

create tighter regulations for privacy and enforce larger penalties for non-compliance [2]. Another solution involves a new programming language that enforces privacy policies by constraining certain variables that correspond to user data [3]. These are good solutions, and they are important steps towards giving the client more control. One of the biggest issues with the current state of privacy, though, is that even if a user edits their privacy settings, often they are still left unclear as to how their data is managed, and frankly can't tell if it was mismanaged anyway. Figure 1 below shows how clients lose their control after sending it o to a service.




Figure 1: Current Client/Server Relationship

The approach investigated here will not go so far as to ensure the data is used properly, but it will create a 'promise' that the service will implicitly make with every instance of the data being used. This scheme uses a type of Key-Policy encryption called Attribute-Based Encryption. It requires a client to share its data with a server, which will perform some action with the data, and a third-party authority to create, store and distribute keys necessary for the data to be encrypted and decrypted. The prototype for such a system has been built as a web app called privateBook. In privateBook, clients create an account and set a privacy policy, then can write posts that they can view if signed in to their own account. The prototype shows a practical way of ensuring data is encrypted for all parties, and can only be seen as a result of the service adhering to a user's privacy policy.

2 Data Encryption

At a high level, Attribute-Based Encryption (ABE) is a form of Public-Key cryptography which encrypts data with an access policy. A party attempting to decrypt the data must present an attribute list along with the private key. This attribute list must satisfy the access policy of the encrypted data in order for decryption to occur. We will discuss exactly how ABE uses policies and how to encrypt and decrypt data, and then how the Python implementation uses this type of encryption.

2.1 Public-Key Cryptography

In many systems of encryption there is one key that is used to encrypt and decrypt messages, called a cipher. This is analogous to one key that may lock or unlock a door. Public-Key cryptography is also called asymmetric cryptography, due to its use of two separate keys, one used to decrypt and one used to encrypt a message. In this form, there is both a public key and private key. The public key is used to encrypt data, while the private key is used to decrypt. The keys are generated with some function that creates a public key which does not allow a foreign party to discover the private key. This is based on certain problems which are essentially impossible to solve computationally, such as large integer factorization problems.

2.2 Attribute-Based Encryption

There are four portions of the ABE scheme: the parameter setup, encryption, key generation, and decryption.

Setup The setup simply generates random groups, which are stored as G , the public parameters, and g , a preliminary version of the private key. G is the ABE version of a public key. The two keys here are bilinear group generators raised to randomly chosen exponents. These are stored as Python dictionaries.

Encryption Encryption uses G , g , the access structure, along with the string message meant to be encrypted M . A dictionary E is constructed with the following structure: $E = \{ (A, B) \}$ where A is a list of attributes and B is the encrypted message.

Key Generation Key Generation uses a list of attributes A , and MK to generate a decryption key, D . This is the ABE version of a private key, used to decrypt. A is the list of attributes that must satisfy in order for D to successfully decrypt E .

Decryption Decryption takes E , the ciphertext containing C , PK and D . It applies D to E in order to decrypt the message. If A satisfies P , in the sense of the attributes fulfilling the boolean expression included in D , the message will decrypt.

2.3 Example Usage

An example following the paper describing ABE's fine-grained access capabilities can demonstrate how this system would play out if used in exactly the way described above [5].

If there were a system for storing activity logs on a network, this data would need to be protected by an encryption scheme that would vary for different types of information and different types of activity. Here, ABE encryption would be useful, as each log could be stored with a specific access policy: say "user is Bob or Alice AND the date is between September 2010 and May 2014 AND the activity is related to updating or changing the financial information of projects". This information would then be encrypted with this policy. Anyone investigating the logs could then be given a secret key with a specific access list with their properties: user name, title, security clearance, activity type, date, etc. ABE would then only allow the analyst to access information if their information fit with the policy: all other data, which is not pertinent to their work, will be inaccessible.

3 ABE Privacy Prototype

The implementation of ABE is provided by Charm, a Python framework providing many different crypto systems [4]. It is used to power the encryption in privateBook. This can apply to any service which uses client data to perform some function, whether it be like Facebook, just storing and displaying messages, text or multimedia, or something like Google Maps, which uses client location info to display a map or directions. In a sense, the prototype applies ABE in a backwards fashion: rather than encrypting with a policy, and

users attempt to decrypt by presenting their attributes, here the attribute list is created and clients present their policies. Figure 2 displays the process.




