

PRShare: A Framework for Privacy-Preserving, Interorganizational Data Sharing

Lihl Idan
Yale University
lihi.idan@yale.edu

Joan Feigenbaum
Yale University
joan.feigenbaum@yale.edu

ABSTRACT

We consider the task of interorganizational data sharing, in which data owners, data clients, and data subjects have different and sometimes competing privacy concerns. One real-world scenario in which this problem arises is law-enforcement use of phone-call metadata: The data owner is a phone company, the data clients are law-enforcement agencies, and the data subjects are individuals who make phone calls. A key challenge in this type of scenario is that each organization uses its own set of proprietary intraorganizational attributes to describe the shared data; such attributes cannot be shared with other organizations. Moreover, data-access policies are determined by multiple parties and may be specified using attributes that are not directly comparable with the ones used by the owner to specify the data. We propose a system architecture and a suite of protocols that facilitate dynamic, efficient, and privacy-preserving interorganizational data sharing, while allowing each party to use its own set of proprietary attributes. We introduce the novel technique of *Attribute-Based Encryption With Oblivious Attribute Translation (OTABE)*, which plays a crucial role in our solution and may be of independent interest.

CCS CONCEPTS

• Security and privacy → Access control; Privacy-preserving protocols; • Theory of computation → Cryptographic protocols.

KEYWORDS

Attribute-Based Encryption; Privacy-preserving data sharing

ACM Reference Format:

Lihl Idan and Joan Feigenbaum. 2020. PRShare: A Framework for Privacy-Preserving, Interorganizational Data Sharing. In *19th Workshop on Privacy in the Electronic Society (WPES'20), November 9, 2020, Virtual Event, USA*. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3411497.3420226>

1 INTRODUCTION

As the amount, complexity, and value of data available in both private and public sectors has risen sharply, data management and access control have challenged many organizations. Even more

challenging are management and access control in *interorganizational* data sharing. Each organization would like to minimize the amount of sensitive information disclosed to other organizations, including both information about the data and information about the organization's work methodologies and role structure.

1.1 Problem description

We consider scenarios in which multiple organizations need to share data while each organization uses its own set of *proprietary metadata* to describe the shared data. In these scenarios, data records contain a *payload*, which is the actual data, and a set of *metadata attributes* that describe the payload. Although organizations may agree to share the payload, each uses a different set of metadata attributes, taken from its own professional domain, to describe this payload. Data must be shared in a controlled manner that protects the confidentiality of each organization's proprietary attributes and prevents unauthorized users from accessing the payload.

Typically, one organization, the *data owner*, maintains a set of data records that are potentially useful to other organizations, called the *data clients*. Each data record contains sensitive information about an individual, the *data subject*. *Data users*, who are employees of a data client, may need access to data records stored by the data owner to perform their assigned tasks; each user must have the proper authorization to access the payloads of the specific set of records needed for a given task. Our framework also features a third type of organization, *data intermediaries*, that enrich data with additional information that is needed for the client's tasks but is available only to the intermediary. Each organization ORG_i maintains its own vocabulary VOC_i that contains the overall set of domain-specific, intraorganizational attributes used in its operations. VOC_i includes both proprietary, sensitive attributes and attributes that can be shared with other organizations. ORG_i uses a different set of attributes, $ATT_{i,j} \subseteq VOC_i$, to describe each shared payload p_j .

For example, the data owner may be an email service provider (ESP). The data records represent email messages. Each email record is composed of a payload, which is the content of the email message, and metadata attributes about the payload such as sender, receiver, and date. Some attributes, e.g., the email message's receiver, are sensitive; therefore, the ESP will share them with other organizations only when required to do so and only in a controlled manner. Each email message is created by one of the ESP's customers, who are the data subjects; it is then stored and cataloged using attributes that represent the message's metadata as collected by the ESP. Clients may be law-enforcement (LE) agencies, in which agents (data users) need access to email records in order to perform investigations. Intermediaries may include government agencies such as the IRS,

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.

WPES'20, November 9, 2020, Virtual Event, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8086-7/20/11...\$15.00

<https://doi.org/10.1145/3411497.3420226>

which could provide tax records associated with the email addresses that appear in the messages' metadata attributes.

Design goals: Each organization wishes to maintain its proprietary view of the shared data and to keep that view confidential. This means that the set $ATT_{i,j}$ of attributes that ORG_i maintains on each shared payload must be hidden from the other organizations.

Another requirement that must be accommodated is the use of multiple vocabularies. The owner uses vocabulary VOC_1 to store and query the shared data, an intermediary uses a different vocabulary VOC_2 to enrich the shared data, and the client uses a third vocabulary VOC_3 to query and process the data, to manage access control, and to issue data-access authorizations to its employees. Therefore, our framework must provide a mechanism that *dynamically and obliviously* transforms attributes of shared data from one vocabulary to another. Note that that this problem cannot be solved simply by requiring any set of organizations that may need to share data to agree on a shared, standard vocabulary. Such a standardization effort would require the organizations to know both the names and values of attributes used by other organizations; however, our premise is that the values of many attributes used internally by organizations are sensitive and cannot be exposed to other organizations. Furthermore, in many natural use cases (see Sec. 2.2), transformations require auxiliary information, such as up to date statistics or lists. Such information is known only at the point at which a user requests a specific data record and may need to be supplied by an intermediary that is not known by the data owner at the time that the owner encrypts the data.

Finally, because attributes could reveal sensitive aspects of organizations' activities, regulators and data subjects should expect sharing of both payloads and attributes to be kept to a minimum. To facilitate minimal exposure of sensitive information, an interorganizational data-sharing framework should offer a data-centric access-control mechanism. Such a mechanism will allow a user to access a payload only if it is essential for the completion of one of her tasks; in addition, it will allow the user to learn only the subset of that payload's attributes that are needed for the task.

1.2 Starting point: attribute-based encryption

Attribute-based encryption (ABE) is a natural starting point in the design of our framework. In our terminology, the encryptor is the data owner, users are data clients' employees (data users), and trusted authorities (TAs) both inside and outside the data client determine users' access policies. An ABE scheme grants an individual user a key that permits him to decrypt a ciphertext if and only if the key matches certain attributes specified during the ciphertext's creation. ABE enables fine-grained access control, which is essential in a privacy-preserving data-sharing framework. It provides one-to-many encryption, which can significantly increase the scalability of encryption and key management – properties that are necessary for interorganizational data sharing. ABE policy-access formulae are highly expressive, because they can be specified with binary or multivalued attributes, using AND, OR, and threshold gates.

Existing ABE schemes, however, have several properties that make them unsuitable for our framework.

In existing ABE schemes, encryptors, users, and TAs all use the same vocabulary. This means that these schemes cannot be

used off-the-shelf in our framework, where a crucial element of the problem description is that participating organizations may belong to different business sectors or professional domains and thus use different vocabularies. In particular, a data client's TAs and employees use a different vocabulary from that of the data owner. In ABE terms, this implies that attributes used in access policies (and keys) issued by the TAs to data users might belong to a different vocabulary from the one used by the owner to encrypt and store ciphertexts. Unless a suitable transformation is made between the keys and the ciphertexts, decryption will fail even if the ciphertext satisfies the user's access policy. Such a transformation must separately consider each attribute in the ciphertext and change it into a valid attribute from the users' keys' vocabulary. To protect both data subjects' privacy and organizations' proprietary views, the original attribute must remain hidden from the user and the new attribute must remain hidden from the encryptor. Existing ABE schemes cannot support this requirement.

Moreover, existing ABE schemes are generally used for role-based access control and thus have user-centric vocabularies (attributes that describe decryptors' traits) that reflect organizational structure and roles. The use of user-centric attributes, coupled with the single-vocabulary assumption, implies that the encryptor (data owner) must be exposed to the roles of potential decryptors (clients' data users) and the organizational structure that they fit into. Many organizations are reluctant to share such sensitive information.

1.3 Main contributions

We present a new system architecture and a suite of protocols for interorganizational data sharing that support privacy of both data (*payload hiding*) and organizational vocabularies (*attribute hiding*). We introduce *Attribute-Based Encryption With Oblivious Attribute Translation (OTABE)*, in which a semi-trusted proxy *translates* the attributes under which a data record's payload was encrypted into the attributes under which it can be decrypted by authorized users. The proxy performs the translation without learning the underlying plaintext data. Moreover, translation is *oblivious* in the sense that the attributes under which the record is encrypted remain hidden from the proxy. We provide a concrete OTABE scheme and prove it selectively secure in the standard model, resulting in an efficient, expressive, and flexible interorganizational data-sharing framework that we call **PRShare**. Its desirable features include:

Data centrality: It uses data-centric vocabularies. This means that decryption policies are determined according to traits of the *data requested by a query* rather than the *queriers' roles*. Data centrality protects, to the extent possible, the data subjects' privacy and the client's internal organizational structure.

Division of trust: We design and implement a *multi-authority* OTABE scheme (MA-OTABE). Each authority determines its own decryption policy. A data user is able to decrypt a record only if its attribute set satisfies all of the decryption policies.

Hidden access policy: The set of attributes associated with each ciphertext is hidden from the proxies, the users, and the cloud servers on which the owner stores the encrypted data.

Multi-vocabulary: Each organization can use its own set of attributes to describe the shared data.

Attribute privacy: The metadata information that each organization maintains on the shared data remains confidential to the extent that the organization requires.

Dynamically reconfigurable attributes: Translation is performed using dynamically updated information, and authorized users obtain up-to-date attributes. This feature does not require re-encryption of the records.

Offline delegation: The owner does not need to authorize or serve clients' data-access queries.

Direct revocation: We leverage our attribute-translation technique to create a modular and efficient revocation mechanism that enables *direct revocation*; this means that revoking the keys of a set U of users does not affect the keys of users not in U .

