

# Attribute-based Key Exchange with General Policies\*

Vladimir Kolesnikov  
Bell Labs  
kolesnikov@research.bell-labs.com

Hugo Krawczyk  
IBM Research  
hugo@ee.technion.ac.il

Yehuda Lindell  
Bar-Ilan University  
yehuda.lindell@biu.ac.il

Alex J. Malozemoff<sup>†</sup>  
Galois  
amaloz@galois.com

Tal Rabin  
IBM Research  
talr@us.ibm.com

## ABSTRACT

Attribute-based methods provide authorization to parties based on whether their set of attributes (e.g., age, organization, etc.) fulfills a policy. In attribute-based encryption (ABE), authorized parties can decrypt, and in attribute-based credentials (ABCs), authorized parties can authenticate themselves. In this paper, we combine elements of ABE and ABCs together with garbled circuits to construct attribute-based key exchange (ABKE). Our focus is on an interactive solution involving a client that holds a certificate (issued by an authority) vouching for that client's attributes and a server that holds a policy computable on such a set of attributes. The goal is for the server to establish a shared key with the client but only if the client's certified attributes satisfy the policy. Our solution enjoys strong privacy guarantees for both the client and the server, including attribute privacy and unlinkability of client sessions.

Our main contribution is a construction of ABKE for *arbitrary circuits* with high (concrete) *efficiency*. Specifically, we support general policies expressible as boolean circuits computed on a set of attributes. Even for policies containing hundreds of thousands of gates the performance cost is dominated by two pairing computations per policy input. Put another way, for a similar cost to prior ABE/ABC solutions, which can only support small formulas efficiently, we can support *vastly* richer policies.

We implemented our solution and report on its performance. For policies with 100,000 gates and 200 inputs over a realistic network, the server and client spend 957 ms and 176 ms on computation, respectively. When using offline preprocessing and batch signature verification, this drops to only 243 ms and 97 ms.

\*The full version of this paper is available at <https://eprint.iacr.org/2016/518>

<sup>†</sup>Portion of work done while at University of Maryland and Bell Labs.

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

CCS'16, October 24-28, 2016, Vienna, Austria

© 2016 ACM. ISBN 978-1-4503-4139-4/16/10...\$15.00

DOI: <http://dx.doi.org/10.1145/2976749.2978359>

## 1. INTRODUCTION

The increasing need and complexity of authentication in the digital world, alongside ever growing privacy concerns, has given rise to encryption and authentication mechanisms that combine privacy aspects (anonymity, unlinkability, etc.) with credentials that go well beyond asserting an identity of a client but rather vouch for a full set of attributes (age, rank, role, etc.). These mechanisms allow for authentication and encryption that build on authorization policies computed on the provided set of attributes. For example, in a hospital setting, access to a patient's records can be provided to the patient, her doctor, nurses while on duty, or to the director of the hospital ward, and this can be formalized as a policy.

The prime examples of these mechanisms are attribute-based credentials (ABCs) [23, 17, 19, 20, 7, 3, 4] and attribute-based encryption (ABE) [36, 30, 9, 38]. The former is mainly directed towards identification and access control based on a set of attributes, and emphasizes privacy aspects such as anonymity, unlinkability, and attribute privacy. While ABCs generally assume an interactive setting between a client and a server, ABE focuses on (non-interactive) encryption where authorization is enforced through an encryption scheme that ties a ciphertext to a policy and a decryption key that ensures that only clients that have attributes that satisfy the policy can decrypt (note that we focus on the "ciphertext policy" setting of ABE [9]). Since ABE schemes are non-interactive (and the party decrypting is completely passive), they have certain implicit privacy properties of ABC solutions such as unlinkability and attribute privacy. An essential requirement of the above primitives is that of collusion resistance. This means that different clients of the system cannot combine their attributes in order to pass policy verification that neither could have individually passed.

In many practical settings simply communicating a message to a party or just establishing rights (such as validating a function of attributes) is not enough. For example, authentication to an online service is usually followed with further communications that also need to be protected. In other words, the goal in these systems is the use of credentials to bootstrap a key exchange protocol that provides the parties with keys to protect a session.

In this paper we combine elements of ABE and ABCs to build attribute-based key exchange (ABKE) where our focus is on an interactive solution involving a client that holds a certificate (issued by an authority CA) vouching for the client's attributes and a server that holds a policy com-

putable on the set of attributes. The goal is for the server to establish a shared key with the client if and only if the client's certified attributes satisfy the policy.<sup>1</sup>

Of course, the above goal is easy to achieve if the client is willing to reveal its attributes to the server. The objective of our work is to enable the ABKE functionality while keeping the attributes of the client private alongside ensuring additional important properties. The main features of our ABKE solution are summarized next.

*General policies.* We support any policy expressible as a polynomial-size boolean circuit computed on a set of attributes.

*Attribute privacy.* Client attributes are never disclosed. Of course, the server learns whether the key exchange succeeded and thus learns that the client's attributes fulfill the policy used in the exchange. However, nothing beyond this fact is revealed.

*Unlinkability.* Multiple communications with the same client (with one or more servers) cannot be linked together.

*Collusion resistance.* It is not possible for an adversary given keys associated with multiple clients with different attributes (certified by the CA) to succeed in an exchange in which no single client with its associated attributes fulfills the policy. In particular, attributes from different clients cannot be mixed-and-matched.

## 1.1 Overview of Our Solutions

Our main contributions are a definition (cf. §4) and realization (cf. §6) of attribute-based key exchange (ABKE) for public (circuit-based) policies.

**ABKE using garbled circuits.** Our construction uses garbled circuits in order to achieve ABKE. The use of garbled circuits enables us to obtain a solution that both supports arbitrarily-complex policies (without requiring heavy machinery like multilinear maps or fully homomorphic encryption) and is concretely efficient. In our approach, the server generates a garbled circuit and sends it to the client. The client then obtains the garbled values on the input wires of the circuit, depending on its attributes. This is achieved by encrypting the garbled values on the input wires using a type of encryption that enables the client to decrypt only those values associated with its attributes. We call this notion *attribute selective encryption (ASE)* (cf. §5). The main technical difficulty comes with ensuring that the client obtains input labels corresponding to its credentials in a *private*, *unlinkable*, and *collusion-free* manner. At a high level, we construct such an encryption scheme using a rerandomizable set of public keys and a rerandomizable signature binding the public keys together. The client then presents a set of rerandomized keys (along with a signature on them), and the server encrypts the garbled labels knowing that the client can only decrypt the appropriate set. We introduce and utilize the notion of *extractable linearly homomorphic (ELH) signatures* (cf. §7) to construct two instantiations of ASE: one based on identity-based encryption (cf. §8) and the other

built directly from ELH signatures (cf. §9). The *extractability* requirement ensures that a simulator can extract the *original* message that was signed, even though the adversary presents a *randomized* message. We prove this extractability property using the knowledge-of-exponent (KEA) assumption.

Our use of garbled circuits is a careful adaptation of the zero-knowledge-using-garbled-circuits approach of Jawurek et al. [34]. As shown in their work, we can use a *single* garbled circuit while still achieving malicious security; this is discussed in more detail in §6.

**Concrete performance.** At a cost similar to that of prior ABE/ABC solutions, which only run efficiently on (small) formulas, we can support vastly richer policies represented by large circuits. Specifically, we instantiate our construction over bilinear groups requiring a number of pairings proportional to the number of *input* attributes to the policy circuit. Then a garbled circuit computation of the policy circuit is performed with cost that is not noticeable for policies of even relatively large circuit size<sup>2</sup>.

To directly measure the performance of our scheme, we implemented it and ran various experiments; see §10. For example, in our implementation, the server and client computation time for a 1,000-gate policy and 10 attributes is **67 ms** and **11 ms**, respectively; for a 100,000-gate policy and 200 attributes the times are only **957 ms** and **176 ms**. We also note that much of the computation can be moved offline and we can use batch signature verification on the server side. Again with a 100,000-gate policy and 200 attributes, this optimized time is only around **243 ms** for the server, when assuming the server is batching ten messages in its signature verification, and **97 ms** for the client.

**Additional features.** Our construction can be easily extended to provide additional useful features, as detailed below.

- *Credential expiration*, by having attributes encode the expiration date.
- *Delegation of attributes*. This follows directly from the projectability property of our ASE definition (cf. §5).
- *Multi-authority*. This can be achieved generically by having credentials from different CAs encode, as a sequence of attributes, a unique certified serial number which is verified to be the same during ABKE. A more efficient alternative is offered by our ELH-based ASE construction (cf. §9) by using a common value  $u$  in the clients' public keys in lieu of a unique serial number.
- *Unlinkability with respect to CA*. Our IBE-based construction for ASE (cf. §8) provides information theoretic unlinkability, which implies unlinkability even against a colluding server and CA. Such unlinkability is also achieved by our ELH-based solution provided that the public key components  $g$  and  $h$  are generated jointly between the client and CA.

**Future work.** In this work we consider public policies only. However, our techniques can be used to provide some notion of *private* policies and we leave this for future work. Likewise, our focus here has been on achieving *practical efficiency*, and we achieve this using the KEA assumption and

<sup>1</sup>Note that we focus on the client-server setting where the client authenticates to the server. Server authentication usually happens with regular public key certificates that identify the server and can use standard tools such as TLS. Extensions of our system to the mutual authentication setting are possible but not treated here.

<sup>2</sup>For example, we can garble (resp., evaluate) an AND gate in roughly 46 (resp., 28) cycles per gate using privacy-free garbled circuits [25, 39].

the random-oracle model. The goal of achieving comparable efficiency under standard assumptions only and without a random oracle is important and we leave it for future work.

## 2. RELATED WORK

Our ABKE notion relates to ciphertext-policy attribute-based encryption (CP-ABE) and attribute-based credentials (ABCs). CP-ABE gives rise to a single-message key exchange (KE) solution in which a session key is encrypted under ABE and hence is implicitly authenticated by clients that can decrypt. Since the same key is distributed to any client with a set of attributes satisfying the policy, multiple clients may share the same key. This is the solution proposed by Gorantla et al. [28], who provide a game-based definition of attribute-based authenticated key exchange (under the abbreviation AB-AKE) and note that such a scheme is more in line with group key exchange than standard AKE.

ABE-based AKE requires several public-key operations per gate of the policy formula. Recent solutions to ABE for general circuits [26, 29], while sufficient to show feasibility, are mainly of theoretical interest due to the use of heavy underlying primitives. By using garbled circuits, our protocol costs are dramatically lower than either of the above ABE-based solutions.

