2 y_2, \dots, x_{i_m} \text{op}_m y_m)$, where $1 \leq i_1, i_2, \dots, i_m \leq n$, each op_i is one of the following predicates $\{=, \neq, >, <, \geq, \leq\}$, and p is a disjunctive (or conjunctive) normal form in which the number of disjuncts (or conjuncts) is no more than λ .

If the server chooses a function f from the family $\mathcal{F}(\ell, n, m, \lambda)$ of functions, the client should not be able to distinguish f from any other functions in the family. For instance, consider Example 1 in Section 1, Bob (the bank) can set $n = 3, m = 8, \lambda = 4$, and the policy function is of the form:

$$f(x_1, x_2, x_3) = (x_1 \geq 30 \wedge x_2 \geq 43000 \wedge x_3 > 6) \vee (x_1 \geq 25 \wedge x_2 \geq 45000 \wedge x_3 > 12),$$

where x_1 denotes age, x_2 denotes annual income in dollars, and x_3 denotes length of current employment in months. Alice learns that x_1, x_2 , and x_3 are used for comparison at most 8 times, she would not learn information such as which values they are compared with, and how many times each attribute is compared.

If Bob builds a circuit for $f(x_1, x_2, x_3)$ in a naive fashion, Alice can learn from the topology of the circuit how many times each x_i is compared, what these comparison operators are, and some information about the structure of the policy function. One technical difficulty in hiding such information is that each attribute may be used in multiple comparisons, and we want to hide the number of times it is used. A straightforward way to do this is to use m circuits, each of which select one input from the n inputs. This is not efficient as it needs $O(nm)$ gates. Our construction uses results from the literature on permutation and multicast switching networks (see, for example [46, 49, 2, 58, 54]). Some of these networks may be useful for constructing circuits for families of functions beyond the ones considered in this paper.

Basic circuit components

We introduce three basic circuit components that will be used in our construction. We depict them in Figure 3.

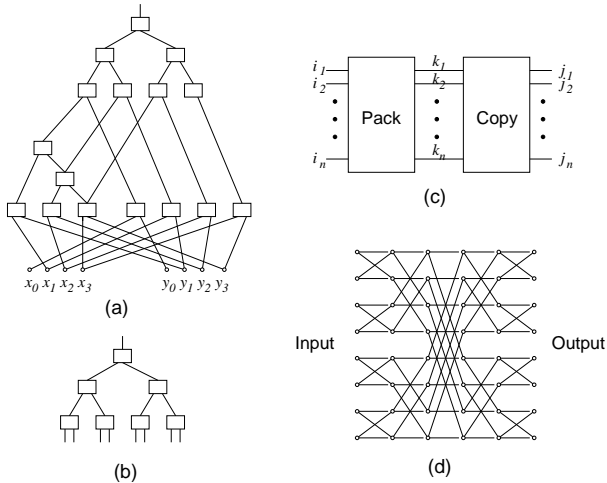


Figure 3: Basic circuit components: (a) the structure of 4-bit comparison circuits, (b) the structure of 8-input logical operation circuits, (c) the high-level schema for a generalizer circuit, (d) an (8,8)-generalizer,

1. *Comparison circuit.* Given two ℓ -bit integers x and y , the comparison circuit computes $x = y$, $x \neq y$, $x > y$, or

$x < y$. Observe that $x \geq y$ and $x \leq y$ can be represented as $x > y - 1$ and $x < y + 1$, respectively. Let $x_{\ell-1} \dots x_1 x_0$ be the binary representation of x and $y_{\ell-1} \dots y_1 y_0$ be the binary representation of y .

- Circuit for $x > y$ is $\bigvee_{i=0}^{\ell-1} (x_i \wedge \neg y_i \wedge \bigwedge_{j=i+1}^{\ell-1} (x_j = y_j))$
- Circuit for $x < y$ is $\bigvee_{i=0}^{\ell-1} (\neg x_i \wedge y_i \wedge \bigwedge_{j=i+1}^{\ell-1} (x_j = y_j))$
- Circuit for $x = y$ is $\bigwedge_{i=0}^{\ell-1} (x_i = y_i)$
- Circuit for $x \neq y$ is $\bigvee_{i=0}^{\ell-1} (x_i \neq y_i)$

Note that the circuits for $x > y$ and $x < y$ have the same topology. To make the structure of all comparison circuits uniform, we modify the circuits for $x = y$ and $x \neq y$ by adding some “dummy” gates. For example, the comparison circuit for $x = y$ could be $\bigwedge_{i=0}^{\ell-1} ((x_i = y_i) \wedge \bigwedge_{j=i+1}^{\ell-1} g(x_j, y_j))$ where $g(x_i, y_i)$ always outputs 1. Figure 3(a) shows the structure of 4-bit comparison circuits. Note that each ℓ -bit comparison circuit requires $O(\ell)$ gates ($5\ell - 4$ gates).