Key-abuse prevention and collusion resistance: Although decryption policies in PRShare are based on data-centric attributes, the decryption keys do include a user-specific component. This personalization of both keys and ciphertexts yields a collusion-resistant scheme and prevents key-abuse attacks.

Before proceeding to our technical results, we note that our approach is not suitable for *all* data-sharing applications. For example, it is not intended for scenarios in which the data subject participates directly in the user's request for data about her and could be asked to grant explicit consent. In general, data subjects in the scenarios we consider will not even be aware of the specific uses that are made of data about them. Similarly, our approach is not intended for scenarios in which there are clear, efficiently decidable, and universal rules that govern which users can access which portions of the data; existing access-control mechanisms suffice in such scenarios. Our techniques are useful in scenarios in which there are legally mandated, general principles that govern access to sensitive data, but instantiating those principles in the form of efficiently decidable rules requires data-specific and dynamically changing knowledge. We give two examples in Sec. 2.2.

Preliminary version: In this preliminary version of the paper, we provide simplified versions of our constructions and omit many technical details; full, detailed constructions and proofs will be given in our forthcoming journal paper and can be found in our technical report [24].

2 BACKGROUND AND MOTIVATION

2.1 Related work

Existing privacy-preserving data-sharing schemes fall into two general categories: centralized and decentralized. The former category includes the works of Jajodia *et al.* [25], De Capitani di Vermicati *et al.* [12], Dong *et al.* [13], X. Liu *et al.* [32], Popa *et al.* [36] and Vinayagamurthy *et al.* [43]. The major advantage of the centralized approach is efficiency; disadvantages include single points of failure and the lack of division of trust. Decentralized solutions can be found in the work of Fabian *et al.* [15], Froelicher *et al.* [19], C. Liu *et al.* [31], and Nayak *et al.* [33]. Decentralized solutions avoid single points of failure, but they often have limited efficiency or scalability.

The original motivation for PRShare was enhancement of privacy protections in surveillance processes. Previous work in this area includes that of Kamara [26] and Kroll *et al.* [27]; they proposed cryptographic protocols that protect the privacy of known surveillance targets. Segal *et al.* [41, 42] focused on unknown (*i.e.*, not yet

identified) targets and provided cryptographic protocols that protect privacy of innocent bystanders in two commonly used surveillance operations: set intersection and contact chaining. Frankle *et al.* [18] used secure, multiparty computation and zero-knowledge protocols to improve the accountability of electronic surveillance.

Attribute-based encryption was introduced by Sahai and Waters [40]. Their work was followed by many ciphertext-policy ABE and key-policy ABE constructions, including those in [4, 6, 20, 28, 35, 38]. Chase [10] introduced multi-authority ABE, and Nishide *et al.* [34] introduced ABE with hidden access policy. ABE has been applied in a wide range of domains, including fine-grained data-access control in cloud environments [44], health IT [2, 29], and security of blockchains and Internet-Of-Things devices [37, 45].

We now give a high-level explanation of some crucial differences between the role of proxies in OTABE and their roles in previous works.

An OTABE scheme provides an algorithm, *Translate()*, which allows a semi-trusted proxy to translate one or more of the attributes under which a data record's payload is encrypted without learning the underlying plaintext. Moreover, translation can be done obliviously, in the sense that the attributes under which the payload is encrypted remain hidden from the proxy who translates them; the proxy learns only the attributes' new values.

Two common responsibilities of proxies in ABE are *outsourced decryption*, which was introduced by Green *et al.* [22], and *revocation management*, which was used by Yu *et al.* [46, 47]. In both cases, proxies are used for efficiency; they assume much of the computational cost of decryption or revocation and lighten other parties' loads. The attribute-translation protocols in OTABE are *not* designed to reduce the client's or the owner's computational loads. Similarly, outsourced-decryption and revocation-management proxies are not designed to enable oblivious translation between organizational vocabularies or to support dynamically reconfigurable attributes – two of OTABE's primary goals. Simply put, proxies used for outsourced decryption and revocation management and those in OTABE serve completely different primary purposes.¹

The use of proxies for *ciphertext delegation* was introduced by Sahai *et al.* [39]. Proxies in this scenario take ciphertexts that are decryptable under policy P_1 and transform them into ciphertexts that are decryptable under policy P_2 . However, P_2 must be stricter than and use the same vocabulary as P_1 ; here, "stricter" means that P_2 permits the decryption of a subset of the ciphertexts that could be decrypted under the original policy P_1 used by the encryptor. Neither of these restrictions applies to the proxies in OTABE.

In *attribute-based proxy re-encryption* (ABPRE), which was introduced by Liang *et al.* [30], a proxy re-encrypts a ciphertext encrypted under access structure AS_1 to one that can be decrypted under access structure AS_2 without learning the plaintext. There is a surface similarity between ABPRE and OTABE in that proxies in both transform ciphertexts encrypted by data owners under AS_1 into ciphertexts decryptable by clients under AS_2 . However, the entity that issues re-encryption keys to proxies in ABPRE requires knowledge of the vocabularies of both owner and client; to create re-encryption keys she must know AS_1 and AS_2 . Thus, unlike

¹A direct-revocation mechanism, partially managed by the proxy, is a natural byproduct of attribute translation, as described in Sec. 4, but it is not the primary goal of OTABE.

OTABE, ABPRE does not support multiple vocabularies and can not provide attribute privacy.

In an ABPRE scheme, re-encryption keys are issued to a proxy on a per-access-policy basis. In order to perform re-encryption, the entire access policy must be changed so that the new policy contains no attributes that appear in the original policy. Neither of these restrictions applies to OTABE, in which re-encryption-key issuing and re-encryption itself can be done on a per-attribute basis. The responsibility for determining the new attribute set and performing the re-encryption is divided among multiple parties from different trust domains. Each party performs a partial re-encryption that uses only the attributes that belong to its trust domain and does so in a controlled manner that results in a final, full re-encryption that satisfies the data owner's requirements. This decentralized approach allows OTABE to support multiple vocabularies, provide attribute privacy, and enable dynamically reconfigurable translation policies that do not require re-initialization of the system or re-encryption of records by the owner.

Finally, in ABPRE, the proxy must know the ciphertext's original access policy in order to perform the re-encryption. OTABE proxies, by contrast, perform *oblivious* translation and re-encryption; they do not learn the original set of attributes or the original access structure under which the plaintext was encrypted.

2.2 Use cases

In order to motivate the notion of OTABE and illustrate its applicability in real-world scenarios, we provide two examples.

Law-enforcement agencies: The Electronic Communications Privacy Act (ECPA) [14] was passed to protect the privacy rights of ISPs' customers with respect to disclosure of their personal information. The ECPA limits LE access to email and other communication records in a manner that is consistent with the Fourth Amendment. However, it has several "loopholes." For example, the ECPA classifies an email message that is stored on a third party's server for more than 180 days as "abandoned." As a result, LE agencies can request that both the metadata and the content of those email messages be turned over without the need for judicial review.

Unrestrained government access to communication data is clearly undesirable. However, given national-security and public-health concerns, expecting LE and intelligence agencies never to access *any* data held by communication companies such as ESPs is unrealistic. A more realistic goal is to deploy a policy that restricts such data sharing to the minimum needed in order to perform the task at hand, as defined by multiple trusted entities. OTABE provides a mechanism that can enforce such policies and protect the confidential information of all organizations and agencies that participate in the data-sharing protocols.

In OTABE terms, the data owner is the ESP, and the data subjects are people who send and receive email messages. The data are email records. Each email record contains a payload, which is the content of an email message, encrypted under a set of metadata attributes, *e.g.*, sender's and receiver's email addresses, date, subject line, *etc.* The client is an LE agency, such as the FBI or a municipal police department, and the intermediaries may be other LE agencies, non-LE government agencies, or private companies. The data users are LE agents employed by the client.

Clearly, email records can be useful to LE agencies, but an agent should be able to decrypt only those records whose metadata attributes constitute *probable cause* in the context of a specific investigation. The entities who determine probable cause on a per-investigation basis are the TAs. Each TA is motivated by a different set of interests and goals. A TA may be part of the judicial branch, the ESP, the LE agency, or another external entity.

Not all of the attributes used by the ESP to store email records can be shared with the LE agency, because some of them reveal both private information about the ESP's customers or proprietary information of the ESP itself. Similarly, the attributes used by the LE agency to access and process records and to issue access policies cannot be shared with the ESP, because they reveal confidential information about the LE agency's investigations. Furthermore, some of the attributes that are used by the parties do not belong to the same vocabulary; for instance, the attribute "appears-in-watchlist" is frequently used in keys issued to LE agents, but it is meaningless to the ESP. Such attributes must undergo dynamic adaptation to ensure that agents' keys match an email message's attributes. OTABE allows the ESP and LE agency to use their own vocabularies while keeping the email messages' content and metadata confidential.

TAs are likely to grant an agent who is investigating a crime access to email records in which either the sender or the receiver is on the agency's watchlist. The LE agency's proxy can translate the ESP's sender and receiver attributes into the LE agency's "on-watchlist" attribute in an oblivious fashion, thus maintaining both the confidentiality of the watchlist and the privacy of data subjects' email addresses. In addition, an agent might want to check whether the sender or receiver appears on other agencies' lists, *e.g.*, a list of investigations ongoing at LEA-2, which is another LE agency. Because details of LEA-2's ongoing investigations cannot be shared with the client, the translation of the attributes sender and receiver will be done obligiously by LEA-2's intermediary proxy.

Similarly, the access policy of an agent who is investigating cyber fraud may enable access to email records in which the subject line matches a "suspicious" pattern. The definition of "suspicious" may be determined according to a dynamically updated list of keywords. Using this keyword list, the client's proxy can obliviously translate the attribute "subject line," maintained by the ESP, into the attribute "is-suspicious-subject," maintained by the client and used in the agent's access policy. Neither the agent nor the proxy is able to read the actual subject line, and the data subject's privacy is maintained.