Since most key exchange settings allow for interaction (the session that they protect is in itself typically interactive), our work leverages interaction to improve policy expressiveness as well as performance. In this sense we are closer to the ABC setting, where clients own the credentials they use in an interaction with a verifier. Our work inherits many of the challenges of ABCs, particularly in the area of client privacy, with properties such as attribute-privacy and unlinkability being central to our work. We note, however, two important differences.

First and foremost, prior ABC protocols and systems focus on (but are not limited to) small *formula-based policies* [23, 17, 19, 20, 7, 3, 4, 1] due to the high cost of needing several public-key operations per gate. Besides the cost, difficulty of policy design and analysis of non-trivial hand-generated small formulas is the reason that today's deployed systems mainly implement conjunction policies. In this work, we dramatically increase the computation power of the policy by enabling its implementation via garbled circuits. We believe that in addition to improving efficiency of existing ABC use cases, our work enables a much larger application scope for ABCs, due to the ability to run (large) policies auto-generated from easy-to-understand high-level code.

Secondly, the ABC literature focuses on verification of credentials and not on bootstrapping an authenticated session. In general, the ability to verify client credentials (i.e., a yes/no result) is insufficient for authenticating a session, even if the communication is carried over a server-authenticated channel (e.g., TLS). The relationship between credentials and key exchange is explicitly studied by Camenisch et al. [18], but their implementations do not cover rich policies, and do not outperform ABCs.

ABCs provide several practical features which we regard as future work, such as credential revocation and CA-verifier collusion. Other features, such as non-boolean and multi-authority credentials can be easily and cheaply built within our system (cf. §1.1).

In a concurrent and independent work, Chase et al. [22]

approach the problem of ABCs for non-boolean attributes by relying on garbled circuits to represent policies and, as a consequence, allow general circuit-based policies. The method of delivery of wire labels to the prover (in our notation, the client) is indeed the technical core of both of our approaches. Chase et al. allow the prover to enter arbitrary inputs to the garbled circuit, requiring a zero-knowledge proof that its garbled circuit inputs are consistent with arithmetic committed values, which, in turn, are consistent with the credential vector on which there exists a valid CA signature. This results in a number of exponentiations per boolean attribute, even in cases where a small subset of them are used in the policy. Chase et al. offer an alternative algorithm to reduce the number of public-key operations at the expense of message authentication code computation inside the garbled circuit, which introduces a significant communication overhead but may be a worthy trade-off for provers with many attributes. This approach, too, scales with the total number of attributes. In contrast, in our approach the client needs to only compute public-key operations per *policy* attribute (rather than over all the client's attributes as is required by Chase et al.), which may be significantly faster in many settings. However, our improved performance is a trade-off for using stronger assumptions. Additionally, we present the first implementation of general circuit ABKE (and hence ABCs), and report on its concrete performance. Finally, the construction of [22] does not support delegation, and it is not immediately clear how to enable it there.

Finally, Sakai et al. [37] very recently proposed attribute-based signatures for circuits based on bilinear maps. In their setting, only signers satisfying a certain policy on their attributes could successfully sign a message. Their scheme could be a basis for an ABC solution; however, they require several public-key operations and about 1 Kb of data sent per circuit gate; our garbled circuit-based solution is much more efficient (16 bytes and several symmetric key operations per circuit gate).

## 3. PRELIMINARIES

Let  $P_1, \dots, P_\ell$  and  $S_1, \dots, S_t$  be the set of clients and servers, respectively, and let  $A = \{a_1, \dots, a_m\}$  be the universe of all possible attributes. We associate an  $m$ -bit string  $\chi_i = \chi_i[1] \cdots \chi_i[m] \in \{0, 1\}^m$  with each  $P_i$  such that  $\chi_i[j] = 1$  if and only if  $P_i$  has attribute  $a_j$ . A policy is a (polynomial sized) circuit  $C$  with  $m$  inputs and a single-bit output. We say that  $P_i$  satisfies policy  $C$  if and only if  $C(\chi_i) = 1$ .

**Garbled circuits.** One of our main building blocks is garbled circuits. As the circuit description is public and only one party has input, we can utilize privacy-free garbled circuits [25], which are more efficient than standard garbled circuits. We use the garbled circuit notation of Bellare et al. [8], with one function (verification) introduced by Jawurek et al. [34]. We only consider circuits with a single bit of output, as this is all that is needed in our setting.

We define a *verifiable garbling scheme* by a tuple of functions  $\mathcal{G} = (\text{Gb}, \text{Ev}, \text{Ve})$  with each function defined as follows:

- *Garbling* algorithm  $\text{Gb}(1^n, C)$ : A randomized algorithm which takes as input the security parameter and a circuit  $C : \{0, 1\}^m \rightarrow \{0, 1\}$  and outputs a tuple of strings  $(\text{GC}, \{X_j^0, X_j^1\}_{j \in [m]}, \{Z^0, Z^1\})$ , where GC is the garbled circuit, the values  $\{X_j^0, X_j^1\}_{j \in [m]}$  denote

the input-wire labels, and the values  $\{Z^0, Z^1\}$  denote the output-wire labels.

- *Evaluation* algorithm  $\text{Ev}(\text{GC}, \{X_j\}_{j \in [m]})$ : A deterministic algorithm which evaluates garbled circuit GC on input-wire labels  $\{X_j\}_{j \in [m]}$ .
- *Verification* algorithm  $\text{Ve}(\text{C}, \text{GC}, \{X_j^0, X_j^1\}_{j \in [m]})$ : A deterministic algorithm which takes as input a circuit C, garbled circuit GC, and input-wire labels  $\{X_j^0, X_j^1\}_{j \in [m]}$ , and outputs **accept** if GC is a valid garbling of C and **reject** otherwise.

A verifiable garbling scheme must satisfy three security properties: (1) *correctness*, (2) *authenticity*, and (3) *verifiability*. The definitions for correctness and authenticity are standard: correctness enforces that a correctly garbled circuit, when evaluated, outputs the correct output of the underlying circuit; authenticity enforces that the evaluator can only learn the output label that corresponds to the value of the function. *Verifiability* [34] allows one to check that the garbled circuit indeed implements the specified plaintext circuit C.

***t*-KEA assumption.** We recall the *t*-KEA assumption used in our implementation of extractable linearly homomorphic signature from §7. The assumption was formulated in [11, 10]. See these papers and [32] for a good discussion and further references related to this assumption and its recent use. See also [5, 31] for a proof of security for *t*-KEA in the generic (bilinear) group model. The formulation below is simplified by not including an auxiliary input that, if present, is the same for both algorithms *E* and *E'*. We will use the plain acronym KEA when referring to the 1-KEA assumption.

**DEFINITION 3.1.** (*t*-KEA [11, 10]) *Let  $G$  be a cyclic group of prime order  $q$ . Consider algorithms that on input  $t$  random elements  $g_1, \dots, g_t$  in  $G$  and  $t$  values  $g_1^x, \dots, g_t^x$  for  $x \in_R \mathbb{Z}_q$ , output a pair  $(f, f')$  in  $G^2$ . Such an algorithm  $E$  is said to be a *t*-KEA algorithm if with non-negligible probability (over the choice of inputs to  $E$  and  $E$ 's random coins)  $E$  outputs  $(f, f')$  such that  $f' = f^x$ . We say that the ***t*-KEA assumption holds over  $G$**  if for every efficient *t*-KEA algorithm  $E$  in  $G$  there exists another efficient algorithm  $E'$  for which the following property holds except for a negligible probability: Let  $g_1, \dots, g_t, g_1^x, \dots, g_t^x$  be an input to  $E$  and  $\rho$  a vector of random coins for  $E$  on which  $E$  outputs  $(f, f' = f^x)$  then on the same inputs (and random coins)  $E'$  outputs a vector  $(f, f' = f^x, x_1, \dots, x_n)$  such that  $f = g_1^{x_1} \cdots g_n^{x_n}$ .*

**Auxiliary functionalities.** Our construction makes use of two (standard) functionalities for commitments ( $\mathcal{F}_{\text{com}}$ ) and secure coin-tossing ( $\mathcal{F}_{\text{cointoss}}$ ).

**Anonymous channels.** Our protocol assumes that the parties interact over anonymous channels. In practice, the anonymity provided by the network used by the clients is the level of anonymity that they achieve. For the purpose of proving security, we assume a *perfect* anonymous channel. In the simple-UC framework [21], which we use in this work, all messages to and from functionalities have **public headers** consisting of the type of operation, and the **private content** itself; the public header is revealed to the adversary but not the private content. However, the adversary is always given

the identity of the party sending the message to the functionality and the identity of the party receiving the message from the functionality. Thus, in order to model anonymous channels, all parties must send and receive together. (This actually makes sense since in principle, an adversary who can view the entire network can break anonymity unless every party interacts in each round. Nevertheless, here we use this simply as a way to model the requirements.) The  $\mathcal{F}_{\text{anon}}$  functionality appears in Figure 3.1. In the functionality all parties send a message to all other parties in each round. Note that if a party has no message at all to send, or it only needs to send to some parties, then it can simply use an empty message. We stress again that in practice not all parties need to interact in each round; this is merely for the purpose of modeling.

## 4. SECURITY DEFINITION

All of our definitions and proofs are in the simple-UC (SUC) model [21]. As was shown in the aforementioned work [21], any protocol that is secure in the SUC framework is also secure in the full UC framework.

**Attribute-based key exchange.** We present a functionality  $\mathcal{F}_{\text{abke}}$  for attribute-based key exchange supporting attribute privacy, unlinkability, and collusion resistance. The functionality is initialized with a set of attribute vectors  $\{\chi_i\}$ , where  $\chi_i$  corresponds to the attribute vector of client  $P_i$ . The functionality begins by waiting for a message from a server  $S_j$  that contains a circuit C representing  $S_j$ 's policy. The functionality stores this information and broadcasts a notification to all parties  $P_1, \dots, P_\ell$  that a policy is available. Upon receiving a response by one of the parties, say,  $P_i$ , the functionality proceeds as follows. If  $C(\chi_i) = 1$ , the policy is satisfied and so the functionality forwards a random key  $k$  to both  $P_i$  and  $S_j$ . If  $C(\chi_i) = 0$ , then  $\mathcal{F}_{\text{abke}}$  sends  $\perp$  to both  $P_i$  and  $S_j$ . The full description of  $\mathcal{F}_{\text{abke}}$  can be found in Figure 4.1.

Attribute privacy is captured by the fact that  $S_j$  never receives the attribute vector  $\chi_i$  of client  $P_i$ . Collusion resistance is handled by the fact that each party's attribute vector is fixed upon functionality initialization and cannot be changed. Thus, parties cannot use any attribute vector that differs from their initial ones. This implies that collusions between parties to effectively use a different attribute vector are impossible. Finally, unlinkability follows since the functionality does not pass on the identity of the client  $P_i$  to the server  $S_j$  at any time. We note that we do not provide server anonymity in our definition, since it does not seem to be required for the ABKE setting. Thus, the server's identity is revealed in the functionality definition.