2. *Logical operation circuit.* Given m Boolean inputs a_1, \dots, a_m , the logical circuit computes $\bigvee_{i \in S} a_i$ or $\bigwedge_{i \in S} a_i$ where $S \subseteq \{1, 2, \dots, m\}$. We can use a binary tree structure to implement the m -input logical circuit. For example, to compute the logical formula $\bigvee_{i \in S} b_i$, every gate in the binary tree computes \vee ; if $i \in S$ we give the corresponding wire value a_i , otherwise, set value 0. Figure 3(b) shows a 8-bit logical operation circuit. Note that the m -input logical circuits require $O(m)$ gates ($m - 1$ gates).
3. *Generalizer circuit.* An (n, n) -generalizer is an n -input and n -output switching network, it passes each input i to zero or more outputs. The existence of (n, n) -generalizer with $O(n)$ gates is demonstrated nonconstructively by Pipenger [46]. Ofman [43] gives a construction of a generalizer using the schema shown in Figure 3(c). In his construction, the network consists of two parts: a pack network and a copy network. The pack network packs those inputs having requests to consecutive positions. The copy network copies inputs to multiple outputs. The network proposed by Ofman [43] requires $3n \log n$ gates. Thompson [49] improved Ofman’s work and gives a construction using $2n \log n$ gates. The Thompson’s construction uses a reversed butterfly network concatenated with a butterfly network. Figure 3(d) is the Thompson’s construction of a $(8, 8)$ -generalizer.

Our construction

Our construction takes the following three stages.

1. *Copy Stage.* The copy stage takes n ℓ -bit integers x_1, \dots, x_n and outputs m ℓ -bit integers in which each x_i is copied to output v_i times where $v_i \geq 0$ and $\sum v_i = m$. To build the copy stage in circuit, we construct ℓ identical (n, m) -generalizers, one for each bit. A (n, m) -generalizer can be implemented by $\lceil \frac{m}{n} \rceil$ numbers of (n, n) -generalizer. This stage needs $O(\ell m \log n)$ gates (about $2\ell m \log n$ gates).
2. *Comparison Stage.* The comparison stage takes m ℓ -bit integers and makes m comparisons. This stage consists of m comparison circuits, one for each (x, y) pair. This stage needs $O(\ell m)$ gates (about $5\ell m$ gates).

3. *Logical Computation Stage.* Observe that all the disjunctive normal forms where the number of conjunctions is no more than λ can be expressed as $\bigvee_{j=1}^{\lambda} (\bigwedge_{i \in S_j} a_i)$, where $S_1, S_2, \dots, S_m \subseteq \{1, 2, \dots, m\}$. Such disjunctive normal forms can be implemented using λ m -input logical operation circuits and one λ -input logical operation circuits. For each m -input logical operation circuit, the input consists of the m output bits from the comparison stage, the output is connected to the input wire of the last λ -input logical operation circuit. The conjunctive normal forms can be implemented analogously. This stage needs $O(\lambda m)$ gates (about λm gates).

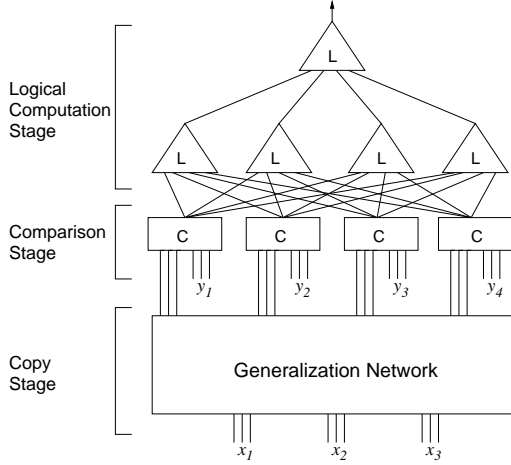


Figure 4: An example circuit structure for the family \mathcal{F} of functions with parameters $\ell = 3$, $n = 3$, $m = 4$, and $\lambda = 4$. There are 4 comparison circuits in the comparison stage, and 5 logical operation circuits in the logical computation stage.

Figure 4 shows the structure of circuits that can compute the family $\mathcal{F}(3, 3, 4, 4)$ of functions. For the family $\mathcal{F}(\ell, n, m, \lambda)$ of functions, our circuit construction needs $O(\ell m \log n + \lambda m)$ gates (around $(2 \log n + 5)\ell m + \lambda m$ gates).

6. A COMMITTED-INTEGER BASED OBLIVIOUS TRANSFER (CIOT) PROTOCOL

To build a CIPPE protocol using the scrambled circuit protocol, we have to ensure that the client gets the keys of the input wires corresponding to her committed input. We present a Committed-Integer based Oblivious Transfer (CIOT) protocol to achieve this. A CIOT protocol involves a sender and a receiver. The receiver has a committed ℓ -bit integer x , the sender has ℓ pairs of values $(k_1^0, k_1^1), \dots, (k_\ell^0, k_\ell^1)$, and both the sender and receiver share the commitment of x . In the end of the protocol, the receiver learns exactly one key in each pair; furthermore, the keys she learns correspond to the bits in x . The main idea of CIOT is as follows. Using the commitment of x , the receiver generates ℓ new commitments, one for each bit of x . Then the sender and receiver run a modified version of non-interactive oblivious transfer protocol [4, 41] for each commitment.