Note that, in both of these investigations, dynamic translations are needed, because watchlists and lists of suspicious keywords change over time. They enforce the requirement that an agent cannot access payloads without probable cause, but they do not reveal to the ESP confidential information about watchlists and ongoing investigations.

Insurance companies: Consumer reporting agencies (CRAs) collect and share credit-related information about consumers. This information is used by credit-card issuers, mortgage lenders, insurance companies, *etc.* to assess creditworthiness of consumers. The three largest CRAs in the US are Experian, TransUnion, and Equifax.² The Fair Credit Reporting Act (FCRA) [16] regulates the collection,

²In September of 2017, Equifax announced a data breach that exposed the personal information of 147 million people and cost the company hundreds of millions of dollars in compensation to affected people [5, 17].

Table 1: Summary of notations and symbols.

Notation	Description	Notation	Description
$(M)_S$	encryption of M under a set S of attributes	$[x]_y$	encryption of x under the key y
ORG_S	set of proxies involved in translation of $C = (M)_S$	DEC_S	set of parties involved in decryption of $C = (M)_S$
P_{org_j}	the proxy operating on behalf of organization org_j	S_m	mutable attributes
S_{im}	immutable attributes	S_p	the set of attributes' labels that P_{org_p} is allowed to translate
$pub_{\Pi}(x)$	the public key of entity x , created by a public-key scheme Π	K_x	a symmetric shared key between org_{owner} and organization org_x
$org(k)$	the organization who is allowed to translate attribute att_k	$E_j(L)$	encryption of auxiliary information L by organization org_j
$F(K, x)$	pseudorandom function keyed with symmetric key K	$F(x)[0]$	the first argument of the output of the evaluation of F on x

dissemination, and use of credit-related information. The FCRA gives companies the right to access consumers' credit reports. This access is not limited to reports on the company's customers; it may include reports on large sets of potential customers. In order to create pre-screened offers and market them to potential customers, an insurance company is allowed to access consumers' credit reports and to share information with credit-card issuers, banks, other insurance companies, etc. However, access rights to credit reports are limited by the FCRA to information for which an insurance company has a *permissible purpose*. OTABE can be used to formalize and enforce this vague concept in a manner that protects both consumers' privacy and proprietary information of insurance companies and CRAs.

In OTABE terms, the data owner is a CRA, and the data subjects are consumers. Data records are credit reports, owned by the CRA. Each record is encrypted under the set of attributes that describe the report, e.g., the phone number, credit score, and driver's license number (DLN) of the data subject, credit-utilization ratio, date and time of the report's creation, CRA-internal statistics, etc.

Insurance companies are the data clients. Data users are insurance-company employees who use credit reports to make decisions about which insurance products to offer consumers and how to price them. In order to comply with the FCRA's "permissible-purpose" requirement, employees should only access credit reports on a "need-to-know" basis. An employee can only access those records whose associated attributes are relevant to her task, as determined by a set of TAs. TAs may include the CRA, a government entity, or various parties within the insurance company. Other organizations, such as credit-card issuers, government entities, banks, and other insurance companies may serve as intermediaries by "enriching" data supplied by a CRA in a privacy-preserving manner.

As in the LE scenario, each organization wants to protect the confidentiality of its proprietary information. For instance, the CRA does not want to reveal unnecessary identifying information about its customers, an insurance company does not want to reveal how it makes business decisions regarding which consumers are considered "qualified" for pre-screened offers, etc. Also as in LE, different organizations may use different vocabularies. Consider the attribute "number of accidents," which is used by insurance companies to screen potential customers. This attribute cannot be used by CRAs, because they do not maintain such information in their credit reports. OTABE supports all of these requirements.

Assume that each report is encrypted under the following attributes: CREDIT-UTILIZATION-RATIO, CREDIT-SCORE, PHONE-NUMBER, DLN, and DATE. Employee U in the car-insurance department is assigned the task of finding qualified potential customers

and tailoring pre-screened car-insurance offers, using information found in their credit reports.

The TAs determine that, for this task, a qualified customer is defined by the following policy:

$CREDIT-SCORE > X \wedge \#ACCIDENTS < Y \wedge IS-BLACKLISTED = FALSE \wedge IS-CREDIT-RATIO-LESS-THAN-AVERAGE = TRUE$

The intermediaries in this case are financial business partners of the insurance company, e.g., banks and credit-card issuers, and the Department of Motor Vehicles (DMV).

To complete her task, U submits to the CRA a query that requests the reports of all consumers whose credit scores are greater than X . The CRA then sends each matching record to two intermediaries: the DMV and a credit-card issuer.

For each record, the DMV's proxy obviously translates the DLN attribute into $\#ACCIDENTS$, which is found in the data subject's driving record. The credit-card issuer's proxy obviously translates the numeric CREDIT-UTILIZATION-RATIO attribute into a binary attribute IS-CREDIT-RATIO-LESS-THAN-AVERAGE by obviously comparing the consumer's utilization ratio with the average utilization ratio of the issuer's customers. The insurance company's proxy obviously translates the PHONE-NUMBER attribute into the attribute IS-BLACKLISTED, using a dynamically updated list of individuals who were blacklisted by the insurance company or one of its business associates for, e.g., failure to pay.

When U receives a record, she will be able to decrypt the credit report, read its contents, and learn the subjects' identifying information if and only if the record's post-translation attributes satisfy her access policy.

Data privacy is achieved, because only authorized users can decrypt a given credit report. Attribute privacy is achieved, because attributes used by each organization remain hidden to the extent required. Moreover, sensitive information about consumers whose records are decrypted is also protected; for example, a user may learn that a consumer's number of accidents is below a certain threshold but not learn the exact number. Finally, these translations demonstrate OTABE proxies' ability to translate *dynamically*, because the list and the average change over time, and *obliviously*, because neither the attributes nor the data are revealed to them.

3 ATTRIBUTE-BASED ENCRYPTION WITH OBLIVIOUS ATTRIBUTE TRANSLATION

3.1 Terminology

Attributes: Our scheme uses multi-valued attributes, denoted by $\langle label, operator, value \rangle$. Note that this representation is different from the ones found in typical ABE schemes, which use "descriptive"

(essentially binary) attributes. We denote by att_k^L and att_k^V the label and value of an attribute att_k and show how to implement ABE with non-descriptive attributes in [24]. Translation of an attribute can be done either by changing the attribute's value (*i.e.*, replacing $value$ with $value^*$) or by replacing both the attribute's label and its value with $label^*$ and $value^*$, respectively.

In PRShare, attributes' labels are partitioned into two sets: mutable, denoted S_m , and immutable, denoted S_{im} . Immutable attributes are ones that cannot be translated by any party in the system. Intuitively, they are the attributes that are shared by the owner and the client. Mutable attributes, on the other hand, are ones that can be translated by a semi-trusted proxy at some point after their initialization by the owner.

Hidden access policy: We introduce an OTABE scheme with hidden access policy by ensuring that the set of attributes used to encrypt a message is hidden from the CSP, the proxies, and the data users. We use the term "hidden access policy" for compatibility with the terminology used in existing CP-ABE work, in which access policies are attached to the ciphertexts.

In such a scenario, a data user cannot learn the attributes that are attached to a ciphertext but is able to determine which attributes are needed to perform the decryption. The hidden access policy feature is used to enhance privacy; however, if the owner and client wish to reveal the ciphertexts' attributes to the users or wish to speed up decryption at the expense of some privacy, they can turn off this feature without having to alter the encryption, translation, and decryption operations. This follows from the modular design of the system, as discussed in Sec. 5. Note that the hidden access policy feature does not enable the creation of trivial policies (*i.e.*, those that always allow a user to decrypt every record she receives). This is because a key must satisfy *all* TAs' policies in order to succeed in decrypting, and the data owner can always serve as a TA or delegate to a TA that it trusts not to permit decryptions that it wishes to forbid.

In general, PRShare is designed to achieve a high level of privacy while allowing flexible and expressive data-sharing protocols. In real-world scenarios, however, organizations have different priorities. Some may favor privacy, but others may favor functionality and thus prefer to allow their data users broader access to information about the shared data at the expense of privacy. PRShare is able to support both approaches: It is highly modular, and each privacy guarantee relies on a different low-level feature that can be removed or changed to fit the organization's privacy-functionality trade-offs while maintaining the rest of the privacy guarantees.

(Informal) Definition: Let M be a data record's payload encrypted under a set $S \subseteq \mathcal{U}_1$ of attributes, resulting in a ciphertext C . We refer to the set S as the set of original attributes under which M is encrypted. Let $T : \mathcal{U}_1 \rightarrow \mathcal{U}_2$ be a translation function from the universe \mathcal{U}_1 of attributes to the universe \mathcal{U}_2 of attributes, and let Q_j be the set of original attributes that a semi-trusted proxy j is allowed to translate. An ABE scheme supports **oblivious attribute translation by semi-trusted proxy** j if, given C , Q_j , and T , for all $s \in Q_j$, the proxy is able to compute $T(s)$ without:

- learning anything about M ,
- learning anything about the attributes in $S \setminus Q_j$, or
- learning the labels or the values of attributes in $S \cap Q_j$.

Formal security definitions are given in Sec. 5.

3.2 Algorithms

An MA-OTABE scheme consists of the following algorithms:

GlobalSetup(λ) \Rightarrow (PK): The global-setup algorithm takes as input a security parameter λ and outputs global parameters PK .

AuthoritySetup(PK) \Rightarrow (PK_i, MSK_i): Each authority runs the authority-setup algorithm with PK as input to produce its own public key PK_i and master secret key MSK_i .