We also introduce a functionality  $\mathcal{F}_{\text{setup}}$  for providing each party with the keys used in our protocol construction; see Figure 4.2.

## 5. ATTRIBUTE SELECTIVE ENCRYPTION

We introduce the notion of **attribute selective encryption** (ASE). ASE is related to ABE in the sense that clients' keys and decryption capabilities are related to the attributes they possess. In ASE a plaintext is comprised of a *set* of messages, and a client's credentials determine *which subset* can be decrypted. In more detail, each client has an  $m$ -bit vector  $\chi \in \{0, 1\}^m$  representing a set of attributes:  $\chi[j]$  is set to 1 *if and only if* the client possesses the  $j$ th attribute. The

$\mathcal{F}_{\text{anon}}$  works with clients  $P_1, \dots, P_\ell$  as follows:

1. Upon receiving a message (**send**,  $\text{sid}$ ,  $P_i$ ,  $(m_1^i, \dots, m_\ell^i)$ ) from  $P_i$ ,  $\mathcal{F}_{\text{anon}}$  stores the message.
2. After  $\mathcal{F}_{\text{anon}}$  receives a **send** message from *every* client  $P_1, \dots, P_\ell$ ,  $\mathcal{F}_{\text{anon}}$  sends (**receive**,  $\text{sid}$ ,  $P_j$ ,  $(m_j^1, \dots, m_j^\ell)$ ) to client  $P_j$  for  $j = 1, \dots, \ell$ .

The public header of each message is (**send**,  $\text{sid}$ ,  $P_i$ ) and (**receive**,  $\text{sid}$ ,  $P_j$ ), respectively, for send and receive messages. The private contents is the vector of messages.

**Figure 3.1:** Anonymous communications functionality  $\mathcal{F}_{\text{anon}}$ .

$\mathcal{F}_{\text{abke}}$  runs with clients  $P_1, \dots, P_\ell$  with attribute vectors  $\chi_1, \dots, \chi_\ell \in \{0, 1\}^m$ , and servers  $S_1, \dots, S_t$ , and works as follows:

1. Upon receiving (**policy**,  $\text{sid}$ ,  $C$ ) from some  $S_j$ , where  $C$  is either a circuit  $C' : \{0, 1\}^m \rightarrow \{0, 1\}$  or  $\perp$ , send (**policy**,  $\text{sid}$ ,  $S_j$ ,  $C$ ) to all  $P_1, \dots, P_\ell$ . If  $C = \perp$  then halt, and otherwise store (**policy**,  $\text{sid}$ ,  $S_j$ ,  $C$ ).
2. Upon receiving (**exchange**,  $\text{sid}$ ,  $S_j$ ) from  $S_j$  and (**exchange**,  $\text{sid}$ ,  $S_j$ ,  $P_i$ ) from  $P_i$ , if some message (**policy**,  $\text{sid}$ ,  $S_j$ ,  $C$ ) is stored, then:
  - If  $C(\chi_i) = 1$  then choose  $k \in_R \{0, 1\}^n$  and send (**completed**,  $\text{sid}$ ,  $k$ ) to  $P_i$  and  $S_j$ .
  - If  $C(\chi_i) = 0$  then send (**completed**,  $\text{sid}$ ,  $\perp$ ) to  $P_i$  and  $S_j$ .
3. Upon receiving (**abort**,  $\text{sid}$ ) from  $\text{Sim}$ , clear any message (**policy**,  $\text{sid}$ ,  $S_j$ ,  $C$ ) that is stored, send (**abort**,  $\text{sid}$ ) to  $P_i$  and  $S_j$ , and halt.

The public header of each message is: (**policy**,  $\text{sid}$ ,  $S_j$ ,  $C$ ), (**exchange**,  $\text{sid}$ ,  $S_j$ ), and (**completed**,  $\text{sid}$ ); all other content is private.

**Figure 4.1:** Attribute-based key exchange functionality  $\mathcal{F}_{\text{abke}}$  with attribute privacy, unlinkability and collusion resistance.

client holds public and secret keys associated with  $\chi$ . Anyone can encrypt a set of  $2m$  messages  $\begin{pmatrix} x_{1,0} & \cdots & x_{m,0} \\ x_{1,1} & \cdots & x_{m,1} \end{pmatrix}$  under the client's public key, and ASE enforces an OT-like property where the client can decrypt using its secret key only one of each  $(x_{i,0}, x_{i,1})$ , depending on  $\chi[j]$ . That is, the client decrypts the messages  $x_{1,\chi[1]}, \dots, x_{m,\chi[m]}$ , and nothing else. We stress that ASE, unlike ABE, encrypts under a specific client's public key, and only that client can decrypt.

Besides the basic semantic security notion of ASE, we consider four additional properties: *attribute privacy*, *collusion resistance*, *unlinkability*, and *projectability*. Each property on its own is easy to achieve, yet the combination, especially that of collusion resistance and unlinkability, makes the construction challenging.

**Attribute privacy.** This property requires that the public key  $pk_\chi$  hides attribute vector  $\chi$ . That is, publishing  $pk_\chi$  does not reveal which attributes the client holds. The following trivial solution achieves attribute privacy: generate a set of  $2m$  public keys and define the secret key to consist of only one of the secret keys in each pair.

**Collusion resistance.** A set  $\mathcal{P}$  of clients with attribute vectors  $\mathcal{X} = \{\chi_i\}_{P_i \in \mathcal{P}}$  and corresponding keypairs must not be able to construct a keypair representing  $\chi \notin \mathcal{X}$  (or  $\chi$  representing a subset of attributes not implied by  $\mathcal{X}$  — cf. *projectability* below). Collusion resistance can be achieved by combining the trivial solution from above with a secure signature scheme; i.e., by providing a signature on the set of the client's public keys. This prevents clients from mixing and matching the individual keys in their public keys, giving collusion resistance.

**Unlinkability.** Unlinkability is the inability to link between different uses of the same public key. Specifically, we require that it be possible to *randomize* a public key using some algorithm **Unlink** so that the pair  $(pk_\chi, \text{Unlink}(pk_\chi))$  looks like two independent public keys. Without the requirement

of collusion resistance, unlinkability is easy to achieve (e.g., by using ElGamal keys). However, as we are interested in collusion resistance we thus need to enable the creation of a signature on the randomized key. For this we need to use *homomorphic signatures*. However, existing signature schemes do not provide the capabilities that are needed for our schemes. Thus, combining signatures with unlinkability is not straightforward.

**Projectability.** We require that given a keypair associated with a vector  $\chi$ , one can generate a keypair that is associated with any orthogonal projection of  $\chi$  onto some subset  $S \subseteq [m]$ . We stress that the subset  $S$  is *explicit* in the projection (otherwise, the encrypting party cannot know what the projection is, and this could be used to obtain unauthorized decryptions). As each public key needs to be certified, this implies that the certificate for the new key also needs to be derived from the certificate of  $pk_\chi$ .

**Committing encryption.** We require that ASE encryption is committing.

## 5.1 Formal Definition

Let  $n$  denote the security parameter and let  $m$  be the length of the attribute vector. We assume for simplicity that the client receives a public key on the entire attribute vector.

**DEFINITION 5.1.** An attribute selective encryption (ASE) scheme with *attribute privacy*, *collusion resistance*, and *unlinkability* is a tuple of probabilistic-polynomial time algorithms (**Setup**, **GenCert**, **Vrfy**, **Enc**, **Dec**, **Unlink**, **Project**) as follows:

- **Setup**( $1^n, m$ ) takes as input an attribute set size  $m$ , and outputs a master verification key and a master secret key ( $\text{mVK}, \text{mSK}$ ) along with public parameters **PP**. All the following algorithms implicitly take **PP** as input.

Upon initialization with length parameter  $m$ ,  $\mathcal{F}_{\text{setup}}$  runs  $(\text{PP}, \text{mVK}, \text{mSK}) \leftarrow \text{Setup}(1^n, m)$  and stores  $(\text{mVK}, \text{mSK})$ .

1. Upon receiving  $(\text{generate}, \text{sid}, \chi_i)$  from player  $P_i$ ,  $\mathcal{F}_{\text{setup}}$  checks if there exists a record  $(i, \cdot, \cdot, \cdot)$ . If so,  $\mathcal{F}_{\text{setup}}$  sends  $(\text{result}, \text{sid}, \perp)$  to  $P_i$ . Otherwise,  $\mathcal{F}_{\text{setup}}$  runs  $(pk, sk) \leftarrow \text{GenCert}(\text{mSK}, \chi_i)$ , records  $(i, \chi_i)$  and sends  $(\text{result}, \text{sid}, \text{mVK}, pk, sk)$  to  $P_i$ .

The public header of each message is:  $(\text{generate}, \text{sid})$  and  $(\text{result}, \text{sid})$ ; all other content is private.

**Figure 4.2:** Setup functionality  $\mathcal{F}_{\text{setup}}$ .

- $\text{GenCert}(\text{mSK}, \chi)$  takes as input the master secret key and attribute vector  $\chi \in \{0, 1\}^m$ , and outputs a certified keypair  $(pk_\chi, sk_\chi)$  associated with  $\chi$ .
- $\text{Vrfy}(\text{mVK}, pk_\chi)$  takes as input the master verification key and a public key  $pk_\chi$ , and outputs 1 if and only if  $pk_\chi$  is a valid public key.
- $\text{Enc}(pk_\chi, \vec{x})$  takes as input a public key  $pk_\chi$ , and a vector  $\vec{x}$ , where  $\vec{x} = \begin{pmatrix} x_{1,0} & \cdots & x_{m,0} \\ x_{1,1} & \cdots & x_{m,1} \end{pmatrix}$  is a series of  $2m$  messages. The function outputs an encryption  $c$ . For simplicity, we assume that each  $x_{i,b}$  is of length  $n$  (this suffices for our use).
- $\text{Dec}(sk_\chi, c)$  takes as input a secret key  $sk_\chi$  and a ciphertext  $c$ , and outputs a set of  $m$  plaintexts based on  $\chi$ .
- $\text{Unlink}(pk_\chi, sk_\chi)$  takes as input a public key  $pk_\chi$  and its associated private key  $sk_\chi$ , and outputs a new keypair  $(pk'_\chi, sk'_\chi)$  for the same  $\chi$ .
- $\text{Project}(pk_\chi, sk_\chi, S)$  takes as input a public key  $pk_\chi$ , its associated secret key  $sk_\chi$ , and a set  $S \subseteq \{0, 1\}^m$  which defines  $\chi'$  by specifying which attributes of  $\chi$  are to be preserved.  $\text{Project}$  outputs a keypair  $(pk'_{\chi'}, sk'_{\chi'})$  on the projected attribute vector  $\chi'$ .