PROTOCOL 1 (CIOT PROTOCOL). Let $\langle p, q, g, h \rangle$ be the public parameters of the Pedersen commitment scheme. All arithmetic in this section is mod p unless specified otherwise. Let x

be an integer in $[0..2^\ell - 1]$, and $x_{\ell-1} \dots x_1 x_0$ be the binary representation of x , i.e., $x = x_0 2^0 + x_1 2^1 + \dots + x_{\ell-1} 2^{\ell-1}$. Let $c = \text{commit}(x, r) = g^x h^r$ be the commitment of x with a random $r \in \mathbb{Z}_q$.

Input The receiver has x and r , and the sender has ℓ pairs of integers $(k_0^0, k_0^1), \dots, (k_{\ell-1}^0, k_{\ell-1}^1)$. Both the sender and receiver have c .

Output The receiver learns $k_0^{x_0}, \dots, k_{\ell-1}^{x_{\ell-1}}$. The sender learns nothing.

1. The receiver decomposes c into ℓ commitments, one for each bit of x . More specifically, the receiver randomly picks $r_1, \dots, r_{\ell-1} \in \mathbb{Z}_q$ and sets $r_0 = r - \sum_{i=1}^{\ell-1} 2^i r_i \text{ mod } q$. The receiver computes $c_i = \text{commit}(x_i, r_i) = g^{x_i} h^{r_i}$ for $i = 0, 1, \dots, \ell - 1$, and gives them to the sender. The sender checks that $\prod_{i=0}^{\ell-1} (c_i)^{2^i} = c$. Observe that $\prod_{i=0}^{\ell-1} (c_i)^{2^i} = \prod_{i=0}^{\ell-1} (g^{x_i} h^{r_i})^{2^i} = g^{\sum_{i=0}^{\ell-1} x_i 2^i} h^{\sum_{i=0}^{\ell-1} r_i 2^i} = g^x h^r = c$.
2. For $i = 0, 1, \dots, \ell - 1$, the sender calculates $K_i^0 = \langle p, q, h, c_i \rangle$ and $K_i^1 = \langle p, q, h, c_i g^{-1} \rangle$. Using the ElGamal encryption scheme [22] (modified to have messages from a subgroup [50]), the sender sends to the receiver two ciphertexts $E_{K_i^0}(k_i^0) = (h^{y_i}, k_i^0 c_i^{y_i})$ and $E_{K_i^1}(k_i^1) = (h^{z_i}, k_i^1 (c_i g^{-1})^{z_i})$, where y_i and z_i are chosen at uniform random from \mathbb{Z}_q by the sender. The receiver can obtain $k_i^{x_i}$ as follow: If x_i equals 0, then $c_i = h^{r_i}$, the receiver knows the private key corresponding to K_i^0 (the private key is r_i), therefore she can decrypt $E_{K_i^0}(k_i^0)$ to recover k_i^0 . If x_i equals 1, then $c_i g^{-1} = h^{r_i}$, the receiver knows the private key corresponding to K_i^1 , she can decrypt $E_{K_i^1}(k_i^1)$ to recover k_i^1 .

Both the sender and receiver need $O(\ell)$ modular exponentiation, i.e., the sender needs 2ℓ modular exponentiation, and receiver needs 4ℓ modular exponentiation. The security properties of the CIOT protocol are given by the following theorems.

THEOREM 1. *The sender does not learn anything from the CIOT protocol.*

PROOF. The CIOT protocol consists of two phases: a bit-commitment phase and an oblivious transfer phase. The sender learns nothing about x from the oblivious transfer phase, as the receiver does not send any information to the sender during that phase. Thus, all the information the sender learns about x is from the bit-commitment phase. In the bit-commitment phase, the sender learns the commitments $c_0, \dots, c_{\ell-1}$. Observe that $c_0 = c \prod_{i=1}^{\ell-1} (c_i)^{-2^i}$; therefore, c_0 can be computed from $c, c_1, \dots, c_{\ell-1}$ and does not leak any additional information. Recall that $r, r_1, \dots, r_{\ell-1}$ are chosen uniformly randomly from \mathbb{Z}_q ; the distributions of $c, c_1, \dots, c_{\ell-1}$ are exactly the same as the distribution of any commitment under the Pedersen commitment scheme. Thus $c, c_1, \dots, c_{\ell-1}$ leak nothing about their corresponding committed values $x, x_1, \dots, x_{\ell-1}$. Therefore, the sender does not learn anything about x . In other words, for any $x, y \in [0..2^\ell - 1]$ (let c_x, c_y be the corresponding commitments), and for any adversary executing the sender's part, the views that the adversary sees when the receiver inputs (x, c_x) and when the receiver inputs (y, c_y) are perfectly indistinguishable. \square

THEOREM 2. *Under the DDH assumption and the DL assumption on G_q , the order- q subgroup of \mathbb{Z}_p^* , the receiver learns at most one value per (k_i^0, k_i^1) pair.*