Encrypt($M, PK, S, \{PK_i\}_{i \in Aut}$) \Rightarrow (CT): The encryption algorithm takes as input a message M , a set S of attributes, and the public parameters. It outputs the ciphertext CT .

KeyGen(PK, MSK_i, A_i, u, t) \Rightarrow ($SK_{i,u,t}$): The key-generation algorithm takes as input the global parameters, an access structure A_i , a master secret key MSK_i , the global identifier u of a data user who issued the key-generation request, and a task t . It outputs a decryption key $SK_{i,u,t}$.

Distribute(I) \Rightarrow ($\{C^j | j \in DEC_S\}$): This algorithm takes as input a set I of ciphertexts' ids. It outputs a set of partial ciphertexts, $\{C^j | j \in DEC_S\}$.

Translate($PK, j = p, C^p, \{PK_i\}_{i \in Aut}$) \Rightarrow (C^p): The translation algorithm takes as input the global public parameters and the authorities' public parameters, a proxy's index $j = p$, and a partial ciphertext C^p . It outputs a translated partial ciphertext C^p .

Decrypt($PK, \{SK_{i,u,t}\}_{i \in Aut}, C^u, \{C^j | j \in ORG_S\}$) \Rightarrow (M): The decryption algorithm takes as input the global parameters, a set of secret keys $\{SK_{i,u,t}\}_{i \in Aut}$, a partial ciphertext C^u , and a set of translated partial ciphertexts $\{C^j | j \in ORG_S\}$. It outputs the plaintext M .

4 SYSTEM MODEL

Definition of attributes: We define two sets of attributes' labels: S_{owner} represents the set of attributes that the owner uses to encrypt, store, and access data records that it owns. This set is determined by the data owner. S_{client} represents the set of attributes under which keys are generated; they are the attributes that the client uses to access and process the shared data records, and they are chosen by org_{client} . Note that $S_{owner} \cap S_{client} \neq \emptyset$; this means that some attributes are shared by the client and the owner. This enables the users to retrieve data records of potential interest from the CSP using queries that are composed of shared attributes and also enables the data owner, if it wishes, to be one of the TAs. We denote the universes of attributes comprising each set by \mathcal{U}_{owner} and \mathcal{U}_{client} .

For each data intermediary org_j in the system, we define a set of attributes' labels $S_j \subseteq S_m$. It represents the set of attributes that is governed by org_j and hence can be translated by the semi-trusted proxy P_{org_j} that acts on behalf of org_j .

4.1 System participants

Data owner: org_{owner} is responsible for encrypting each of its data records using the set $S \subseteq \mathcal{U}_{owner}$ of attributes that are most likely to appear in future queries.

Data users: Data users are employees of org_{client} who need access to data records stored by org_{owner} in order to perform daily tasks. Each user is assigned a unique global identifier and a list of tasks.

Each task t has a well defined time limit tl_t . The list is dynamic in the sense that tasks can be removed or added to it during the system run. A user issues two types of queries. A *key request* is used to obtain a key that corresponds to a specific access policy. A *data query* is used to obtain data records owned by org_{owner} that are relevant to a specific task in the user's task list.

Cloud-service provider: The CSP stores the ciphertexts outsourced by org_{owner} and responds to queries submitted by data users in org_{client} .

Trusted authorities: TAs are the entities that determine the decryption policy of org_{client} and issue secret keys that are used by data users; they use attributes from \mathcal{U}_{client} . There must be at least two TAs, and they may be entities in org_{owner} , org_{client} , or external organization. We assume that at least one TA belongs to org_{client} and that at least one TA does not.

Proxies: Each proxy P_{org_j} represents a different organization org_j (either an intermediary or a client) and operates on behalf of that organization. The role of a proxy P_{org_j} is to translate a subset of attributes in \mathcal{U}_{owner} under which a ciphertext was encrypted to the corresponding attributes in \mathcal{U}_{client} . To do this, the proxy uses both a generic translation algorithm that is used by all proxies in the system and an organization-specific translation function that is determined by org_j and may involve auxiliary information provided by the organization to its proxy. The generic translation algorithm is public, but the organization-specific translation function and auxiliary information are considered private to org_j and P_{org_j} . We assume that every MA-OTABE scheme includes at least one proxy (the "client proxy") that is responsible for managing org_{client} 's user-level revocation mechanism and for performing vocabulary translations.

Data subjects: Each data record owned by org_{owner} is linked to a certain individual, the data subject. A data record's payload contains personal information about the data subject, including content produced by the data subject. We assume that every data subject has a user id (UID) that varies based on the type of data used in the system. Examples of UIDs include phone numbers and email addresses.

4.2 Revocation mechanism

One major byproduct of OTABE is the ability to implement an efficient and direct revocation mechanism, in which revoking the keys of a set U of users does not affect the keys of users not in U . Using the translation technique, a semi-trusted mediator can transform a ciphertext that was encrypted under a set of data-centric attributes at point A into a "personalized" ciphertext reflecting a specific data query made by a user at point B. The main idea of our revocation mechanism is the addition of global-identifier (GID) and time attributes to each key; we also add a dummy GID and dummy times during encryption. These dummy attributes will be translated to suit the specific data query's time and requester only if a certain criterion is met. This creates an efficient mechanism in which most revocations are enforced automatically.

We assume that every data user receives a unique GID. The data client maintains a revocation list that contains revoked GIDs. Users whose GIDs are on the revocation list are not allowed to access any data record. Revocation-list updates are infrequent and happen

only when a user completely leaves the organization. Furthermore, GIDs can be removed from the revocation list after a relatively short time, because the key-level revocation mechanism ensures that secret keys become invalid within a well known and controlled length of time from the date they were issued.

For the key-level revocation mechanism, we leverage a basic trait of an organizational task: It has a well defined time limit. This time limit is determined by the user's manager and may change while the user is working the task. In our case, the entities who choose the time limit are the TAs; this is an integral part of the per-task "probable-cause" approach. The time limit given to a specific task performed by a user becomes an attribute in the user's key. In addition, the encryptor adds to each ciphertext a dummy "time" attribute. That dummy attribute is translated by the client proxy to the current time at which the data query is submitted by the user, thus making a key-level revocation check an automatic part of any decryption attempt. In our construction, we view a "time limit" as a date; this can easily be extended to include finer-grained notions of time.

We also leverage our attribute-translation technique for the user-level revocation mechanism. It enables us to include a user-specific component in the ciphertext; this component is adjusted according to the specific data user by the client proxy in the data-retrieval phase. Note that we treat the GID as an *additional attribute*. We incorporate the user's GID as an attribute in the user's secret keys and, in parallel, add a "placeholder" GID attribute to each ciphertext. When a user submits a data query, the placeholder attribute is translated to that specific user's GID only if she does not appear in the revocation list. This mechanism provides an efficient user-level revocation mechanism and protects the scheme from collusion attempts and key-abuse attacks.

Details of the translations used in our revocation mechanism are provided in [24].

4.3 Main flows

The system model consists of an encryption flow, a data flow, and a key-generation flow. We assume that the system has already been set up, resulting in the global public parameters PK and a public-key, master-secret-key pair (PK_i, MSK_i) for each trusted authority Aut_i .

Encryption flow: In order to encrypt a data record's payload M , org_{owner} first determines the set S of attributes under which M will be encrypted. $S \subseteq \mathcal{U}_{owner}$ is composed of $|S| - 2$ data-centric attributes that describe the record's metadata and two attributes that serve as "placeholders." The placeholders att_{GID} and att_{TIME} are initialized with random, "dummy" values by org_{owner} and receive their actual values from org_{client} 's proxy. Based on the attributes in S , the encryptor determines the set DEC_S of decryption parties. DEC_S contains all parties involved in the decryption of the ciphertext, *i.e.*, a data user and the set ORG_S of organizations that are allowed to translate attributes in S (represented by their proxies). ORG_S includes the client's proxy and any number of data intermediaries' proxies. After determining DEC_S , org_{owner} encrypts M under S by calling $Encrypt(M, PK, S, \{PK_i\}_{i \in Aut})$ and receives a set $\{C^j\}$ of $|DEC_S|$ partial ciphertexts. $|DEC_S| - 1$ of the partial

ciphertexts correspond to proxies and contain only attribute components. One corresponds to the data user and contains both attribute components and a data component; the latter contains the payload M itself. Note that, for each $C^j, U(C^j) \subseteq \mathcal{U}_{owner}$, where $U(C)$ is the vocabulary of attributes under which a ciphertext C is encrypted. Lastly, org_{owner} computes $Y = \{Obf(att_k) \mid att_k \in S\}$, a set of obfuscated values for immutable attributes in S , and uploads to the cloud the preprocessed ciphertext and the UID that the ciphertext is associated with.