Figure 2: Prototype encryption/decryption process

3.1 Parties Involved

In privateBook, there are three relevant parties involved: the server, which provides the service, a client, who uses the service, and the authority, a third party that creates, distributes and stores the keys for both parties. The server creates default settings for privacy, saving them as an attribute list.

1. Server Registration This is the first step, and only happens once. The server registers with the Authority, naming itself and providing a list of attributes, A , which is a list of the default values created for a privacy policy. The Authority stores A .

2. Key Distribution This happens once, in three parts. The Authority generates PK and MK , the public parameters and master key, which are used to generate $D.A$, PK and D are stored by the Authority. Then, D is sent to the server, which will use it to decrypt data, and PK is sent to the client, who uses it to encrypt data.

3. Client Encryption This will happen with every client request. The client encrypts by providing their data, ρ , their personal privacy policy, and PK . This yields E , a cipher text, which is sent to the server. The server then stores the data as a cipher text.

4. Server Response This happens after every client request. The server stores the encrypted data, and attempts to decrypt it by applying the de-

table, and a homepage, where the client can write their updates and have them displayed; each update is stored by the server as a cipher text with a user ID in a Posted Data table. We will demonstrate each of the four steps in the privateBook implementation.

4.1 privateBook Server Registration

The privateBook server registers its name and default privacy settings with the Authority, which saves them as an attribute list, and then creates PK and D, the public parameters and private key. Figure 3 shows the printed values created from the default privateBook privacy values.

```
The Attribute List is:
[{ADDRESS: 1234, ADDELETED: 1, TARGET: 0, TRACK: 1, INFO: 0, SET: 1}]

-----

The public parameters PK are:
{p: 2147483647, q: 2147483647, g: 2, G: 1, n: 2147483646, h: 1, K: 1,
  x: 1, y: 1, z: 1, w: 1, v: 1, u: 1, t: 1, s: 1, r: 1, r1: 1, r2: 1, r3: 1, r4: 1,
  r5: 1, r6: 1, r7: 1, r8: 1, r9: 1, r10: 1, r11: 1, r12: 1, r13: 1, r14: 1, r15: 1,
  r16: 1, r17: 1, r18: 1, r19: 1, r20: 1, r21: 1, r22: 1, r23: 1, r24: 1, r25: 1,
  r26: 1, r27: 1, r28: 1, r29: 1, r30: 1, r31: 1, r32: 1, r33: 1, r34: 1, r35: 1,
  r36: 1, r37: 1, r38: 1, r39: 1, r40: 1, r41: 1, r42: 1, r43: 1, r44: 1, r45: 1,
  r46: 1, r47: 1, r48: 1, r49: 1, r50: 1, r51: 1, r52: 1, r53: 1, r54: 1, r55: 1,
  r56: 1, r57: 1, r58: 1, r59: 1, r60: 1, r61: 1, r62: 1, r63: 1, r64: 1, r65: 1,
  r66: 1, r67: 1, r68: 1, r69: 1, r70: 1, r71: 1, r72: 1, r73: 1, r74: 1, r75: 1,
  r76: 1, r77: 1, r78: 1, r79: 1, r80: 1, r81: 1, r82: 1, r83: 1, r84: 1, r85: 1,
  r86: 1, r87: 1, r88: 1, r89: 1, r90: 1, r91: 1, r92: 1, r93: 1, r94: 1, r95: 1,
  r96: 1, r97: 1, r98: 1, r99: 1, r100: 1}

-----

The private key D is:
{d: 1, d1: 1, d2: 1, d3: 1, d4: 1, d5: 1, d6: 1, d7: 1, d8: 1, d9: 1, d10: 1,
  d11: 1, d12: 1, d13: 1, d14: 1, d15: 1, d16: 1, d17: 1, d18: 1, d19: 1, d20: 1,
  d21: 1, d22: 1, d23: 1, d24: 1, d25: 1, d26: 1, d27: 1, d28: 1, d29: 1, d30: 1,
  d31: 1, d32: 1, d33: 1, d34: 1, d35: 1, d36: 1, d37: 1, d38: 1, d39: 1, d40: 1,
  d41: 1, d42: 1, d43: 1, d44: 1, d45: 1, d46: 1, d47: 1, d48: 1, d49: 1, d50: 1,
  d51: 1, d52: 1, d53: 1, d54: 1, d55: 1, d56: 1, d57: 1, d58: 1, d59: 1, d60: 1,
  d61: 1, d62: 1, d63: 1, d64: 1, d65: 1, d66: 1, d67: 1, d68: 1, d69: 1, d70: 1,
  d71: 1, d72: 1, d73: 1, d74: 1, d75: 1, d76: 1, d77: 1, d78: 1, d79: 1, d80: 1,
  d81: 1, d82: 1, d83: 1, d84: 1, d85: 1, d86: 1, d87: 1, d88: 1, d89: 1, d90: 1,
  d91: 1, d92: 1, d93: 1, d94: 1, d95: 1, d96: 1, d97: 1, d98: 1, d99: 1, d100: 1}
-----
```

Figure 3: Authority values (Attribute list, PK, D)

These default values are based on the client privacy policy creation page, which is shown in Figure 4.