PROOF. Suppose an adversarial receiver learns both k_i^0 and k_i^1 for some given i , where $0 \leq i \leq \ell - 1$. Under the DDH assumption, the ElGamal encryption scheme is semantically secure [50]. Therefore, the adversary knows the private keys corresponding to the ElGamal public keys $K_i^0 = \langle p, q, h, c_i \rangle$ and $K_i^1 = \langle p, q, h, c_i g^{-1} \rangle$. In other words, the adversary knows r where $h^r = c_i$, and r' where $h^{r'} = c_i g^{-1}$. Thus, the adversary knows r and r' where $h^r = gh^{r'}$; she can effectively compute $\log_g(h) = (r - r')^{-1} \bmod q$, which contradicts the DL assumption. \square

THEOREM 3. *Under the DDH assumption and the DL assumption on G_q , the order- q subgroup of \mathbb{Z}_p^* , if the receiver learns ℓ keys, these values must be $k_0^{x_0}, \dots, k_{\ell-1}^{x_{\ell-1}}$.*

PROOF. By Theorem 2, if an adversarial receiver learns ℓ keys, she learns exactly one key per (k_i^0, k_i^1) pair. Suppose she learns $k_0^{y_0}, \dots, k_{\ell-1}^{y_{\ell-1}}$, where $y_i \in \{0, 1\}$ for $i \in [0.. \ell - 1]$ and there exist at least one j such that $x_j \neq y_j$. Therefore, $\sum_{i=0}^{\ell-1} y_i 2^i \neq \sum_{i=0}^{\ell-1} x_i 2^i = x$. Under the DDH assumption, the adversary knows the private keys corresponding to the ElGamal public keys $K_0^{y_0}, \dots, K_{\ell-1}^{y_{\ell-1}}$; thus she knows t_i for each $i \in [0.. \ell - 1]$ such that $g^{y_i} h^{t_i} = c_i$. As

$$g^x h^r = c = \prod_{i=0}^{\ell-1} (c_i)^{2^i} = \prod_{i=0}^{\ell-1} (g^{y_i} h^{t_i})^{2^i} = g^{\sum_{i=0}^{\ell-1} y_i 2^i} h^{\sum_{i=0}^{\ell-1} t_i 2^i} = g^y h^t,$$

where y denotes $\sum_{i=0}^{\ell-1} y_i 2^i$ and t denotes $\sum_{i=0}^{\ell-1} t_i 2^i \pmod{q}$, the receiver knows x , r , y , and t such that $g^x h^r = g^y h^t$. The receiver can efficiently compute $\log_g(h)$, which contradicts the DL assumption. \square

In the CIOT protocol, a malicious receiver may learn the inputs that do not correspond to bits of her committed x . However, in this case, she cannot learn a key for every input wire of the circuit. Therefore in this case she cannot compute the output of the scrambled circuit.

7. OUR CIPPE PROTOCOL

We now give the CIPPE protocol which follows Definition 1, and specify what each participant does in each step.

PROTOCOL 2 (CIPPE PROTOCOL). The CIPPE protocol involves a client C, a server S, and a trusted CA, and has the following four phases:

CA Setup CA takes a security parameter σ and a setup parameters ℓ as input. CA runs the Pedersen commitment setup algorithm to create $\text{Params} = \langle p, q, g, h \rangle$ such that $2^\ell < q$, and sends it to C and S.

Client-CA Interaction C chooses n integers $x_1, \dots, x_n \in [0..2^\ell - 1]$ and sends them to CA. For each x_i , $1 \leq i \leq n$, CA picks $r_i \in_R \mathbb{Z}_q$ and computes the commitment $c_i = (g^{x_i} h^{r_i} \bmod p)$. CA gives c_i and r_i to C, and c_i to S.

Client-Server Setup S takes three parameters ℓ , n , m , and λ as input, and outputs the family \mathcal{F} of functions as defined in Section 5. S sends the description of \mathcal{F} to C, then chooses a private function $f \in \mathcal{F}$.

Now S has c_1, \dots, c_n , and f . C has $c_1, \dots, c_n, x_1, \dots, x_n$, and r_1, \dots, r_n .

Client-Server Interaction The steps are as follows.

1. *Scrambling the circuit:* S constructs a circuit that computes the function f using the technique described in Section 5, then scrambles the circuit. S gives the scrambled circuit to C.
2. *Committing the output:* Let wire w_t denote the unique output wire of the scrambled circuit, and (k_t^0, k_t^1) denote the corresponding keys of w_t . S sends $\langle \eta_0 = E_{k_t^0} [0^\sigma], \eta_1 = E_{k_t^1} [1^\sigma] \rangle$ to C.
3. *Coding the input:* For each x_i where $1 \leq i \leq n$, there are ℓ corresponding input wires in the scrambled circuit. C and S run the CIOT protocol in which C inputs x_i, r_i , and c_i ; and S inputs c_i and ℓ pairs of keys that correspond to the ℓ input wires. In the end of this step, C learns one key per input wire; furthermore, each key corresponds to a bit in C's committed input.
4. *Evaluating the circuit:* After Step 3, C possesses enough information to evaluate the scrambled circuit independently. C evaluates the circuit and obtains k , the key of the output wire. Recall that C receives (η_0, η_1) from S in step 2, C tries to decrypt η_0 and η_1 using key k . If C fails in decrypting both of them, she outputs \perp and aborts; this happens only when S intentionally misbehave. If C succeeds in decrypting η_0 and gets 0^σ , she outputs 0. Otherwise, if C succeeds in decrypting η_1 and gets 1^σ , she outputs 1.
5. *Notifying the result:* C sends k to S, enabling S to output 0 if $k = k_t^0$ and output 1 if $k = k_t^1$.