Key-generation flow: A user u who belongs to org_{client} sends a key request to the TAs in each of the following cases: Either a new task is inserted to u 's task list, or the time limit for an existing task in u 's task list has expired, and her existing secret key for that task is no longer valid. The request contains a description of the task and the "ideal" access policy that u would like to obtain in the context of that task. Each authority Aut_i creates an access policy A_i based on an examination of the user's request and the nature of the specific task. It creates a GID attribute att_{GID} that contains the user's GID u . Finally, it determines tl_t , which is either a new time limit for t (if t is a new task) or an extended time limit (if t is an existing task and its time limit has expired) and uses tl_t to create a time-limit attribute att_{LIMIT} . The time-limit attribute that is embedded in a secret key must be expressed using the same units (date, month, time stamp, etc.) used in the time attribute att_{TIME} that is attached to the ciphertext. It then creates its secret key $SK_{i,u,t}$ by calling $KeyGen(PK, MSK_i, A'_i, u, t)$, where

$$A'_i = A_i \wedge att_{GID} \wedge att_{LIMIT} = A_i \wedge (GID == u) \wedge (TIME < tl_t).$$

Data flow: A data user u sends a data query to the CSP; it contains a conjunctive query ψ on attributes from $\mathcal{U}_{owner} \cap \mathcal{U}_{client}$. The CSP retrieves the ciphertexts that satisfy the query. For each ciphertext C , it sends $C^{j=u}$ to u and each $C^{j=p}$ to a proxy P_{org_p} . At that point, because u received only a partial ciphertext, she cannot yet use her key for decryption. Each proxy P_{org_p} in ORG_S translates each attribute att_k such that $(att_k \in S) \wedge (att_k^L \in S_p)$ by calling $Translate(PK, j = p, C^p, \{PK_i\}_{i \in Aut})$ and computes an obfuscated value for each new attribute $att_{k'}$ that it added, creating $Y_p = \{Obf(att_{k'})\}$. The client organization's proxy also manages the user-level mechanism by performing a correct translation of att_{GID} and att_{TIME} only if u does not appear in the revocation list. Each proxy P_{org_p} then sends the translated partial ciphertext $C^{j=p}$ and Y_p to the user. At this point, $U(C^j)$ has changed from \mathcal{U}_{owner} to \mathcal{U}_{client} . Because each partial ciphertext is, from the proxy's view, independent of the data component inside the ciphertext, each proxy is able to perform the translations without learning M . Moreover, the structure of each partial ciphertext ensures that P_{org_j} learns nothing about the attributes with labels that do not belong to S_j . All attribute components that correspond to attributes that the proxy can translate contain obfuscations of the attributes, rather than the attributes themselves; thus, each attribute att_k such that $(att_k \in S) \wedge (att_k^L \in S_p)$ remains hidden from the proxy, while the obfuscated value can still be used for various translation operations. The user gathers all the translated partial ciphertexts $\{C^j \mid j \in ORG_S\}$ and her partial ciphertext C^u to create an aggregated ciphertext that she can decrypt using her secret key. Finally, u decrypts the payload by calling

$Decrypt(PK, \{SK_{i,u,t}\}_{i \in Aut}, C^u, \{C^j \mid j \in ORG_S\})$. The decryption succeeds if and only if the following three conditions hold:

- $\forall i \in Aut, TR(S) \models A_i$, where $TR(S) = Y \cup \{Y_j\}_{j \in ORG_S}$ represents the set of translated attributes, created based on the original set S of attributes.
- tl_t , the time limit for task t , has not expired. (Otherwise, att_{LIMIT} cannot be satisfied.)
- u has not been revoked, and no collusion or key-abuse attempt has been made. (Otherwise, att_{GID} cannot be satisfied.)

5 SECURITY

5.1 Security goals and trust relationships

An OTABE-based framework should satisfy three security goals with respect to all PPT adversaries.

Selective security against chosen-plaintext attacks: The adversary cannot learn (in the selective-security model) the plaintext of either an original ciphertext or an aggregated, translated ciphertext.

Security against colluding parties: Let $C = (M)_S$ be a valid MA-OTABE ciphertext. No coalition of at most $|DEC_S| - 1$ parties can learn anything about M .

Attribute secrecy: The trust model that we consider in this paper is different from the standard ABE trust model. Unlike the plaintext, for which we have a single security notion that applies to all the participants, we cannot apply a uniform security criterion to the attributes. Because each party plays a distinct role in the protocol, the set of attributes to which it is allowed to be exposed differs from the sets to which other parties are allowed to be exposed. We define three security requirements to ensure the secrecy of ciphertexts' attributes: hidden access policy, oblivious translation, and attribute privacy.

Hidden access policy: The set of attributes used to encrypt a message cannot be learned by the CSP, the proxies, or the data users.

Oblivious translation: The original attributes that each proxy P_{org_j} translates remain hidden from the proxy. That is, for every attribute s such that $s^L \in S_j$, the proxy P_{org_j} is able to translate s into a new attribute $s' \in \mathcal{U}_{client}$ without learning s .

Attribute privacy: Informally, the attribute-privacy requirement states that organizations that share data must be able to maintain separate views of the data that they share.

Definition 5.1. Given a payload space \mathcal{M} , a universe \mathcal{U}_{owner} of attributes used by the encryptor (org_{owner}) to describe data records it owns, and a universe \mathcal{U}_{client} of attributes used by org_{client} for data usage and authorization management, we define a function $MAP : \mathcal{M} \times \mathcal{U}_{owner} \rightarrow \mathcal{U}_{client}$ that maps attributes in org_{owner} 's vocabulary (corresponding to data records' payloads $M \in \mathcal{M}$) to attributes in org_{client} 's vocabulary. An OTABE scheme achieves **attribute privacy** if and only if:

- For every data record's payload M and every attribute $s \in \mathcal{U}_{owner}$, if s is mutable, the encryptor does not learn $MAP(M, s)$, the translated value of the attribute s with respect to M .
- For every data record's payload M and every attribute $v \in \mathcal{U}_{client}$, if $MAP^{-1}(M, v)$ is mutable, data users and TAs do

not learn $MAP^{-1}(M, v)$, the original value of the attribute v with respect to M .

The following observations about our threat model, which considers external adversaries as well as the parties presented in Sec. 4.1, are natural aspects of the security definitions and results presented in Sec. 5.2.

No organization fully trusts the other organizations. Our framework protects the owner's data records, attributes of the data held by each organization, and auxiliary information held by each organization that is used for attribute translation. We assume that the owner is honest but curious.

No organization fully trusts its proxy server. CSPs and proxies in our framework, which we assume to be honest but curious, are only given encrypted attributes and encrypted auxiliary information. Note that the use of honest but curious proxies is well established in the design of cryptographic protocols [3, 7, 9, 21, 23, 47].

The client organization does not fully trust its data users. Data users in our system, who are assumed to be malicious, can only access records that are relevant to their assigned tasks, as determined by the TAs. We assume that at least one TA is honest. Data users also cannot learn attributes of the shared data records that are held by organizations other than the data client.

5.2 Security definitions and results

We start by presenting the definition of selective security for our MA-OTABE scheme.

Let $E = (Setup, AuthoritySetup, Encrypt, Distribute, KeyGen, Translate, Decrypt)$ be an OTABE scheme for a set of authorities Aut , $|Aut| = K$. Consider the following OTABE game for a PPT adversary \mathcal{A} , a challenger \mathcal{B} , a security parameter λ , an attribute universe \mathcal{U}_{owner} , and an attribute universe \mathcal{U}_{client} .

Init: The adversary chooses the challenge attribute set S , where $S \subseteq \mathcal{U}_{owner}$. Based on S , the adversary chooses the challenge decryption-parties set DEC_S^* , where $DEC_S^* \subseteq DEC_S$. The adversary also chooses a subset of corrupted authorities Aut_c . We assume that all authorities but one are corrupted and denote the honest authority by Aut_h ; thus, $Aut = Aut_c \cup \{Aut_h\}$. The adversary sends Aut_c , Aut_h , S , and DEC_S^* to the challenger.

Setup: The challenger runs the *Setup* algorithm to produce the public parameters PK and, for each authority Aut_i , runs the *AuthoritySetup* algorithm to produce PK_i and MSK_i . If Aut_i is honest, the challenger sends PK_i to the adversary. If Aut_i is corrupted, the challenger sends both PK_i and MSK_i to the adversary.

Phase 1: The adversary chooses a revocation list RL and sends it to the challenger. It may then issue any polynomial number of private key requests for tuples of the form (access structure, GID, task identifier) and send them to the challenger.

Given a request (access structure= $AC \in \mathcal{U}_{client}$, GID= u , task= t), the adversary proceeds as follows. For requests issued for a corrupted authority Aut_i , the adversary runs $SK_{iut} = KeyGen(PK, MSK_i, AC, u, t)$ itself, because it has MSK_i , which was given to it in the setup phase. For requests issued for the honest authority Aut_h , the challenger provides the answer. The challenger extracts the time limit tl_t from the description of task t and creates a time-limit attribute $att_{LIMIT} = \langle DATE, <, tl_t \rangle$. In addition, given the GID

u in the request, the challenger creates a GID attribute $att_{GID} = \langle GID, =, u \rangle$. It then creates $AC' = AC \wedge att_{LIMIT} \wedge att_{GID}$, which is an updated version of AC , and performs:

- If $S \models AC'$ and $u \notin RL$, the challenger aborts.
- If $S \models AC'$ and $u \in RL$, then S must contain $S_{GID} = u$. The challenger picks GID u' , $u' \neq u$, and generates the secret key using $SK_{hu't} = KeyGen(PK, MSK_h, AC, u', t)$.
- If $S \not\models AC'$, the challenger generates the secret key using $SK_{hut} = KeyGen(PK, MSK_h, AC, u, t)$.

Challenge: The adversary submits two messages m_0 and m_1 to the challenger. In addition, for every proxy j in DEC_S^* , it sends a bit a_j to the challenger. (By default, if j represents the user, we assume $a_j = 0$.) The challenger flips a fair coin b and encrypts m_b under S : $CT = Encrypt(m_b, PK, S, \{PK_i\}_{i \in Aut})$. Assuming I_{CT} is the index corresponding to the ciphertext CT , the challenger computes a set $\{C^j | j \in DEC_S^*\}$ of partial ciphertexts using $Distribute(I_{CT})$. For each proxy $j \in DEC_S^*$, if $a_j = 1$, the challenger performs a translation of the corresponding partial ciphertext, $C'^j = Translate(PK, j, C^j, \{PK_i\}_{i \in Aut})$, resulting in a translated partial ciphertext C'^j . Finally, it sends the ciphertext C^* to the adversary:

$$C^* = \bigcup_{j \in DEC_S^*} c_j^* \quad c_j^* = \begin{cases} C'^j & \text{if } a_j = 1 \\ C^j & \text{if } a_j = 0 \end{cases}$$

Phase 2: Phase 1 is repeated.