We require the following properties on the algorithms:

- (Correctness) For every  $(\text{PP}, \text{mVK}, \text{mSK}) \leftarrow \text{Setup}(1^n, m)$ ,  $\chi \in \{0, 1\}^m$ ,  $(pk_\chi, sk_\chi) \leftarrow \text{GenCert}(\text{mSK}, \chi)$ , and  $\vec{x}$ , it holds that  $\text{Dec}(sk_\chi, \text{Enc}(pk_\chi, \vec{x})) = (x_{1,\chi_1}, \dots, x_{m,\chi_m})$ .
- For every  $(pk_\chi, sk_\chi)$ , the output of  $\text{Unlink}(pk_\chi, sk_\chi)$  is distributed identically to the output of  $\text{GenCert}$ .
- For every  $(\text{PP}, \text{mVK}, \text{mSK}) \leftarrow \text{Setup}(1^n, m)$ ,  $\chi \in \{0, 1\}^m$ , and  $(pk_\chi, sk_\chi) \leftarrow \text{GenCert}(\text{mSK}, \chi)$ , it holds that  $\text{Vrfy}(\text{mVK}, pk_\chi) = 1$ .
- The algorithm  $\text{Enc}$  is a committing encryption scheme.

Finally, we require the existence of the following two algorithms, which are used in our security definitions:

- $\text{GenCert}^*(\text{mSK})$  takes as input the master secret key and outputs a certified keypair  $(pk, sk)$  associated with both the 0 and 1 value of each attribute.
- $\text{Dec}^*(sk, c)$  takes as input a secret key  $sk$  generated by  $\text{GenCert}^*$  and a ciphertext  $c$ , and outputs the full set of  $2m$  plaintexts.

We call an ASE scheme *projectable* if:

- For every  $(\text{PP}, \text{mVK}, \text{mSK}) \leftarrow \text{Setup}(1^n, m)$ ,  $\chi \in \{0, 1\}^m$ ,  $(pk_\chi, sk_\chi) \leftarrow \text{GenCert}(\text{mSK}, \chi)$ ,  $S \subseteq \{0, 1\}^m$ , the output of  $\text{Project}(pk_\chi, sk_\chi, S)$  is distributed according to  $\text{GenCert}(\text{mSK}, \chi')$  for  $\chi'$  derived according to  $S$ .
- Correctness holds for every projected attribute vector.

Having defined the syntax, we now define security. We define this via experiments between a challenger  $\mathcal{C}$  and an adversary  $\text{Adv}$  for an ASE scheme  $\pi$ .

**Collusion resistance.** Our collusion resistance experiment guarantees that players can only obtain decryptions authorized by their attribute vectors. The adversary  $\text{Adv}$  is given oracle access to  $\text{GenCert}$  in order to model  $\text{Adv}$  corrupting multiple parties and learning their attribute vectors. Eventually,  $\text{Adv}$  sends a public key to  $\mathcal{C}$ , who responds with a random plaintext  $\vec{x}$  encrypted under this public key. The adversary  $\text{Adv}$  responds with a set of potential plaintext messages. If some subset of this set corresponds to an attribute vector (or any of its projections) that were *not* queried by  $\text{Adv}$  to  $\text{GenCert}$ , then  $\text{Adv}$  wins.

The reason we need to define collusion resistance in this way is that when proving security of our ABKE scheme, we extract the plaintext through the adversary's calls to the random oracle. Namely, the plaintext messages  $x_{i,b}$  are input into the random oracle by the adversary. However, the adversary is not limited to just inputting the proper messages to the random oracle, and thus we need to consider the set of *all* queries to the random oracle, a subset of these which may contain the extracted plaintext.

Note that it is easy for the challenger to check whether such a subset exists as follows. It checks whether each message in  $\mathcal{M}$  is a valid plaintext message  $x_{i,b}$ . Given this set of valid plaintext messages, the challenger can extract an attribute vector (based on the  $(i, b)$  values) and check whether such an attribute vector is unauthorized as per the definition.

**Experiment  $\text{Expt}_{\pi, \text{Adv}}^{\text{collude}}(1^n, m)$ :**

1.  $\mathcal{C}$  computes  $(\text{PP}, \text{mVK}, \text{mSK}) \leftarrow \text{Setup}(1^n, m)$  and sends  $\text{PP}$  and  $\text{mVK}$  to  $\text{Adv}$ .
2.  $\text{Adv}$ , with oracle access to  $\text{GenCert}(\text{mSK}, \cdot)$ , outputs a public key  $pk$ . Let  $\mathcal{X}$  be the set of attribute vectors  $\text{Adv}$  used as input to its oracle.
3.  $\mathcal{C}$  chooses a random plaintext  $\vec{x}$ , as specified by the ASE syntax, and sends  $\text{Enc}(pk, \vec{x})$  to  $\text{Adv}$ .
4.  $\text{Adv}$  outputs a set  $\mathcal{M}$  of potential plaintext messages.
5. The output of the experiment is 1 (and  $\text{Adv}$  wins) *if and only if* the following conditions all hold:
  - (a)  $\text{Vrfy}(\text{mVK}, pk) = 1$ ;
  - (b) There exists some subset  $\mathcal{M}' \subseteq \mathcal{M}$  such that either (1) the strings in  $\mathcal{M}'$  correspond to  $\{x_{i,\chi[i]}\}_{i \in [m]}$  for some attribute vector  $\chi$ , or (2) there exist two strings  $s, s' \in \mathcal{M}'$  such that  $s = x_{i,0}$  and  $s' = x_{i,1}$  for some  $i \in [m]$ .
  - (c)  $\chi \notin \mathcal{X}$ , and  $\chi$  is not a projection of any vector from  $\mathcal{X}$ .

Note that  $\text{Adv}$  does not have oracle access to  $\text{GenCert}$  after Step 3. This models the fact that  $\mathcal{F}_{\text{abke}}$  assumes a static setup after which the clients and their attributes are fixed.

**Attribute privacy.** We now consider an adversary who aims to infer  $\chi$  from  $pk_\chi$ . This follows a standard indistin-

guishability-based formulation. At a high level, the adversary is trying to distinguish a public key generated for some attribute vector  $\chi$  with an “all-powerful” public key generated by  $\text{GenCert}^*$ . Note that the inability to distinguish these two settings implies the inability to distinguish between any two attribute vectors by a simple hybrid argument.

**Experiment**  $\text{Expt}_{\pi, \text{Adv}}^{\text{att-priv}}(1^n, m)$ :

1.  $\mathcal{C}$  computes  $(\text{PP}, \text{mVK}, \text{mSK}) \leftarrow \text{Setup}(1^n, m)$  and sends PP and mVK to Adv.
2. Adv, with oracle access to  $\text{GenCert}(\text{mSK}, \cdot)$ , sends attribute vector  $\chi \in \{0, 1\}^m$  to  $\mathcal{C}$ .
3.  $\mathcal{C}$  chooses  $b \in_R \{0, 1\}$ , and proceeds as follows:
  - If  $b = 0$ , compute  $(pk, sk) \leftarrow \text{GenCert}(\text{mSK}, \chi)$  and send  $pk$  to Adv.
  - If  $b = 1$ , compute  $(pk, sk) \leftarrow \text{GenCert}^*(\text{mSK})$  and send  $pk$  to Adv.
4. Adv outputs a bit  $b'$ .
5. The output of the experiment is 1 (and Adv wins) if and only if  $b' = b$ .

**Unlinkability.** Finally, we define an experiment to formalize the property of unlinkability. The definition is relatively weak in that we only need to prevent an adversary from determining whether a keypair has been run through **Unlink** or not. However, this is sufficient for our purposes. In particular, unlinkability of keys used in our ABKE protocols will hold due to the conjunction of the guarantees of both attribute privacy and unlinkability.

**Experiment**  $\text{Expt}_{\pi, \text{Adv}}^{\text{link}}(1^n, m)$ :

1.  $\mathcal{C}$  computes  $(\text{PP}, \text{mVK}, \text{mSK}) \leftarrow \text{Setup}(1^n, m)$  and sends PP and mVK to Adv.
2. Adv with oracle access to  $\text{GenCert}(\text{mSK}, \cdot)$  eventually sends  $\chi$  to  $\mathcal{C}$ .
3.  $\mathcal{C}$  computes  $(pk_0, sk_0) \leftarrow \text{GenCert}(\text{mSK}, \chi)$  and  $(pk_1, sk_1) \leftarrow \text{Unlink}(pk_0, sk_0)$ .  $\mathcal{C}$  chooses  $b \in_R \{0, 1\}$  and sends  $(pk_b, sk_b)$  to Adv.
4. Adv outputs  $b'$ .
5. The output of the experiment is 1 (and Adv wins) if and only if  $b' = b$ .

Note that we cannot simply set **Unlink** to the identity function as we need the output distribution of **Unlink** to be the same as that of **GenCert**, as required in Definition 5.1.

**Security definition.** We are now ready to define security.

**DEFINITION 5.2.** A (projectable) attribute selective encryption scheme  $\pi$  with *attribute privacy*, *collusion resistance*, and *unlinkability* is **secure** if for every probabilistic-polynomial time adversary Adv there exists a negligible function  $\mu$  such that for every  $n$  and every  $Y \in \{\text{att-priv}, \text{link}\}$  it holds that

$$\Pr \left[ \text{Expt}_{\pi, \text{Adv}}^Y(1^n, m) = 1 \right] \leq \frac{1}{2} + \mu(n)$$

and

$$\Pr \left[ \text{Expt}_{\pi, \text{Adv}}^{\text{collude}}(1^n, m) = 1 \right] \leq \mu(n).$$

**ASE Instantiations.** We present two schemes realizing Definition 5.2 in §8 and §9.

## 6. ABKE FROM ASE

We now construct ABKE for public policies by integrating ASE with garbled circuit-based zero-knowledge proofs [34] and key agreement. Jawurek et al. [34] observed that for zero-knowledge proofs, the verifier-constructed circuit may be opened to the prover post-evaluation since it has no private data. Carefully arranging the prover’s and verifier’s commitments and openings, they ensure that neither can cheat, and only a single garbled circuit needs to be garbled, sent, and evaluated. Specifically, their protocol proceeds by the server using a *sender-committing* oblivious transfer (OT) to transfer the input-wire labels to the client. Given the garbled circuit and input-wire labels, the client can evaluate the garbled circuit and commit the output-wire label to the server. Now, the server can decommit to its inputs of the OT, allowing the client to verify that the garbled circuit was constructed correctly. If so, the client can open the commitment to its output and the server can verify that the client indeed computed the correct output-wire label.