The purpose of committing the output in step 2 is to achieve the fairness of the computation. The client and server need $O(\ell n)$ modular exponentiation and $O(\ell m \log n + \lambda m)$ symmetric key encryptions. More precisely, the server needs around $2\ell n$ modular exponentiation and the client needs around $4\ell n$ modular exponentiation, both the client and server need $(16 \log n + 40)\ell m + 8\lambda m$ symmetric key encryptions.

The CIPPE protocol is complete in the sense that if both C and S follow the protocol, C will get proper keys of the input wires, and will be able to evaluate the scrambled circuit correctly. Due to the space limit, we cannot give the formal security proofs of the CIPPE protocol here; the proofs are given in the full version of this paper [35]. We now briefly discuss the security properties of our CIPPE protocol and the intuitions underlying these properties.

- The CIPPE protocol is secure against honest-but-curious adversaries. This property follows from the fact that the scramble circuit protocol is secure in the honest-but-curious model [55].

An honest-but-curious server cannot learn any information about a client's private input, as the server is not able to obtain any information about the client's committed values from the CIOT protocol. An honest-but-curious client cannot learn any information about the a server's private function f , because the client learns only the topology of the circuit which is public information.

- The CIPPE protocol is secure against weak-honest adversaries. As we described in Section 6, for the client to evaluate the scrambled circuit, the client has to get one input value for each of her input wires to the circuit. Furthermore, if the client gets an input value for each of the input wires, these

input values must correspond to her committed attribute values. Therefore, the client can neither learn the server’s private function f nor change the result of the computation, as long as the server follows the protocol.

For a weak-honest server, it is not possible for her to build a circuit that has different topology than expected, because the client can detect such deviated behavior immediately. The server can build a circuit that computes a function other than f , we denote the function computed by the circuit as f' . Because f is the server’s private input, it is essentially same as the server inputs f' instead of f . In other words, such adversarial behavior in the real model can be simulated by the execution of computing f' in the ideal model.

- The CIPPE protocol is secure against a malicious client. However, a malicious server may learn (at most) one extra bit of information. The server can do this by constructing a scrambled circuit that would fail for some of the client’s input. While executing step 4 of Protocol 2, the client may fail during the evaluation of the circuit or fail to decrypt η_0 and η_1 . At the end of the protocol, the client may send k_t^0 , k_t^1 , or nothing to the server. Thus the server can classify the client’s input into one of three subsets determined by the circuit. When the server is honest, the server can classify the client’s input into only two subsets. In other words, a malicious server may learn less than one bit of extra information. However, the client would know that the server has cheated when the circuit evaluation fails.

In the example scenario we consider, if the bank is being detected to be dishonest by Alice, the bank may suffer from reputation as well as other kinds of damage. This small extra gain does not seem to warrant such malicious behavior in the kind of electronic commerce scenarios we consider.

8. RELATED WORK

Automated Trust Negotiation

Our work is closely related to a growing body of work on Automated Trust Negotiation (ATN) [53, 57, 56, 52] whose goal is to enable clients and servers to establish trust in each other through cautious, iterative, bilateral disclosure of sensitive certificates and policies. Recent works on using cryptographic protocols for ATN include Hidden Credentials [32, 8, 25], Secret Handshakes [3], and Oblivious Signature Based Envelope [36]. While these schemes are useful for scenarios where policies are based on attributes such as secret clearance or memberships in some secret underground movements, they are not suitable for the kind of e-commerce scenarios such as Example 1. Using any of these schemes, the server could send an encrypted message to a client such that the client can decrypt if and only if the client has certificates whose contents are the same as those identified by the server’s policy; at the same time, the server does not know whether the client has those certificates or not. These schemes can implement policy-hiding access control when the servers’ policies have very specific forms. In Example 1, if Bob’s loan approval policy is either Alice’s birthdate is April 1st, 1974 or Alice’s salary is exactly \$60,000, then policy-hiding access control can be achieved using these existing schemes. However, for the kind of policies in Example 1, where many possible attribute values would satisfy a policy, these schemes do not work well.

Secure Function Evaluation

Secure Function Evaluation (SFE) [55, 30, 28] is a powerful and general cryptographic primitive. It allows two or more parties to

jointly compute some function while hiding their inputs to each other. CIPPE may be cast as a special case of 2-SFE problem: (1) take x_1, \dots, x_n as Alice’s private input and make the commitment verification part of the public function, (2) treat the description of Bob’s private function f as part of Bob’s private input, and (3) make the public function to be evaluated a universal circuit that takes f ’s description, x_1, x_2, \dots, x_n , and c_1, c_2, \dots, c_n , and computes $f(x_1, \dots, x_n)$ (this universal circuit is similar in concept to a universal Turing Machine). However, applying the general solution in this case is very inefficient, as the circuit for verifying commitments and the universal circuit are very large.