Guess: The adversary outputs a guess b' of b . The advantage of the adversary in this game is defined as $\Pr[b' = b] - 0.5$.

Definition 5.2. An MA-OTABE scheme is **selectively secure** if all PPT adversaries have negligible advantage with respect to λ in the selective-security game.

In the proof that our MA-OTABE construction is secure, we use a q -type assumption about prime-order bilinear groups: the **decisional q -Bilinear (t, n) -threshold Diffie-Hellman assumption** ((q, t, n) -DBTDH). It is similar to the Decisional q -Bilinear Diffie-Hellman assumption (q -DBDH) used in [38].

The assumption is parameterized by a security parameter λ , a suitably large prime p , two prime-order bilinear groups G_1 and G_2 , a bilinear map $e : G_1 \rightarrow G_2$, and integers q, t , and n , where $n \geq 1$ is polynomial in λ , and $t \leq n$. It is defined by a game between a challenger and an attacker. The attacker chooses a subset $V \subseteq [n]$ of t indices and sends it to the challenger. The challenger picks a group element g uniformly at random from G_1 , $q + 3$ exponents $x, y, z, b_1, b_2, \dots, b_q$ independently and uniformly at random from Z_p , and $n - 1$ additional exponents z_1, \dots, z_{n-1} independently and uniformly at random from Z_p . It sets $z_n = z - \sum_{c=1}^{n-1} z_c$. Then it sends (p, G_1, G_2, e) and the following terms to the attacker:

$$g, g^x, g^y, g^z, g^{(xz)^2} \\ \forall l \in [q] : g^{b_l}, g^{xz b_l}, g^{xz/b_l}, g^{x^2 z b_l}, g^{y/b_l^2}, g^{y^2/b_l^2} \\ \forall l, f \in [q], l \neq f : g^{y b_l/b_f^2}, g^{x y z b_l/b_f^2}, g^{(xz)^2 b_l/b_f}, \Psi_{l,f}$$

where

$$\Psi_{l,f} = \{g^{x z_c (b_l/b_f)} | c \in V\}.$$

The challenger flips a fair coin b . If $b = 0$, it gives the term $e(g, g)^{xy^z}$ to the attacker. Otherwise, it gives the attacker a term R chosen uniformly at random from G_2 . Finally, the attacker outputs its guess b' for the value of b .

Definition 5.3. We say that **the (q, t, n) -DBTDH assumption holds** if all PPT attackers have at most a negligible advantage in λ in the above security game, where the advantage is defined as $\Pr[b' = b] - 1/2$.

LEMMA 5.4. If $n \geq 2$ and $t \leq n$, then (q, t, n) -DBTDH \Rightarrow q -DBDH.

THEOREM 5.5. If (q, n, n) -DBTDH holds, then our MA-OTABE scheme achieves selective security against all PPT adversaries with a challenge attribute set S of size W , where $W \leq q$, and a challenge decryption-parties set DEC_S^* of size P , where $P \leq n$.

THEOREM 5.6. Let $C = (M)_S$ be a MA-OTABE ciphertext. No coalition of at most $|DEC_S| - 1$ parties can learn anything about M .

LEMMA 5.7. Let $C = (M)_S$ be a MA-OTABE ciphertext. The proxies in an MA-OTABE scheme cannot learn anything about M , even if they *all* collude.

THEOREM 5.8. Let F and F_p be two PRFs used in the construction of our MA-OTABE scheme. If F and F_p are secure, then the scheme achieves attribute secrecy; this includes hidden access policy, oblivious translation, and attribute privacy.

6 CONSTRUCTION

We now present a simplified version of our MA-OTABE construction. Specifically, we present the basic techniques that are new to OTABE in Sec. 6.1 and some key features of the *Translate()* function in Sec. 6.2. Recall that the full construction and proofs of all of results given in Sec. 5 can be found in [24].

6.1 Main techniques

Our scheme is inspired by the single-authority, large-universe ABE scheme of Rouselakis and Waters [38], but it is multi-authority.

Ciphertext composition is [38] is given by these equations:

$$C_0 = Me(g, g)^{s\alpha} \quad C_1 = g^s \quad C_{2k} = g^{f_k} \\ C_{3k} = (\theta^{att_k h})^{f_k} (w)^{-s}$$

where M denotes the message to be encrypted, α denotes the authority's private key, (g, θ, h, w) are the global public parameters chosen in the setup phase, and s and f_k are elements of Z_p , for a suitably large prime p .

The ciphertext is composed of a data layer and an attribute layer. We refer to C_0 and C_1 as *data-layer components*, C_2 and C_3 as *attribute-layer components*, and each element in C_3 as an *attribute component*. The data-layer component C_0 in [38] contains the message M masked by the public key $e(g, g)^\alpha$ of the (single) TA. Assuming that M is encrypted under a set S of attributes, the attribute layer contains $2|S|$ components, *i.e.*, two $(C_{2k}$ and $C_{3k})$ for each attribute att_k in the ciphertext. Each pair contains a uniformly randomly chosen term f_k that is local to the specific attribute att_k ; C_{3k} also contains the attribute att_k itself. The two layers are connected by the binder term s .

The basic idea of our construction is as follows. Assume that we have a data owner, a data client, two authorities (denoted Aut_1 and

Aut_2), a data user u , and a client proxy.³ Assume that the keys given to u by Aut_1 and Aut_2 are based on the access structures $att_1 \vee att_2$ and att_4 , respectively.

The data owner wishes to encrypt a record M under a set of attributes $S = att_1, att_3$, where $att_1 \in U_{owner} \cap U_{client}$, but $att_3 \notin U_{owner} \cap U_{client}$. That is, att_3 does not belong to the client's vocabulary and hence needs to undergo translation before it can be used for decryption by u , using the keys she received from the authorities. In this example, we assume that $T(att_3) = att_4$; that is, a correct translation of the attribute $att_3 \in U_{owner}$ is $att_4 \in U_{client}$. In order to encrypt M , the owner produces a two-level ciphertext; it is similar to the one in [38] but differs in the following four respects.

First, instead of creating $|S|$ attribute components C_{3k} , one for each attribute, the owner creates $|S| * |DEC_S|$ attribute components $C_{3k,j}$, one for each pair (attribute, decryption party), where DEC_S is the set of parties that participate in the decryption of the ciphertext (the *decryption-parties set*). In this example, $|DEC_S| = 2$, because the two decryption parties are the user and the client proxy.

Second, we use the binder term s differently from the way it is used by Rouselakis and Waters. In [38], the binder term is used in the data layer and in each attribute component. By contrast, we use secret sharing to break s into $|DEC_S|$ shares; each attribute component $C_{3k,j}$ contains only one share of the binder term – the one that corresponds to the decryption party j . In this example, the two decryption parties, *i.e.*, the user and the client proxy, each receive one share.

Third, each attribute component in [38] contains the actual attribute to which it corresponds. In our OTABE scheme, however, each attribute component contains the output of a given transformation that is applied to the attribute. This enables the proxy to translate the attribute *obliviously* without knowing its label or value. In our construction, the transformation is a keyed PRF, but, as explained in Sec. 6.2, OTABE can accommodate a richer set of transformations in order to better serve each organization's business logic.

Fourth, we use another uniformly randomly chosen term, which we denote by l_k . Like f_k , the term l_k is local to the attribute component in which it appears. It is used to double blind the attribute part $(\theta^{att_k h})$ of each attribute component, using $d_k = f_k * l_k$ as a blinding factor. In this way, f_k can be used by the proxy as a token for oblivious translation.

Because of the composition of the ciphertext, the proxy is able to translate the attribute $att_3 \in U_{owner}$ into a new attribute $att_4 \in U_{client}$. The proxy uses the attribute component $C_{3att_3, proxy}$, an obfuscated version of the original attribute att_3 , the *Translate()* algorithm, and tokens that it obtains from the *Encrypt()* algorithm. In general, determination of the new attribute is done obliviously based on the obfuscated original attribute's label and value.

When the user receives the translated record (corresponding to the new attribute att_4) from the proxy, she combines it with her own attribute-layer components and data-layer components to create the final aggregated ciphertext. She uses the keys that she received from Aut_1 and Aut_2 to decrypt the aggregated ciphertext. Decryption uses secret sharing and the unique structure of the translated attribute component received from the proxy, which

³For clarity, we do not use intermediaries in this simplified construction.

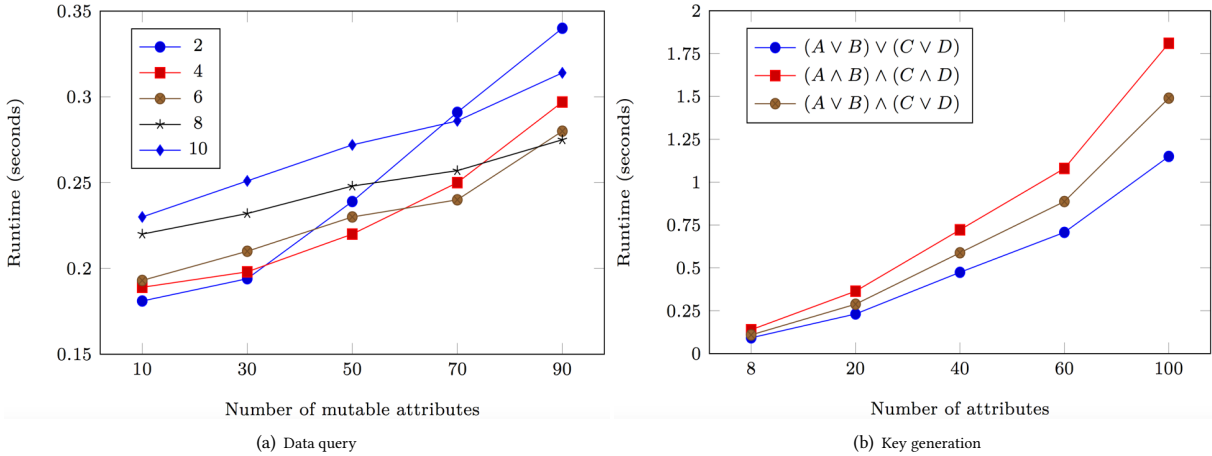


Figure 1: Typical running times in seconds

includes both an obfuscated version of the original attribute att_3 and the new attribute att_4 .