We adapt this protocol to realize  $\mathcal{F}_{\text{abke}}$  by replacing sender-committing OT with ASE. That is, instead of the parties running the OT step in Jawurek et al.’s protocol, the client sends its (randomized) ASE public key to the server, who encrypts each input-wire label of the garbled circuit with ASE, guaranteeing that the client is only able to decrypt labels corresponding to its attribute vector. Next, the client evaluates the garbled circuit and commits to the output-wire label it computed. The server can then open all the encrypted values, allowing the client to verify the circuit was correctly garbled (before the client reveals anything). If the circuit is correct, the client decommits the output-wire label it computed, allowing the server to verify that indeed the client satisfied the policy. The parties then run a secure coin-tossing protocol to derive the shared key. See Figure 6.1 for the full protocol description.

**Garbling scheme for ABKE.** For us to successfully reduce to the collusion experiment in our ASE definition, we need to extract the plaintext from a malicious client to feed to the challenger in the experiment. This plaintext corresponds to the input-wire labels of the garbled circuit. Thus, we need some way to do this extraction. We do this by using a random oracle: we can monitor the inputs to the random oracle and use these as “potential” plaintexts which we can feed to the challenger in the collusion experiment. Thus, we construct a simple modified garbling scheme which allows us to do this extraction.

**ABKE garbling scheme**  $\mathcal{G}_{\text{ABKE}} = (\text{Gb}_{\text{ABKE}}, \text{Ev}_{\text{ABKE}}, \text{Ve}_{\text{ABKE}})$  **for a circuit  $C$  with  $m$  inputs.**

$\text{Gb}_{\text{ABKE}}$  is defined as follows.

1. Generate  $2m$  random labels  $X_i^b \in_R \{0, 1\}^n$ , where  $X_i^0$  and  $X_i^1$  correspond to input wire  $i$ . The set of all  $X_i^b$  form the input-wire labels to be encrypted. Let  $h_i^b := \text{RO}(i \| b \| X_i^b)$ , where RO is a random oracle.
2. Using any secure garbling scheme Gb, generate the garbled circuit, including the  $2m$  input-wire labels  $W_i^b$ .
3. Append to the generated garbled circuit the following input-wire translation tables: for wire  $i$ , append  $\begin{pmatrix} \text{Enc}_{h_i^0}(W_i^0) \\ \text{Enc}_{h_i^1}(W_i^1) \end{pmatrix}$ . Set the input-wire labels to be the set  $\{X_i^b\}$  and the output-wire labels to be those set by Gb.

The protocol  $\Pi_{\text{abke}}$  is between server  $S$  and client  $P$  with attribute vector  $\chi$ . We assume a setup where each client  $P_i$  sends  $(\text{generate}, \text{sid}, S, \chi_i)$  for some set  $S \subseteq [m]$  and attribute vector  $\chi_i$  to  $\mathcal{F}_{\text{setup}}$ , receiving  $(\text{result}, \text{sid}, \text{mVK}, pk, sk)$  in response. We assume all messages are sent/received through the  $\mathcal{F}_{\text{anon}}$  functionality; for simplicity we omit the use of this functionality in the description below. Also for simplicity, we assume the evaluated policy  $C$  uses all  $m$  attributes (otherwise, a corresponding Project operation can be applied to the party's key).

1.  $S$  broadcasts policy circuit  $C$  to all parties. If  $C$  is not a valid policy then  $P$  outputs  $\perp$ .
2.  $S$  runs  $(GC, \{X_j^0, X_j^1\}_{j \in [m]}, \{Z^0, Z^1\}) \leftarrow \text{Gb}_{\text{ABKE}}(1^n, C)$ .
3.  $P$  computes  $(pk', sk') \leftarrow \text{Unlink}(pk, sk)$  and sends  $pk'$  to  $S$ .
4.  $S$  runs  $\text{Vrfy}(\text{mVK}, pk')$ . If the output is zero then  $S$  aborts; otherwise,  $S$  sets  $\vec{x} = \{X_j^0, X_j^1\}_{j \in [m]}$ , computes  $c \leftarrow \text{Enc}(pk', \vec{x})$  and sends  $c$  and  $GC$  to  $P$ .
5.  $P$  computes  $\vec{m} \leftarrow \text{Dec}(sk', c)$ , where  $\vec{m} = \hat{X}_0^{\chi_0^0}, \dots, \hat{X}_m^{\chi_m^0}$  and computes  $Z \leftarrow \text{Ev}_{\text{ABKE}}(GC, \{\hat{X}_i^{\chi_i^j}\}_{i \in [m]})$ ;  $P$  sets  $Z := \perp$  if  $\text{Ev}_{\text{ABKE}}$  fails.
6.  $P$  sends  $(\text{commit}, \text{sid}, 1, Z)$  to  $\mathcal{F}_{\text{com}}$ , which sends  $(\text{committed}, \text{sid}, 1, |Z|)$  to  $S$ .
7.  $S$  sends the wire labels  $\{X_i^b\}$  and the randomness  $r$  used in the encryption to  $P$ , who verifies that the encryptions match the wire labels and then computes  $\text{Ve}_{\text{ABKE}}(C, GC, \{X_i^0, X_i^1\}_{i \in [m]})$ . If either the wire labels did not match the encrypted values or the output of  $\text{Ve}_{\text{ABKE}}$  is reject then  $P$  outputs  $\perp$ . Likewise, if  $C(\chi_j) = 0$  then  $P$  outputs  $\perp$ . Otherwise,  $P$  sends  $(\text{reveal}, \text{sid}, 1)$  to  $\mathcal{F}_{\text{com}}$ , which sends  $(\text{reveal}, \text{sid}, 1, Z)$  to  $S$ .
8.  $S$  checks that  $Z = Z^1$ ; if not, it sends  $\perp$  to  $P$  and halts. Otherwise, the parties both send  $(\text{toss}, \text{sid})$  to  $\mathcal{F}_{\text{cointoss}}$ , receive  $(\text{tossed}, \text{sid}, k)$ , and output  $k$ .

**Figure 6.1:** Protocol  $\Pi_{\text{abke}}$  realizing  $\mathcal{F}_{\text{abke}}$  in the  $(\mathcal{F}_{\text{com}}, \mathcal{F}_{\text{cointoss}}, \mathcal{F}_{\text{anon}})$ -hybrid model.

The  $\text{Ev}_{\text{ABKE}}$  and  $\text{Ve}_{\text{ABKE}}$  functions are defined naturally from  $\text{Ev}$ ,  $\text{Ve}$ , and  $\text{Gb}_{\text{ABKE}}$ .

Clearly, the scheme allows evaluation and verification in the same manner as the underlying garbling scheme once a label per each wire is obtained. At the same time, any party evaluating a garbled circuit must make a call to the random oracle per input-wire label in order to learn the “real” underlying label for the garbled circuit. Thus, the underlying garbling scheme  $\text{Gb}$  cannot be decrypted without a random oracle evaluation on an input-wire label of  $\text{Gb}_{\text{ABKE}}$ , which is exactly the property we need for the reduction to the collusion experiment.

We prove the following in the full version [35].

**THEOREM 6.1.** *Assume that the encryption scheme used in  $\Pi_{\text{abke}}$  is a secure attribute selective encryption scheme. Then  $\Pi_{\text{abke}}$  securely computes  $\mathcal{F}_{\text{abke}}$  in the  $(\mathcal{F}_{\text{com}}, \mathcal{F}_{\text{cointoss}}, \mathcal{F}_{\text{anon}})$ -hybrid model, in the random-oracle model.*

## 7. ELH SIGNATURES

We introduce the notion of *extractable linearly homomorphic (ELH) signatures* and show an implementation using the Boneh-Lynn-Shacham (BLS) [16] signature scheme. ELH signatures play a central role in our ASE constructions detailed in §8 and §9.

**DEFINITION 7.1.** (Linearly homomorphic signatures) *Let  $\text{Sig} = (\text{Sign}, \text{Vrfy})$  be a signature scheme over a space of messages consisting of elements of a group  $G$  of prime order  $q$ , with signatures also lying in this group. We say that  $\text{Sig}$  is linearly homomorphic over  $G$  if for any two elements  $g_1, g_2 \in G$ , it holds that  $\text{Sign}(g_1 g_2) = \text{Sign}(g_1) \text{Sign}(g_2)$ . The scheme is called *unforgeable* if no probabilistic polynomial-time algorithm given  $n$  pairs  $(g_i, \text{Sign}(g_i))$  for random elements  $g_1, \dots, g_n \in G$  and an additional random independent element  $g \in G$ , has non-negligible probability to output  $\text{Sign}(g)$ .*

Note that being linearly homomorphic implies that given  $n$  signed elements  $g_1, \dots, g_n \in G$ , one can compute (without the signing key) the signature of any linear combination (in the exponent) of  $g_1, \dots, g_n$ ; namely, for any  $x_1, \dots, x_n \in \mathbb{Z}_q$  we have that  $\text{Sign}(g_1^{x_1} \dots g_n^{x_n}) = \text{Sign}(g_1)^{x_1} \dots \text{Sign}(g_n)^{x_n}$ . We note that the requirement of the signatures lying in the same group as the message space is not essential but it simplifies notation by using the same group operation for group elements and signatures, and is a property of our implementation using BLS signatures. This notion can be seen as a one-dimensional case of homomorphic signatures for linear spaces [14, 27, 15]. Also note that the unforgeability property holds only with respect to random messages (i.e., random elements in the group).

We now define the property of extractability. It captures the intuition behind linearly homomorphic signatures as allowing limited malleability. That is, anyone can generate signatures on a value  $g$  without possessing the signing key as long as it *knows* a representation of  $g$  as a linear combination (in the exponent) of previously signed elements. Extractability formalizes this knowledge similarly to existing knowledge extractability notions. In spite of being intuitively appealing we are not aware of this form of homomorphic signatures being defined in prior work.

**DEFINITION 7.2.** (Extractable linearly homomorphic signatures) *Let  $G$  be a cyclic group of prime order  $q$  and  $\text{Sig} = (\text{Sign}, \text{Vrfy})$  a linearly homomorphic signature scheme over  $G$ . Consider algorithms that on input  $t$  random elements  $g_1, \dots, g_t$  in  $G$  and corresponding signatures  $\text{Sign}(g_1), \dots, \text{Sign}(g_t)$ , output a pair  $(f, \text{Sign}(f))$  for  $f \in G$  with non-negligible probability (over the choice of  $g_i$ s and the algorithm's random coins). We say that  $\text{Sig}$  is an *extractable linearly homomorphic (ELH) signature scheme* if for every polynomial-time algorithm  $F$  as above there exists another polynomial-time algorithm  $F'$  for which the following property holds, except for with negligible probability: Let  $\{g_i, \text{Sign}(g_i)\}_{i \in [t]}$  be an input to  $F$  on which  $F$  outputs  $(f, \text{Sign}(f))$ , then on the same inputs (and internal random coins)  $F'$  out-*



puts a vector  $(f, \text{Sign}(f), x_1, \dots, x_n)$  with  $x_i \in \mathbb{Z}_q$  such that  $f = g_1^{x_1} \cdots g_n^{x_n}$ .