The idea of committing local inputs before the function evaluation has appeared in the SFE literature (e.g., [17, 30, 27, 10]) to ensure the correctness of the computation. This concept is substantially different from our model, where Alice’s input is certified by a trusted third party in CIPPE. Cachin and Camenisch [10] introduced the notion of fair secure computation where a partially trusted third party T participates to ensure the fairness of the computation. Their work [10] is different from CIPPE in that (1) T ’s job is to achieve fairness instead of certifying Alice’s input (in fact, Alice’s input is even private to T), and (2) the function to compute is public.

Selective Private Function Evaluation was introduced by Canetti et al. [15] whose goal is for Bob to compute a private function $f(x_{i_1}, \dots, x_{i_m})$ over a subset of Alice’s database $x = x_1, \dots, x_n$ without revealing Bob’s function. In their, the authors focused on the case where f and m are public but the m locations in the database are private to Bob.

Abadi and Feigenbaum [1] introduced the notion of Secure Circuit Evaluation. In Secure Circuit Evaluation, Alice has a private input x and Bob has a private circuit C . In the end Alice learns the value $C(x)$ but nothing else about C . Sander et al. [48] improved the previous results and gave an efficient one-round protocol for secure evaluation of circuits that have polynomial size and depth $O(\log n)$. In these protocols, Alice can choose which input value to use in the circuit evaluation. It is not clear how Secure Circuit Evaluation protocols can be applied to CIPPE because the client’s input in CIPPE is committed and certified by a trust CA.

OACerts and Anonymous Credentials

Our work builds directly on OACerts developed in [34]. The ideas of storing commitments of attribute values in certificates appeared in the literature on anonymous credentials [16, 9, 12, 11]. Thus it is possible to replace OACerts and use anonymous credentials in policy-hiding access control and CIPPE. Note that using zero-knowledge proof protocols [19, 7, 9, 12] together with OACerts or anonymous credentials, Alice can prove that her attribute values in her certificates have certain properties without revealing any other information about her attributes. However, in order for Alice to use such techniques to prove that she satisfies Bob’s policy, she needs to know the policy. Therefore, zero-knowledge proof protocols are not suitable for policy-hiding access control.

Oblivious Transfer

Crépeau [20] introduced the notion of Committed Oblivious Transfer (COT). In COT, Bob commits two bits: a_0 and a_1 , and Alice commits a bit b . In the end, Alice learns a_b without learning anything else, while Bob learns nothing. Garay et al. [26] gave an efficient construction of COT in the universal composability framework. The CIOT protocol we propose in this paper differs from the COT protocols in that the receiver’s input in CIOT is a committed integer instead of a bit. Finally, the details of our CIOT protocol are reminiscent of the techniques used in the oblivious transfer proto-

cols [41, 51], zero-knowledge proofs of that a committed number belongs to an interval [40, 21], and anonymous fingerprinting [44].

9. CONCLUSION AND FUTURE WORK

We have presented an efficient and provably secure solution to policy-hiding access control, which enables Bob to decide whether Alice's certified attribute values satisfy Bob's policy, without Bob learning any other information about Alice's attribute values or Alice learning Bob's policy. Our approach uses OACerts and CIPPE. Our construction for CIPPE uses Yao's scrambled circuit protocol and two novel techniques, one is constructing topologically uniform circuits that can compute arbitrary functions in a function family, the other is a committed-integer based oblivious transfer protocol. Future work includes constructing efficient topologically uniform circuits for function families other than the one we studied in Section 5.

Acknowledgement

This work is supported by NSF ITR grant CCR-0325951 and by sponsors of CERIAS. We thank the anonymous reviewers for their helpful comments.