Finally, to enable hidden access policy, we do not attach the actual set of attributes S to the ciphertext. Instead, both the data owner and the proxy compute an obfuscated value of each attribute they add to the ciphertext, based on the PEKS construction given in [8]. Using trapdoors received from the TAs, u is able to perform a “blind intersection” of the obfuscated values received with the ciphertext and her own obfuscated access structure’s attributes received from the TAs. Thus, u is able to determine which attributes are required for decryption without learning their values.

6.2 Oblivious translation

The $Translate()$ algorithm assumes the existence of a set of translation functions. Each translation function T_j is determined by an organization org_j and specifies how to translate attributes in S_j .

Translation of attributes in our scheme is done for two purposes: revocation management (as described in Sec. 4.2) and dynamic adaptation of intraorganizational attributes (in order to express them in other organizations’ vocabularies). One of the main reasons that attribute translation is essential to support multiple vocabularies is that the encryption of a data record’s payload is done once by the owner, but the relevance of the data record to the client changes over time. In ABE terms, this means that the set of attributes under which a ciphertext is encrypted is taken from one vocabulary and does not change, but the question of whether or not this set satisfies a given access policy, which is expressed in another vocabulary, does change over time. Furthermore, deciding whether a ciphertext is relevant to the client at a given point in time is done using external “auxiliary information” that is related to one or more of the owner’s, client’s, and intermediaries’ professional domains. The auxiliary information changes over time, as explained in Sec. 2.2, and thus the decision about whether the set of attributes of a given data record satisfies a given access policy does as well. Values of such attributes with respect to a data record cannot be fully determined at encryption time; they must be dynamically translated

at precisely the time that a user needs to access that data record. OTABE supports such dynamic attributes.

Translation of mutable attributes by a semi-trusted proxy can be done in two ways. The first is by leaving the attribute’s label unchanged but changing its value. This is usually what is done for “dummy” attributes that are given random initial values by the owner. The second is by changing both the label and the value of the attribute; in this case, translation requires auxiliary information (a number, a list, *etc.*) that is provided to the proxy by its organization. The attribute inside each attribute component is encrypted using a specific transformation. In addition, each piece x of auxiliary information is encrypted using the same transformation that is used to encrypt the attribute that x is used to translate. The translation is done by extracting the encrypted attribute from its corresponding attribute component and performing an oblivious operation on both the attribute and the auxiliary information. Since the attribute inside the ciphertext and the organization-specific auxiliary information are encrypted using the same keyed transformation, with a key that the proxy does not know, the proxy can perform the translation without learning the attribute’s value and without learning the contents of the private auxiliary information provided by the organization.

For simplicity, we give a full construction that applies the same transformation to each attribute in the ciphertext, using two PRFs. This construction demonstrates a specific translation operation in which the proxy performs oblivious equality tests and set-membership tests to determine the new attribute. However, PRShare supports the more flexible approach in which different transformations are applied to different attributes in the ciphertext, based on the attributes’ types and sensitivities. For example, if $att_k \in \mathcal{U}_{owner}$ is a numerical attribute, the proxy can translate it into a descriptive attribute $att'_k \in \mathcal{U}_{client}$ by comparing att_k with a threshold that was provided to it by the organization that it represents. It determines the value of the new, descriptive attribute according to the result of that comparison. In such a case, we would choose an order-preserving transformation instead of an equality-preserving

transformation. Based on this modular approach and other PRF-based transformations, PRShare enables a broader set of translation operations that better suit various organizations' translation logic; these operations include oblivious addition, keyword search, and numerical comparison [11].

7 IMPLEMENTATION AND EVALUATION

To assess the feasibility of our framework, we implemented the full version of our OTABE scheme using Charm, a framework developed for rapidly prototyping advanced cryptosystems [1]. Charm was used to develop multiple, prominent existing ABE schemes, including that of Rouselakis and Waters [38]. We use a BN256 curve for pairings.

We consider a setting with three authorities and policies of size ten, where the decryption is always successful, and use oblivious list membership as our translation operation. We present benchmarks for two operations that are done frequently and online. The first is the overall turnaround time of a data query, *i.e.*, the total time between a user's initiation of a query and her receiving the *plaintext* records that satisfy it. We also provide benchmarks for the key-generation algorithm, despite the fact that key requests are significantly less frequent than data queries. The overall runtime, as shown in Figure 1, includes computation, communication, and I/O time. Note that the hidden access policy feature is turned off in our experiments. We do not include the time needed for system initialization or for the data owner to encrypt its data, because these operations are done once and offline.

Recall that each data query entails the following steps. A query is sent to the CSP. The CSP searches for all of the records that satisfy the query. For each ciphertext returned by the search, the CSP sends its partial ciphertexts to the relevant proxies. Each proxy obviously translates the partial ciphertext it received. The user aggregates all partial ciphertexts and decrypts the result to obtain the plaintext.

To enable adequate comparison of our OTABE scheme and previous ABE schemes such as the one in [38], results are given for a single-record data query. When generalizing our results to the multi-record case, it is important to note that our scheme is highly parallelizable. No TA or proxy needs to coordinate its computation with any other TA or proxy; thus they can all proceed in parallel. In order to decrypt, a data user must perform a separate computation for each TA, and all of these computations can be done in parallel. Finally, partial ciphertexts that correspond to different attributes can be translated in parallel.

Figure 1(a) compares the average time of a data query that contains 100 attributes, for different numbers of mutable attributes and various sizes of ORG_S . The runtimes are relatively small; it takes only 314ms to perform a 90-translation data query when $ORG_{SS} = 10$. Although there is an increase in runtime as the number of mutable attributes increases, this increase is significantly more noticeable when ORG_S contains fewer proxies. Figure 1(a) also demonstrates an inherent trade-off between the translation and decryption algorithms: A larger number of proxies results in better load balancing of translation operations, but it also results in more expensive decryption.

Figure 1(b) shows the average time taken by the key generation algorithm for various policies. The times are all under 1.81s; this means that, within less than two seconds from a data user's request for a task-related key, she will receive, from each authority, a key that supports a policy of size 100. Bear in mind that key requests are significantly less frequent than data queries and only occur once per time-limited task.

8 CONCLUSIONS AND OPEN PROBLEM

We have proposed PRShare, an interorganizational data-sharing framework that protects the privacy of data owners, data clients, and data subjects. In designing PRShare, we have introduced the novel concept of *Attribute-Based Encryption With Oblivious Attribute Translation*, which may be of independent interest. In future work, we will consider relaxing one or more assumptions that PRShare relies on; for example, we will explore the use of malicious proxies.

9 ACKNOWLEDGMENTS

We wish to thank Babis Papamanthou and Satyanarayana Vusirikala for their helpful comments. The first author was supported in part by US National Science Foundation grants CNS-1407454 and CNS-1409599 and William and Flora Hewlett Foundation grant 2016-3834. The second author was supported in part by US National Science Foundation grants CNS-1407454 and CNS-1409599 and US Office of Naval Research grant N00014-18-1-2743.

REFERENCES

- [1] Joseph A. Akinyele, Christina Garman, Ian Miers, Matthew W. Pagano, Michael Rushanan, Matthew Green, and Aviel D. Rubin. 2013. Charm: a framework for rapidly prototyping cryptosystems. *Journal of Cryptographic Engineering* 3, 2 (2013), 111–128.
- [2] Joseph A. Akinyele, Matthew W. Pagano, Matthew D. Green, Christoph U. Lehmann, Zachary N. J. Peterson, and Aviel D. Rubin. 2011. Securing Electronic Medical Records Using Attribute-based Encryption on Mobile Devices. In *1st ACM Workshop on Security and Privacy in Smartphones and Mobile Devices*. 75–86.
- [3] Giuseppe Ateniese, Kevin Fu, Matthew Green, and Susan Hohenberger. 2005. Improved Proxy Re-Encryption Schemes with Applications to Secure Distributed Storage. In *12th Network and Distributed System Security Symposium*. 29–43.
- [4] Nuttapon Attrapadung, Benoît Libert, and Elie de Panafieu. 2011. Expressive Key-Policy Attribute-Based Encryption with Constant-Size Ciphertexts. In *14th International Conference on Practice and Theory in Public-Key Cryptography*. Springer LNCS volume 6571, 90–108.
- [5] Tara Siegel Bernard, Tiffany Hsu, Nicole Perlroth, and Ron Lieber. 2017. Equifax Says Cyberattack May Have Affected 143 Million in the U.S. *The New York Times* (Sept. 7, 2017).
- [6] John Bethencourt, Amit Sahai, and Brent Waters. 2007. Ciphertext-Policy Attribute-Based Encryption. In *28th IEEE Symposium on Security and Privacy*. 321–334.
- [7] Matt Blaze, Gerrit Bleumer, and Martin Strauss. 1998. Divertible Protocols and Atomic Proxy Cryptography. In *17th EUROCRYPT*. Springer LNCS volume 1403, 127–144.
- [8] Dan Boneh, Giovanni Di Crescenzo, Rafail Ostrovsky, and Giuseppe Persiano. 2004. Public-Key Encryption with Keyword Search. In *23rd EUROCRYPT*. Springer LNCS volume 3027, 506–522.
- [9] Dan Boneh, Xuhua Ding, Gene Tsudik, and Chi-Ming Wong. 2001. A Method for Fast Revocation of Public Key Certificates and Security Capabilities. In *10th USENIX Security Symposium*. 22–22.
- [10] Melissa Chase. 2007. Multi-authority Attribute Based Encryption. In *4th Theory of Cryptography Conference*. Springer LNCS volume 4392, 515–534.
- [11] Nathan Chenette, Kevin Levi, Stephen A. Weiss, and David J. Wu. 2016. Practical Order-Revealing Encryption with Limited Leakage. In *23rd International Conference on Fast Software Encryption*. Springer LNCS volume 9783, 474–493.
- [12] Sabrina De Capitani di Vimercati, Sara Foresti, Sushil Jajodia, Stefano Paraboschi, and Pierangela Samarati. 2011. Private Data Indexes for Selective Access to