Interestingly, extractability in linearly homomorphic signatures implies unforgeability as shown next (the proofs appear in the full version [35]).

**LEMMA 7.3.** *Let  $\text{Sig} = (\text{Sign}, \text{Vrfy})$  be an ELH signature scheme over a group  $G$  where the discrete logarithm problem is hard. Then  $\text{Sig}$  is unforgeable.*

**LEMMA 7.4.** *Under the hardness of the discrete log over group  $G$ , given  $g_1, \dots, g_n, g \in G \setminus \{1\}$ , it is infeasible to find  $x_1, \dots, x_n \in \mathbb{Z}_q$  such that  $g = g_1^{x_1} \cdots g_n^{x_n}$ . Similarly, finding two representations  $\prod_{i=1}^n g_i^{x_i} = \prod_{i=1}^n g_i^{y_i}$  such that there exists an  $i$  for which  $x_i \neq y_i$  is also infeasible.*

## 7.1 Implementation of ELH Signatures

We now demonstrate an implementation of an ELH signature scheme using the Boneh-Lynn-Shacham (BLS) [16] signature scheme, which we first recall.

**Boneh-Lynn-Shacham (BLS) signature scheme.** The scheme assumes groups  $(G_1, G_2, G_T)$  of prime order  $q$  with a bilinear pairing  $e : G_1 \times G_2 \rightarrow G_T$  where the co-CDH assumption holds (i.e., given  $g \in G_1, h, h^x \in G_2$ , finding  $g^x$  is infeasible). The public/private keypair is  $(h^x, x)$ , where  $x \in \mathbb{Z}_q$ , and  $h \in G_2$ . A signature on message  $m$  is computed as  $H(m)^x$  where  $H$  is a hash function mapping messages to random elements in  $G_1$ . Verification of a signature  $\sigma$  on message  $m$  under public key  $y = h^x$  is performed by checking the following equality:  $e(\sigma, h) = e(H(m), y)$ .

**BLS\* signature scheme.** We define the BLS\* scheme to be the same as BLS but the message space is the group  $G_1$  itself and *no hash function is applied to the messages* (this is sufficient for our application that only requires unforgeability on random group elements).

The following lemma shows that BLS\* leads to an implementation of ELH signatures under the  $t$ -KEA assumption [11, 10]. The proof appears in the full version [35].

**LEMMA 7.5.** *Under the  $t$ -KEA assumption over group  $G_1$ , BLS\* is an unforgeable extractable linearly homomorphic signature scheme.*

## 8. ASE USING IBE

We now construct an attribute selective encryption scheme from identity based encryption (IBE) and extractable linearly homomorphic (ELH) signatures. The security of the protocol is based on the security of the underlying IBE and ELH signature schemes. In addition, we require that the master public key of the IBE scheme be from a group so that it can be rerandomized and that the ELH signature scheme works over the same group.

We note that ASE can be constructed from IBE in a generic way if it satisfies these additional requirements. Both the Boneh-Franklin [13] and Boneh-Boyen [12] IBE schemes can be used and they yield different efficiency and computational requirements from the parties. However, this presentation will *not* be generic but rather at points we will be specific to the Boneh-Franklin IBE scheme.

We first give a high level overview of how we use the IBE scheme in our construction. Recall that in an IBE scheme a central authority chooses a master secret key and publishes

the correlated master public key. The master public key is used as part of the encryption key for all clients. Each client has an identity which is known to all and in addition each client receives a private secret key that is computed using the master secret key and its identity. A message is encrypted using the master public key and the identity of the client for whom the message is intended. The client uses its secret key to decrypt.

The first switch that we make in our scheme is that the “identities” are associated with the attributes. Thus, if a client has an attribute then it receives the secret key relating to that attribute. However, that clearly is not sufficient as collusions can take place. A client receiving the secret key for attribute  $j_1$  can collude with another client who has the secret key for  $j_2$ , enabling them to decrypt an unauthorized set of messages. Thus, we introduce our second switch which is that the center creates a “personalized” public master key for each client (by choosing a different master secret key) and modifies the secret keys of the client to relate to the personalized public key. Now this additional change prevents the clients from colluding as their secret keys relate to different master public keys. While this is the basic intuition, there are additional details that need to be added to satisfy all the requirements, which we describe below.

To present our scheme, we first recall the general construction of IBE and define its terminology. An IBE scheme is comprised of four parts: setup, key generation, encryption, and decryption.

- **Setup<sub>IBE</sub>( $1^n$ ):** Takes as input the security parameter. We split the setup into two parts: Part 1 outputs the public parameters  $\text{PP}_{\text{IBE}}$ , and Part 2 outputs the master public key  $\text{mpk}_{\text{IBE}}$  and the master secret key  $\text{msk}_{\text{IBE}}$ . The public parameters are common to all key pairs, and are implicitly given to all algorithms below.
- **KeyGen<sub>IBE</sub>( $\text{msk}_{\text{IBE}}, \text{ID}$ ):** Takes as input the master secret key and the identity  $\text{ID}$  of the client, and outputs the client’s secret key  $\text{sk}_{\text{IBE}}$ .
- **Enc<sub>IBE</sub>( $\text{mpk}_{\text{IBE}}, m, \text{ID}$ ):** Takes as input the master public key, the plaintext message  $m$ , and the identity of a client, and outputs the ciphertext  $c$ .
- **Dec<sub>IBE</sub>( $c, \text{sk}_{\text{IBE}}$ ):** Takes as input the ciphertext and the client’s secret key and outputs the message.

Due to lack of space, we assume familiarity with the definition of IBE. Here we informally state that encrypting with IBE is a secure encryption.

We are now ready to present our construction. Our IBE-ASE scheme is defined as follows.

- **Setup( $1^n, m$ ):**
  1. Run Part 1 of **Setup<sub>IBE</sub>**, receiving  $\text{PP}_{\text{IBE}}$ .
  2. Choose  $2m$  random strings,  $g_{j,0}, g_{j,1}$ , for  $j \in [m]$ , where  $g_{j,0}$  corresponds to not having attribute  $j$ , and  $g_{j,1}$  corresponds to having the attribute. These will be the “identities” of the system.
  3. Run a key generation protocol for an ELH signature scheme, receiving and setting  $\text{mVK}$  to the public verification key and  $\text{mSK}$  to the secret signing key.
  4. Output  $\text{PP} := (\text{PP}_{\text{IBE}}, \{g_{j,0}, g_{j,1}\}_{j \in [m]})$ ,  $\text{mVK}$ , and  $\text{mSK}$ .
- **GenCert( $\text{mSK}, \chi$ ):**

1. Run Part 2 of  $\text{Setup}_{\text{IBE}}$ , creating  $\text{msk}_{\text{IBE}}$  and  $\text{mpk}_{\text{IBE}}$ . Set  $pk_\chi := (\text{mpk}_{\text{IBE}}, \sigma = \text{Sign}(\text{mpk}_{\text{IBE}}))$ . In what follows we sometimes abuse notation and refer to  $pk_\chi$  only as the public key and sometimes as both the public key and its signature.
2. For each attribute  $j \in [m]$ , call  $\text{KeyGen}_{\text{IBE}}$  on input  $\text{msk}_{\text{IBE}}$  and identity  $g_{j,\chi[j]}$ . This returns a secret key  $sk_j$ .
3. Set the client's secret key to  $sk_\chi := (sk_1, \dots, sk_m)$ .

Note that now we can discard the master secret key  $\text{msk}_{\text{IBE}}$ . We use the signature to compensate for the fact that  $\text{mpk}_{\text{IBE}}$  is not one of the public parameters of the system.

- $\text{Vrfy}(\text{mVK}, pk_\chi)$ : Output a bit attesting to the validity of the public key  $pk_\chi$  by checking the signature  $\sigma$ .
- $\text{Enc}(pk_\chi, \{x_{j,0}, x_{j,1}\}_{j \in [m]})$ :
  1. Verify that  $\text{Vrfy}(\text{mVK}, pk_\chi) = 1$ . If not then abort.
  2. For  $j \in [m]$ , compute  $c_{j,0} \leftarrow \text{Enc}_{\text{IBE}}(pk_\chi, x_{j,0}, g_{j0})$  and  $c_{j,1} \leftarrow \text{Enc}_{\text{IBE}}(pk_\chi, x_{j,1}, g_{j1})$ .
  3. Output  $c := c_{1,0}, c_{1,1}, \dots, c_{m,0}, c_{m,1}$ .
- $\text{Dec}(sk_\chi, c)$ : Output  $x_j := \text{Dec}_{\text{IBE}}(c_{j,\chi[j]}, sk_{\chi[j]})$  for all  $j \in [m]$ .
- $\text{Unlink}(pk_\chi, sk_\chi)$ : In our implementation using the Boneh-Franklin IBE scheme, the public key  $pk_\chi$  has the form  $g^z$  for some value  $z$ . We implement the unlink operation by raising  $pk_\chi$  to a random exponent  $s$ . All other values are also raised to  $s$ , including the ELH signature and every component of the secret key.
- $\text{Project}(pk_\chi, sk_\chi, S)$ : The project function in the IBE case is trivial; all that needs to be done is to remove from the secret key the elements whose index is not in the set  $S$ . The public key remains the same.
- $\text{GenCert}^*(\text{mSK})$ : Exactly as in  $\text{GenCert}$ , except now the secret key  $sk$  contains secret keys  $sk_{j,b}$  for all  $j \in [m]$  and  $b \in \{0, 1\}$ .  
This can be trivially achieved in two manners: first, by giving additional secret keys exactly in the format as given in  $\text{GenCert}$  or even simpler by just giving  $x$  to the client.
- $\text{Dec}^*(sk, c)$ : Same as  $\text{Dec}$ , except here the full set of  $2m$  plaintexts are returned. This is completely straightforward, as the secret key now includes decryption keys for all messages.

The proof of the following theorem is presented in the full version [35].