10. REFERENCES

- [1] Martín Abadi and Joan Feigenbaum. Secure circuit evaluation: A protocol based on hiding information from an oracle. *Journal of Cryptology*, 2(1):1–12, 1990.
- [2] S. Arora, T. Leighton, and B. Maggs. On-line algorithms for path selection in a nonblocking network. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pages 149–158. ACM Press, 1990.
- [3] Dirk Balfanz, Glenn Durfee, Narendar Shankar, Diana Smetters, Jessica Staddon, and Hao-Chi Wong. Secret handshakes from pairing-based key agreements. In *Proceedings of the IEEE Symposium and Security and Privacy*, pages 180–196, May 2003.
- [4] Mihir Bellare and Silvio Micali. Non-interactive oblivious transfer and applications. In *Advances in Cryptology: CRYPTO '89*, volume 435 of *Lecture Notes in Computer Science*, pages 547–557. Springer, 1989.
- [5] Matt Blaze, Joan Feigenbaum, and Jack Lacy. Decentralized trust management. In *Proceedings of the 1996 IEEE Symposium on Security and Privacy*, pages 164–173. IEEE Computer Society Press, May 1996.
- [6] Sharon Boeyen, Tim Howes, and Patrick Richard. Internet X.509 Public Key Infrastructure LDAPc2 Schema. IETF RFC 2587, June 1999.
- [7] Fabrice Boudot. Efficient proofs that a committed number lies in an interval. In *Advances in Cryptology: EUROCRYPT '00*, volume 1807 of *Lecture Notes in Computer Science*, pages 431–444. Springer, May 2000.
- [8] Robert Bradshaw, Jason Holt, and Kent Seamons. Concealing complex policies with hidden credentials. In *Proceedings of 11th ACM Conference on Computer and Communications Security*, October 2004.
- [9] Stefan A. Brands. *Rethinking Public Key Infrastructures and Digital Certificates: Building in Privacy*. MIT Press, August 2000.
- [10] Christian Cachin and Jan Camenisch. Optimistic fair secure computation. In *Advances in Cryptology: CRYPTO '00*, volume 1880 of *Lecture Notes in Computer Science*, pages 93–111. Springer, 2000.
- [11] Jan Camenisch and Els Van Herreweghen. Design and implementation of the idemix anonymous credential system. In *Proceedings of the 9th ACM Conference on Computer and Communications Security, CCS '02*, pages 21–30. ACM, nov 2002.
- [12] Jan Camenisch and Anna Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In *Advances in Cryptology: EUROCRYPT '01*, volume 2045 of *Lecture Notes in Computer Science*, pages 93–118. Springer, 2001.
- [13] Ran Canetti. Security and composition of multiparty cryptographic protocols. *Journal of Cryptology*, 13(1):143–202, 2000.
- [14] Ran Canetti, Uri Feige, Oded Goldreich, and Moni Naor. Adaptively secure multi-party computation. In *STOC '96: Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 639–648. ACM Press, 1996.
- [15] Ran Canetti, Yuval Ishai, Ravi Kumar, Michael K. Reiter, Ronitt Rubinfeld, and Rebecca N. Wright. Selective private function evaluation with applications to private statistics. In *Proceedings of the twentieth annual ACM symposium on Principles of distributed computing*, pages 293–304. ACM Press, 2001.
- [16] David Chaum. Security without identification: Transaction systems to make big brother obsolete. *Communications of the ACM*, 28(10):1030–1044, 1985.
- [17] David Chaum, Ivan Damgård, and Jeroen van de Graaf. Multiparty computations ensuring privacy of each party's input and correctness of the result. In *Advances in Cryptology: CRYPTO '87*, volume 293 of *Lecture Notes in Computer Science*, pages 87–119. Springer, 1988.
- [18] Dwaine Clarke, Jean-Emile Elien, Carl Ellison, Matt Fredette, Alexander Morcos, and Ronald L. Rivest. Certificate chain discovery in SPKI/SDSI. *Journal of Computer Security*, 9(4):285–322, 2001.
- [19] Ronald Cramer and Ivan Damgård. Zero-knowledge proof for finite field arithmetic, or: Can zero-knowledge be for free? In *Advances in Cryptology: CRYPTO '98*, volume 1462 of *Lecture Notes in Computer Science*, pages 424–441. Springer, 1998.
- [20] Claude Crépeau. Verifiable disclosure of secrets and applications (abstract). In *Advances in Cryptology: EUROCRYPT '89*, volume 434 of *Lecture Notes in Computer Science*, pages 150–154. Springer, 1990.
- [21] Glenn Durfee and Matt Franklin. Distribution chain security. In *Proceedings of the 7th ACM Conference on Computer and Communications Security*, pages 63–70. ACM Press, 2000.
- [22] Taher El Gamal. A public key cryptosystem and a signature scheme based on discrete logarithms. In *Advances in Cryptology: CRYPTO '84*, volume 196 of *Lecture Notes in Computer Science*, pages 10–18. Springer, 1985.
- [23] Carl Ellison, Bill Frantz, Butler Lampson, Ron Rivest, Brian Thomas, and Tatu Ylonen. SPKI certificate theory. IETF RFC 2693, September 1999.
- [24] Stephen Farrell and Russell Housley. An internet attribute certificate profile for authorization. IETF RFC 3281, April 2002.
- [25] Keith B. Frikken, Mikhail J. Atallah, and Jiangtao Li. Hidden access control policies with hidden credentials. In *Proceedings of the 3rd ACM Workshop on Privacy in the Electronic Society*, October 2004.
- [26] Juan Garay, Philip MacKenzie, and Ke Yang. Efficient and