- Outsourced Data. In *10th ACM Workshop on Privacy in the Electronic Society*. 69–80.
- [13] Xin Dong, Jiadi Yu, Yuan Luo, Yingying Chen, Guangtao Xue, and Minglu Li. 2013. Achieving an Effective, Scalable and Privacy-preserving Data Sharing Service in Cloud Computing. *Computers and Security* 42 (2013), 151–164.
- [14] ECPA 1986. Electronic Communications Privacy Act, Public law 99-508. <https://it.ojp.gov/PrivacyLiberty/authorities/statutes/1285>.
- [15] Benjamin Fabian, Tatiana Ermakova, and Philipp Junghanns. 2015. Collaborative and secure sharing of healthcare data in multi-clouds. *Information Systems* 48 (2015), 132–150.
- [16] FCRA 1970. Fair Credit Reporting Act, Public law 91-508. <https://www.consumer.ftc.gov/articles/pdf-0111-fair-credit-reporting-act.pdf>.
- [17] Federal Trade Commission. 2017. Equifax Data Breach Settlement. <https://www.ftc.gov/enforcement/cases-proceedings/refunds/equifax-data-breach-settlement>.
- [18] Jonathan Frankle, Sunoo Park, Daniel Shaar, Shafi Goldwasser, and Daniel Weitzner. 2018. Practical Accountability of Secret Processes. In *27th USENIX Security Symposium*. 657–674.
- [19] David Froelicher, Patricia Egger, Joao Sa Sousa, Jean Louis Raisaro, Zhicong Huang, Christian Mouchet, Bryan Ford, and Jean-Pierre Hubaux. 2017. UnLynx: A Decentralized System for Privacy-Conscious Data Sharing. *Proceedings on Privacy Enhancing Technologies* 2017, 4 (2017), 232–250.
- [20] Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. 2006. Attribute-Based Encryption for Fine-Grained Access Control of Encrypted Data. In *13th ACM Conference on Computer and Communications Security*. 89–98.
- [21] Matthew Green and Giuseppe Ateniese. 2007. Identity-Based Proxy Re-encryption. In *5th International Conference on Applied Cryptography and Network Security*. Springer LNCS volume 4521, 288–306.
- [22] Matthew Green, Susan Hohenberger, and Brent Waters. 2011. Outsourcing the Decryption of ABE Ciphertexts. In *20th USENIX Security Symposium*. 523–538.
- [23] Luan Ibraimi, Milan Petkovic, Svetla Nikova, Pieter H. Hartel, and Willem Jonker. 2009. Mediated Ciphertext-Policy Attribute-Based Encryption and Its Application. In *10th international conference on information security applications*. 309–323.
- [24] Lihi Idan and Joan Feigenbaum. 2020. PRShare: A Framework for Privacy-Preserving Interorganizational Data Sharing, Technical Report YALEU/DCS/TR-1554. <https://cpsc.yale.edu/sites/default/files/files/tr1554.pdf>.
- [25] Sushil Jajodia, Witold Litwin, and Thoms Schwarz. 2011. Privacy of Data Outsourced to a Cloud for Selected Readers through Client-Side Encryption. In *10th ACM Workshop on Privacy in the Electronic Society*. 171–176.
- [26] Seny Kamara. 2014. Restructuring the NSA Metadata Program. In *2nd Financial Cryptography Workshop on Applied Homomorphic Cryptography and Encrypted Computing*. Springer LNCS volume 8438, 235–247.
- [27] Joshua A. Kroll, Edward W. Felten, and Dan Boneh. 2014. Secure protocols for accountable warrant execution. <https://www.cs.princeton.edu/~felten/warrant-paper.pdf>.
- [28] Allison Lewko, Tatsuaki Okamoto, Amit Sahai, Katsuyuki Takashima, and Brent Waters. 2010. Fully Secure Functional Encryption: Attribute-Based Encryption and (Hierarchical) Inner-Product Encryption. In *29th EUROCRYPT*. Springer LNCS volume 6110, 62–91.
- [29] Ming Li, Shucheng Yu, Kui Ren, and Wenjing Lou. 2010. Securing Personal Health Records in Cloud Computing: Patient-Centric and Fine-Grained Data Access Control in Multi-owner Settings. In *6th International ICST Conference on Security and Privacy in Communication Networks*. Springer LNICST volume 50, 89–106.
- [30] Xiaohui Liang, Zhenfu Cao, Huang Lin, and Jun Shao. 2009. Attribute based proxy re-encryption with delegating capabilities. In *4th ACM Symposium on Information, Computer, and Communications Security*. 276–286.
- [31] Chang Liu, Xiao Shaun Wang, Kartik Nayak, Yang Huang, and Elaine Shi. 2015. OblivVM: A Programming Framework for Secure Computation. In *36th IEEE Symposium on Security and Privacy*. 359–376.
- [32] Xuefeng Liu, Yuqing Zhang, Boyang Wang, and Jingbo Yan. 2013. Mona: Secure Multi-Owner Data Sharing for Dynamic Groups in the Cloud. *IEEE Transactions on Parallel and Distributed Systems* 24, 6 (2013), 1182–1191.
- [33] Kartik Nayak, Xiao Shaun Wang, Stratis Ioannidis, Udi Weinsberg, Nina Taft, and Elaine Shi. 2015. GraphSC: Parallel Secure Computation Made Easy. In *36th IEEE Symposium on Security and Privacy*. 377–394.
- [34] Takashi Nishide, Kazuki Yoneyama, and Kazuo Ohta. 2008. Attribute-Based Encryption with Partially Hidden Encryptor-Specified Access Structures. In *6th International Conference on Applied Cryptography and Network Security*. Springer LNCS volume 5037, 111–129.
- [35] Rafail Ostrovsky, Amit Sahai, and Brent Waters. 2007. Attribute-Based Encryption with Non-Monotonic Access Structures. In *14th ACM Conference on Computer and Communications Security*. 195–203.
- [36] Raluca Popa, Catherine M. S. Redfield, Nikolai Zeldovich, and Hari Balakrishnan. 2011. ACrypDB: Protecting Confidentiality with Encrypted Query Processing. In *23rd ACM Symposium on Operating Systems Principles*. 85–100.
- [37] Yogachandran Rahulamathavan, Raphael C.-W. Phan, Muttukrishnan Rajarajan, Sudip Misra, and Ahmet Kondoz. 2017. Privacy-preserving blockchain based IoT ecosystem using attribute-based encryption. In *11th IEEE International Conference on Advanced Networks and Telecommunications Systems*.
- [38] Yannis Rouselakis and Brent Waters. 2013. Practical constructions and new proof methods for large universe attribute-based encryption. In *20th ACM Conference on Computer and Communications Security*. 463–474.
- [39] Amit Sahai, Hakan Seyalioglu, and Brent Waters. 2012. Dynamic Credentials and Ciphertext Delegation for Attribute-Based Encryption. In *32nd CRYPTO*. Springer LNCS volume 7417, 199–217.
- [40] Amit Sahai and Brent Waters. 2005. Fuzzy Identity-Based Encryption. In *24th EUROCRYPT*. Springer LNCS volume 3494, 457–473.
- [41] Aaron Segal, Joan Feigenbaum, and Bryan Ford. 2016. Open, privacy-preserving protocols for lawful surveillance. *CoRR* abs/1607.03659 (2016). <http://arxiv.org/abs/1607.03659>
- [42] Aaron Segal, Joan Feigenbaum, and Bryan Ford. 2016. Privacy-Preserving Lawful Contact Chaining [Preliminary Report]. In *15th ACM Workshop on Privacy in the Electronic Society*. 185–188.
- [43] Dhinakaran Vinayagamurthy, Alexey Gribov, and Sergey Gorbunov. 2019. StealthDB: a Scalable Encrypted Database with Full SQL Query Support. *Proceedings on Privacy Enhancing Technologies* 2019, 3 (2019), 370–388.
- [44] Guojun Wang, Qin Liu, and Jie Wu. 2010. Hierarchical Attribute-based Encryption for Fine-grained Access Control in Cloud-Storage Services. In *17th ACM Conference on Computer and Communications Security*. 735–737.
- [45] Xuanxia Yao, Zhi Chen, and Ye Tian. 2015. A lightweight attribute-based encryption scheme for the Internet of Things. *Future Generation Computer Systems* 49 (2015), 104–112.
- [46] Shucheng Yu, Cong Wang, Kui Ren, and Wenjing Lou. 2010. Achieving Secure, Scalable, and Fine-grained Data Access Control in Cloud Computing. In *29th IEEE Conference on Computer Communications*. 534–542.
- [47] Shucheng Yu, Cong Wang, Kui Ren, and Wenjing Lou. 2010. Attribute-based data sharing with attribute revocation. In *5th ACM Symposium on Information, Computer, and Communications Security*. 261–270.