**THEOREM 8.1.** *The above scheme is a secure ASE scheme when instantiated with the Boneh-Franklin IBE and an ELH signature scheme. The ELH signature scheme needs to work over the same group  $G$ , as defined in Boneh-Franklin.*

## 9. ASE USING ELH SIGNATURES

We present an instantiation of an attribute-selective encryption (ASE) scheme (cf. §5) based on extractable linearly homomorphic (ELH) signatures (cf. §7). The key generation and certification mechanism uses the homomorphic property of the signatures while their extractability properties are used to prove security.

The ELH-ASE scheme is defined as follows.

- $\text{Setup}(1^n, m)$ : Let  $\text{Sig} = (\text{Sign}, \text{Vrfy})$  be an ELH signature scheme over a group  $G$  of order  $q$ .  $\text{Setup}$  uses

$\text{Sig}$  to sign several elements in each client's public key (defined below). Each element type  $(g, h, u)$  has a dedicated signature key and there is also a per-attribute signature key to sign elements of the form  $ue_j$ . For readability, we denote the above four types of signatures by  $\text{Sign}_g, \text{Sign}_h, \text{Sign}_u$  and  $\text{Sign}_j$  for  $j \in [m]$ . The set of public verification keys for the above signatures form the master verification key  $\text{mVK}$  and the corresponding secret signing keys form the master secret key  $\text{mSK}$ .

We refer to the party running the  $\text{Setup}$  function as the CA.

- $\text{GenCert}(\text{mSK}, \chi)$ : A public key  $pk_\chi$  associated with an attribute vector  $\chi = (\chi_1, \dots, \chi_m)$  is generated as follows:

1. Choose random elements  $g, h, u \in G \setminus \{1\}$  and compute signatures  $\text{Sign}_g(g)$ ,  $\text{Sign}_h(h)$ , and  $\text{Sign}_u(u)$ .
2. Choose  $r_1, \dots, r_m \in_R \mathbb{Z}_q^*$  and set  $e_j = g^{r_j}$  if  $\chi[j] = 0$  and  $e_j = h^{r_j}$  if  $\chi[j] = 1$ ; compute signatures  $\text{Sign}_j(ue_j)$  for  $j \in [m]$ .

Set the public key  $pk_\chi$  to  $(g, h, u, \{e_j\}_{j \in [m]}, \text{Sign}_g(g), \text{Sign}_h(h), \text{Sign}_u(u), \{\text{Sign}_j(ue_j)\}_{j \in [m]})$  and the secret key  $sk_\chi$  to  $\{r_j\}_{j \in [m]}$ .

- $\text{Vrfy}(\text{mVK}, pk_\chi)$ : Check that  $g, h, u \in G \setminus \{1\}$  and use  $\text{mVK}$  to check all signatures.
- $\text{Enc}(pk_\chi, \{x_{j,0}, x_{j,1}\}_{j \in [m]})$ : For  $j \in [m]$ , choose  $s_j, t_j \in_R \mathbb{Z}_q$  and set  $c_{j,0} := (g^{s_j}, e_j^{s_j} \cdot x_{j,0})$  and  $c_{j,1} := (h^{t_j}, e_j^{t_j} \cdot x_{j,1})$ . The ciphertext is the sequence  $\{(c_{j,0}, c_{j,1})\}_{j \in [m]}$ . (Note: We assume for simplicity that the random values  $x_{j,b}$ , that correspond to input wire labels  $X_j^b$  in protocol  $\Pi_{\text{abke}}$  of Fig. 6.1 are random elements in  $G$  (these values are later hashed into strings  $h_j^b$  as part of the garbling scheme  $\mathcal{G}_{\text{ABKE}}$ .)
- $\text{Dec}(sk_\chi, \{(c_{j,0}, c_{j,1})\}_{j \in [m]})$ : For  $j \in [m]$ , set  $(C_1, C_2)$  to the pair  $c_{j,\chi_j}$  and compute  $x_{j,\chi_j} := C_2/C_1^{r_j}$ .
- $\text{Unlink}(pk_\chi, sk_\chi)$ : Choose a random value  $r \in \mathbb{Z}_q$  and raise every element of  $pk$  to the power of  $r$ ; output:

$$pk'_\chi := (g^r, h^r, u^r, \{e_j^r\}_{j \in [m]}, (\text{Sign}_g(g))^r, (\text{Sign}_h(h))^r, (\text{Sign}_u(u))^r, \{(\text{Sign}_j(ue_j))^r\}_{j \in [m]})$$

and  $sk'_\chi := \{r \cdot r_j\}_{j \in [m]}$ .

- $\text{Project}(pk_\chi, sk_\chi, S)$ : Output a new public key by omitting any component  $e_j$  for  $j \notin S$ , and set the corresponding secret key to  $\{r_j\}_{j \in S}$ .
- $\text{GenCert}^*(\text{mSK})$ : Generation of the pair  $(pk_\chi, sk_\chi)$  is the same as for  $\text{GenCert}$  except that  $h$  is set to  $g^\tau$  for known  $\tau \in_R \mathbb{Z}_q$  and  $e_j$  is set to  $g^{r_j}$  for all  $j \in [m]$  (as in the case  $\chi = 0^m$ ). The secret key  $sk$  is comprised of the set  $\{r_j\}_{j \in [m]}$  and the value  $\tau$ . This enables  $\text{Dec}^*$  as follows.
- $\text{Dec}^*(sk, \{(c_{j,0}, c_{j,1})\}_{j \in [m]})$ : For  $j \in [m]$ , set  $(C_1, C_2)$  to the pair  $c_{j,0}$  and compute  $x_{j,0} := C_2/C_1^{r_j}$ . Then, for  $j \in [m]$ , set  $(C_1, C_2)$  to the pair  $c_{j,1}$  and compute  $x_{j,1} := C_2/C_1^{\tau \cdot r_j}$ .

The proof of the following theorem is presented in the full version [35].

**THEOREM 9.1.** *If  $\text{Sig}$  is an extractable linearly homomorphic signature scheme over a DDH group  $G$  then under the DDH and KEA assumptions over  $G$ , the ELH-ASE scheme*

has the properties of collusion resistance, attribute privacy, unlinkability, projectability and committing encryption.

## 10. PERFORMANCE


As the bottleneck in terms of computation is the pairings and exponentiations (versus garbling and evaluating the policy), we now give a concrete count of the number of pairings and exponentiations required for each of the two ASE schemes presented in §8 and §9.

- *IBE-based scheme (§8).* For concreteness we calculate the cost using the Boneh-Franklin IBE scheme [13].
  - The client computes two exponentiations to randomize both its “master public key” and its associated signature.
  - The server computes two pairings to verify the signature of the client’s “master public key”. To encrypt  $2m$  messages, the server computes  $2m$  pairings and  $2m$  exponentiations.
  - The client computes  $m$  pairings and exponentiations to decrypt  $m$  messages.
- *ELH signature-based scheme (§9).*
  - The client computes a total of  $6 + 2m$  exponentiations to randomize both the basis and its associated signature (6) and the public keys and their associated signatures ( $2m$ ). These operations can all be done offline.
  - The server computes  $3 + m$  signature verifications, which requires  $6 + 2m$  pairings (this can be sped up by batch verification of signatures; see below). To encrypt the  $2m$  messages it computes  $4m$  exponentiations.
  - The client computes  $m$  exponentiations to decrypt  $m$  messages.

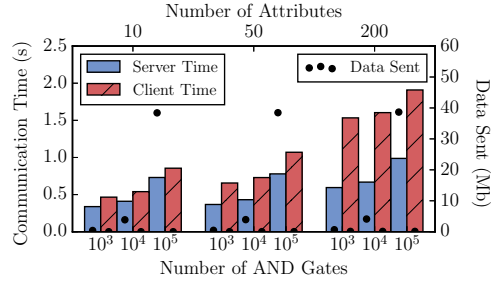
We note several important points regarding the performance of the ELH signature-based scheme. First, the scheme requires the client to only compute exponentiations as opposed to pairings. This could be meaningful in a setting where the client is a small computing device. Second, the server can batch multiple signature verifications from different clients. The CA’s signature keys for  $g, h, u$  and the attributes are the same for all clients. Using techniques of Ferrara et al. [24] for batching pairing-based signatures can help us achieve better amortized run-times.

**Implementation and results.** We implemented the scheme described in Figure 6.1 using the ELH signature-based ASE scheme (cf. §9) utilizing all the optimizations mentioned above. We instantiate the coin-tossing and commitment functionalities using SHA-1, and use the privacy-free garbling technique of Zahur et al. [39]. The code as well as all the scripts for generating our experimental results are available at <https://github.com/amaloz/abke>.


All experiments were conducted on an Intel Core i5-4210H CPU running at 2.90 GHz, and we use the RELIC library [6] for pairings, using the optimal ate pairing on the BN-P256 curve. On our benchmark machine, pairings take roughly 1.88 ms and exponentiations in  $G_1$  take roughly 160  $\mu$ s (76  $\mu$ s when using preprocessing). All experiments were run over localhost; however, to emulate a WAN environment we used the `tc` command in Linux to set the latency



**Figure 10.1:** Computation time of the server and client for various numbers of attributes and sizes of the policy. The bottom x-axis gives the number of gates in the policy, the top x-axis gives the number of attributes, and the y-axis gives the computation time in seconds.



**Figure 10.2:** Communication time of the server and client for various numbers of attributes and sizes of the policy circuit. The bottom x-axis gives the number of gates in the policy, the top x-axis gives the number of attributes, the left y-axis gives the computation time (in seconds), and the right y-axis gives the number of bits sent (in Mb).



**Figure 10.3:** Optimized computation time (i.e., pushing the cost of randomizing the public key and garbling the policy to an offline stage, along with batching of the key verification) of the server and client for various numbers of attributes and sizes of the policy. The bottom x-axis gives the number of gates in the policy, the top x-axis gives the number of attributes, and the y-axis gives the computation time in seconds.

to 33 ms (the average latency in the United States [2]) and the bandwidth to 200 Mbps. For each measurement, we ran 10 iterations of 10 runs, taking the mean of the medians from each run.

Figures 10.1 and 10.2 show the results of our experiments. We varied the number of attributes  $m$  between 10, 50, and 200, and varied the size of the policy (comprised of only AND gates) between 1,000, 10,000, and 100,000. Figure 10.1 depicts the *computation* time of the server and client, whereas Figure 10.2 depicts the *communication* time. We also list the number of bits sent by the server and client in Figure 10.2.