- universally composable committed oblivious transfer and applications. In *Theory of Cryptography, TCC 2004*, volume 2951 of *Lecture Notes in Computer Science*, pages 297–316. Springer, 2004.
- [27] Rosario Gennaro, Michael O. Rabin, and Tal Rabin. Simplified vss and fast-track multiparty computations with applications to threshold cryptography. In *PODC '98: Proceedings of the seventeenth annual ACM symposium on Principles of distributed computing*, pages 101–111. ACM Press, 1998.
- [28] Oded Goldreich. Secure multi-party computation, October 2002.
- [29] Oded Goldreich. Cryptography and cryptographic protocols. *Distributed Computing*, 16(2-3):177–199, 2003.
- [30] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game. In *Proceedings of the nineteenth annual ACM conference on Theory of computing*, pages 218–229, May 1987.
- [31] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18:186–208, feb 1989.
- [32] Jason E. Holt, Robert W. Bradshaw, Kent E. Seamons, and Hilarie Orman. Hidden credentials. In *Proceedings of the 2nd ACM Workshop on Privacy in the Electronic Society*, October 2003.
- [33] Russell Housley, Warwick Ford, Tim Polk, and David Solo. Internet X.509 Public Key Infrastructure Certificate and CRL Profile. IETF RFC 2459, January 1999.
- [34] Jiangtao Li and Ninghui Li. OACerts: Oblivious attribute certificates. In *Proceedings of the 3rd Conference on Applied Cryptography and Network Security (ACNS)*, volume 3531 of *Lecture Notes in Computer Science*. Springer, June 2005.
- [35] Jiangtao Li and Ninghui Li. Policy-hiding access control in open environment. Technical Report CERIAS-TR-2005-31, Center for Education and Research in Information Assurance and Security, Purdue University, May 2005.
- [36] Ninghui Li, Wenliang Du, and Dan Boneh. Oblivious signature-based envelope. In *Proceedings of the 22nd ACM Symposium on Principles of Distributed Computing (PODC)*. ACM Press, July 2003.
- [37] Ninghui Li, John C. Mitchell, and William H. Winsborough. Design of a role-based trust management framework. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy*, pages 114–130. IEEE Computer Society Press, May 2002.
- [38] Ninghui Li, William H. Winsborough, and John C. Mitchell. Distributed credential chain discovery in trust management. *Journal of Computer Security*, 11(1):35–86, February 2003.
- [39] Dahlia Malkhi, Noam Nisan, Benny Pinkas, and Yaron Sella. Fairplay – secure two-party computation system. In *Proceedings of the 13th USENIX Security Symposium*, pages 287–302. USENIX, 2004.
- [40] Wenbo Mao. Guaranteed correct sharing of integer factorization with off-line shareholders. In *Public Key Cryptography: PKC'98*, volume 1431 of *Lecture Notes in Computer Science*, pages 60–71. Springer, February 1998.
- [41] Moni Naor and Benny Pinkas. Efficient oblivious transfer protocols. In *Proceedings of SODA 2001 (SIAM Symposium on Discrete Algorithms)*, pages 448–457, January 2001.
- [42] Moni Naor, Benny Pinkas, and Reuban Sumner. Privacy preserving auctions and mechanism design. In *Proceedings of the 1st ACM conference on Electronic commerce*, pages 129–139. ACM Press, 1999.
- [43] Ju P. Ofman. A universal automaton. *Transactions of the Moscow Math Society*, 14:200–215, 1965.
- [44] Joon S. Park and Ravi Sandhu. Smart certificates: Extending X.509 for secure attribute services on the web. In *Proceedings of the 22nd National Information Systems Security Conference*, October 1999.
- [45] Torben P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In *Advances in Cryptology: CRYPTO '91*, volume 576 of *Lecture Notes in Computer Science*, pages 129–140. Springer, 1991.
- [46] Nicholas J. Pippenger. Generalized connectors. Technical Report RC-6532, IBM Res. Rep., 1977.
- [47] Ronald L. Rivest and Bulter Lampson. SDSI — a simple distributed security infrastructure, October 1996. Available at <http://theory.lcs.mit.edu/~rivest/sdsi11.html>.
- [48] Tomas Sander, Adam Young, and Moti Yung. Non-interactive cryptocomputing for nc1. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science (FOCS)*, page 554. IEEE Computer Society, 1999.
- [49] Clark D. Thompson. Generalized connection networks for parallel processor intercommunication. *IEEE Transactions on Computers*, 27(12):1119–1125, December 1978.
- [50] Yiannis Tsiounis and Moti Yung. On the security of elgamal based encryption. In *Proceedings of the First International Workshop on Practice and Theory in Public Key Cryptography*, pages 117–134. Springer, 1998.
- [51] Wen-Guey Tzeng. Efficient 1-out-n oblivious transfer schemes. In *PKC '02: Proceedings of the 5th International Workshop on Practice and Theory in Public Key Cryptosystems*, number 2274 in *Lecture Notes in Computer Science*, pages 159–171. Springer, 2002.
- [52] William H. Winsborough and Ninghui Li. Safety in automated trust negotiation. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 147–160, May 2004.
- [53] William H. Winsborough, Kent E. Seamons, and Vicki E. Jones. Automated trust negotiation. In *DARPA Information Survivability Conference and Exposition*, volume I, pages 88–102. IEEE Press, January 2000.
- [54] Yuanyuan Yang and Gerald M. Masson. The necessary conditions for clos-type nonblocking multicast networks. *IEEE Transactions on Computers*, 48(11):1214–1227, 1999.
- [55] Andrew C. Yao. How to generate and exchange secrets. In *Proceedings of the 27th IEEE Symposium on Foundations of Computer Science*, pages 162–167. IEEE Computer Society Press, 1986.
- [56] Ting Yu and Marianne Winslett. Unified scheme for resource protection in automated trust negotiation. In *Proceedings of IEEE Symposium on Security and Privacy*, pages 110–122. IEEE Computer Society Press, May 2003.
- [57] Ting Yu, Marianne Winslett, and Kent E. Seamons. Supporting structured credentials and sensitive policies through interoperable strategies for automated trust negotiation. *ACM Transactions on Information and System Security (TISSEC)*, 6(1):1–42, February 2003.
- [58] Ellen Witte Zegura. Evaluating blocking probability in generalized connectors. *IEEE/ACM Transactions on Networking*, 3(4):387–398, 1995.