Step	S	S [opt]	P	P [opt]
2 (Gb)	5 ms	—	—	—
3 (Unlink)	—	—	78 ms	—
4 (Vrfy)	857 ms	159 ms	—	—
4 (Enc)	82 ms	82 ms	—	—
5 (Dec)	—	—	28 ms	28 ms
5 (Ev)	—	—	3 ms	3 ms
6 (commit)	—	—	<1 ms	<1 ms
7 (Enc)	—	—	42 ms	42 ms
7 (Ve)	—	—	8 ms	8 ms
8 (cointoss)	<1 ms	<1 ms	<1 ms	<1 ms
Total	944 ms	241 ms	159 ms	81 ms

**Table 10.1:** Breakdown of server (S) and client (P) computation times for the various steps of  $\Pi_{\text{abke}}$  for a 100,000 gate policy with 200 attributes. [opt] denotes the optimized computation time (i.e., pushing computation to an offline stage and batching verification). See Figure 6.1 for a description of each step. The total cost is slightly less than that reported elsewhere due to rounding errors and not accounting for initialize/cleanup steps.

As we can see, the computation time is fairly consistent for a fixed  $m$ , but grows as  $m$  increases. This validates our claim that the pairings and exponentiations account for most of the overhead as opposed to the garbling and evaluating of the policy. The computation time varies from 67 ms for the server and 11 ms for the client for a 1,000 gate policy with 10 attributes, to 957 ms for the server and 176 ms for the client for a 100,000 gate policy with 200 attributes.

Looking at the case of a 100,000 gate policy with 200 attributes (see also Table 10.1), we see that most of the overhead on the server side comes from verifying the public key sent by the client (857 ms), due to the  $2m$  pairings needed. The next largest operation is encryption, which accounts for 82 ms. Meanwhile, garbling the policy takes only 5 ms. Regarding the client, the costliest operation is checking that the encryption sent by the server is correct, which requires re-encrypting the  $m$  unopened wire-labels (42 ms), followed by randomizing its public key (78 ms). Decryption is relatively cheap, requiring 28 ms. Meanwhile, evaluating the garbled circuit takes 3 ms, again demonstrating that the garbled circuit is not the bottleneck (at least with regards to computation).

Looking at the communication time (cf. Figure 10.2), we see that as both the number of attributes and number of gates grows so does the running time. We stress that this growth impacts previous ABE formula-based solutions to a much greater degree. Importantly, for all but extremely large policies of millions of gates, communication will typically not be a bottleneck of our system. Most of the server’s communication time is spent sending the garbled circuit, whereas most of the client’s time is spent receiving the garbled circuit and the ciphertext, this latter case due to the client blocking while the server verifies the (randomized) public key. We note that our network bandwidth of 200 Mbps is pessimistic, and running our protocol on Amazon EC2 or other networks with 1 Gbps bandwidth will all but eliminate the communication overhead of sending/receiving the garbled circuit (e.g., when running over localhost, the communication time is essentially the time spent blocking waiting for the other party to complete some computation).

Note that with regards to computation, most of the ex-

Operation	Cycles
BLS* sign	522,767
BLS* verify	12,316,919
BLS* batch verify	22,635,625

**Table 10.2:** Benchmarking BLS\* signing and verification, along with the batch verification approach of Ferrara et al. [24] for  $ten$  messages.

pensive operations (such as randomizing and verifying the public key) can be either done offline or batched. Thus, we also calculated an *optimized* computation time; see Figure 10.3. In these experiments, we ignore the cost of the client randomizing its public key and the server garbling its policy, as both of these can be done in an offline stage. To account for the batching optimization, we implemented and benchmarked the batching techniques of Ferrara et al. [24], see Table 10.2. We see a roughly  $5.4\times$  improvement when batch verifying ten messages. Thus, in our experiments we model a server operating over ten clients at a time by dividing the public key verification time by 5.4. We see upwards of a  $4.4\times$  and  $2\times$  improvement in running time for the server and client, respectively. This makes sense in light of the fact that randomizing and verifying the public key are the two most expensive operations.

**A policy cost example.** In light of the above results, we provide a rough calculation of the cost of a realistic policy, where the client succeeds if its geolocation  $(x_U, y_U)$  is within distance  $d$  of the server’s location  $(x_S, y_S)$ . The client’s geolocation credential may be issued with a certified timestamp, which may simultaneously be checked by a policy. Such a policy would require a circuit computing  $(x_U - x_S)^2 + (y_U - y_S)^2 <_? d^2$  and a (much smaller) circuit verifying that the timestamp is in an acceptable time interval. Using a 64-bit input and the CBMC circuit compiler [33], we can compile this function as a circuit containing approximately 20,000 gates. Thus, as demonstrated by our performance results, the cost of the corresponding garbled circuit would be unnoticeable relative to the public key operations required by the server and client. In contrast, an ABE-based solution would require converting the policy circuit into a (very large) formula, and performing pairings proportional to its size, which is not practical in most settings.

## Acknowledgments

This work was supported by the Office of Naval Research (ONR) contract number N00014-14-C-0113. This work was done in part while the authors were visiting the Simons Institute for the Theory of Computing, supported by the Simons Foundation and by the DIMACS/Simons Collaboration in Cryptography through NSF grant #CNS-1523467. Work of Alex J. Malozemoff conducted in part with Government support through the National Defense Science and Engineering Graduate (NDSEG) Fellowship, 32 CFG 168a, awarded by DoD, Air Force Office of Scientific Research.

## 11. REFERENCES

- [1] ABC4Trust EU project. <https://www.abc4trust.eu>.
- [2] Global IP network latency. [http://ipnetwork.bgtmo.ip.att.net/pws/network\\_delay.html](http://ipnetwork.bgtmo.ip.att.net/pws/network_delay.html). Retrieved February 8, 2016.

- [3] Identity mixer. <http://idemix.wordpress.com>.
- [4] Microsoft U-Prove. <http://www.microsoft.com/uprove>.
- [5] M. Abe and S. Fehr. Perfect NIZK with adaptive soundness. In *TCC*, 2007.
- [6] D. F. Aranha and C. P. L. Gouvêa. RELIC is an Efficient LIBrary for Cryptography. <https://github.com/relic-toolkit/relic>.
- [7] M. Belenkiy, J. Camenisch, M. Chase, M. Kohlweiss, A. Lysyanskaya, and H. Shacham. Randomizable proofs and delegatable anonymous credentials. In *Crypto*, 2009.
- [8] M. Bellare, V. T. Hoang, and P. Rogaway. Foundations of garbled circuits. In *CCS*, 2012.
- [9] J. Bethencourt, A. Sahai, and B. Waters. Ciphertext-policy attribute-based encryption. In *S&P*, 2007.
- [10] N. Bitansky, R. Canetti, A. Chiesa, S. Goldwasser, H. Lin, A. Rubinfeld, and E. Tromer. The hunting of the SNARK. Cryptology ePrint Archive, Report 2014/580, 2014.
- [11] N. Bitansky, R. Canetti, A. Chiesa, and E. Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In *ITCS*, 2012.
- [12] D. Boneh and X. Boyen. Efficient selective-ID secure identity based encryption without random oracles. In *Eurocrypt*, 2004.
- [13] D. Boneh and M. K. Franklin. Identity-based encryption from the Weil pairing. In *Crypto*, 2001.
- [14] D. Boneh, D. Freeman, J. Katz, and B. Waters. Signing a linear subspace: Signature schemes for network coding. In *PKC*, 2009.
- [15] D. Boneh and D. M. Freeman. Homomorphic signatures for polynomial functions. In *Eurocrypt*, 2011.
- [16] D. Boneh, B. Lynn, and H. Shacham. Short signatures from the Weil pairing. *Journal of Cryptology*, 17(4):297–319, Sept. 2004.
- [17] S. Brands. *Rethinking Public Key Infrastructures and Digital Certificates: Building in Privacy*. MIT Press, 2000.
- [18] J. Camenisch, N. Casati, T. Groß, and V. Shoup. Credential authenticated identification and key exchange. In *Crypto*, 2010.
- [19] J. Camenisch and A. Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In *Eurocrypt*, 2001.
- [20] J. Camenisch and A. Lysyanskaya. Signature schemes and anonymous credentials from bilinear maps. In *Crypto*, 2004.
- [21] R. Canetti, A. Cohen, and Y. Lindell. A simpler variant of universally composable security for standard multiparty computation. In *Crypto*, 2015.
- [22] M. Chase, C. Ganes, and P. Mohassel. Efficient zero-knowledge proof of algebraic and non-algebraic statements with applications to privacy preserving credentials. In *Crypto*, 2016.
- [23] D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, 1981.
- [24] A. L. Ferrara, M. Green, S. Hohenberger, and M. Ø. Pedersen. Practical short signature batch verification. In *CT-RSA*, 2009.
- [25] T. K. Frederiksen, J. B. Nielsen, and C. Orlandi. Privacy-free garbled circuits with applications to efficient zero-knowledge. In *Eurocrypt*, 2015.
- [26] S. Garg, C. Gentry, S. Halevi, A. Sahai, and B. Waters. Attribute-based encryption for circuits from multilinear maps. In *Crypto*, 2013.
- [27] R. Gennaro, J. Katz, H. Krawczyk, and T. Rabin. Secure network coding over the integers. In *PKC*, 2010.
- [28] M. C. Gorantla, C. Boyd, and J. M. González Nieto. Attribute-based authenticated key exchange. In *ACISP*, 2010.
- [29] S. Gorbunov, V. Vaikuntanathan, and H. Wee. Predicate encryption for circuits from LWE. In *Crypto*, 2015.
- [30] V. Goyal, O. Pandey, A. Sahai, and B. Waters. Attribute-based encryption for fine-grained access control of encrypted data. In *CCS*, 2006.
- [31] J. Groth. Short pairing-based non-interactive zero-knowledge arguments. In *Asiacrypt*, 2010.
- [32] D. Gupta and A. Sahai. On constant-round concurrent zero-knowledge from a knowledge assumption. In *Indocrypt*, 2014.
- [33] A. Holzer, M. Franz, S. Katzenbeisser, and H. Veith. Secure two-party computations in ANSI C. In *CCS*, 2012.
- [34] M. Jawurek, F. Kerschbaum, and C. Orlandi. Zero-knowledge using garbled circuits: how to prove non-algebraic statements efficiently. In *CCS*, 2013.
- [35] V. Kolesnikov, H. Krawczyk, Y. Lindell, A. Malozemoff, and T. Rabin. Attribute-based key exchange with general policies. <https://eprint.iacr.org/2016/518>.
- [36] A. Sahai and B. R. Waters. Fuzzy identity-based encryption. In *Eurocrypt*, 2005.
- [37] Y. Sakai, N. Attrapadung, and G. Hanaoka. Attribute-based signatures for circuits from bilinear map. In *PKC*, 2016.
- [38] B. Waters. Ciphertext-policy attribute-based encryption: An expressive, efficient, and provably secure realization. In *PKC*, 2011.
- [39] S. Zahur, M. Rosulek, and D. Evans. Two halves make a whole - reducing data transfer in garbled circuits using half gates. In *Eurocrypt*, 2015.