Full Name:	<input type="text"/>
Who can view your info?	<input checked="" type="radio"/> Advertisers <input type="radio"/> PrivateBook <input type="radio"/> Only Me
	<input checked="" type="radio"/> Targeted Ads <input type="radio"/> Site Statistics
	<input checked="" type="radio"/> Yes <input type="radio"/> No
	Can we determine your GPS location?

Figure 4: Client Policy creation page

4.2 Authority Key Storage and Distribution

The Authority then serializes the dictionary keys and stores them, shown in Figure 5.

App name:	privateBook
App list:	
P key:	
D key:	gAJ9cQAoVQNEanBxAX1xAihVA0

Figure 5: Authority key storage

These keys are then provided to the server and the clients by way of database

4.3 Client Privacy Policy Creation

The client, to become a privateBook user, registers themselves with a name and their privacy policy, as shown in Figure 4. They are then given their own ID, simply the order in which they were created, as their password, and username of their full name. These design choices are meant to simplify the account creation process: the focus of the prototype is not concerned with creating users as much as maintaining the privacy of their data, so privateBook has foregone security measures that should be in place in practical applications. Two policies are shown in Figure 6. They are strings of boolean expressions. The first is a policy that has the least privacy restrictions as given by the service, and the second has the strictest settings. The first includes an 'or' for every type of privacy setting which includes all possible attributes; this allows for any default setting in the attributes list to satisfy the policy. The second, stricter policy includes the fewest attributes, making it the most difficult to satisfy.

```
The user policy is:  
((ME) and (1) and (SITE) and (KEEP) and (NOBACK) and (NOLOCATE))  
The user policy is:  
((ME) and (1) and (SITE) and (KEEP) and (NOBACK) and (NOLOCATE))
```

Figure 6: Least strict and most strict user privacy policies

This method of creating a default attribute list was to allow the server to fail. In practice, a service that intends to have wide use should never fail, and thus an early idea was to simply let each attribute be the type of privacy setting it was, and store the values elsewhere. Then the server would still need to decrypt the data and thus inherently read the privacy policy, but that would offer no real difference from what exists currently, as the actual privacy settings would still be stored somewhere and potentially ignored, as they are now. So a difficulty arose in finding an appropriate method to allow one attribute list to successfully satisfy many different policies. This method was chosen because it is based on an idea of having different levels of privacy, where a baseline is chosen by the service and clients can be more or less strict; more, rendering the service useless for that user, and less, allowing for full use.

4.4 Client Encryption and privateBook Response

The client then makes a status update, much like a Facebook one, and submits that. The data is encrypted with the update, a string, along with the user's privacy policy and the public key PK. This results in a long cipher text, which is stored in privateBook's database as a serialized dictionary. The results of this process are shown in Figure 7. E is stored in the database, then the page is reloaded, displaying all of the user's updates. The server, upon loading of the user's homepage, first attempts to decrypt their data. If it is successful, the date, time and actual content are shown. If it is unsuccessful, the content is replaced with: "Your status could not be displayed: this service does not support your privacy policy!" Thus this service will not work for the client if their privacy settings are too strict, and therefore their policy is not satisfied by privateBook's default attribute list.

5 Issues with the Prototype

Though the `privateBook` prototype offers a solution, there are, as can be expected, some faults. The biggest issue is that the encrypted data, E , is

Acknowledgements

Professor Robert Muller and Stefan Saroiu at Microsoft Research

References

- [1] Lawler, Ryan. "QuizUp Sends Personal User Info To Strangers, Company Says Bug Contributed To Weakened Security." TechCrunch, 25 Nov. 2013.
- [2] Bajaj, Vikas. "Imagine if Companies Had to Ask Before Using Your Data." Taking Note. The New York Times, 13 Mar 2014.
- [3] J. Yang, K. Yessenov, A. Solar-Lezama. A Language for Automatically Enforcing Privacy Policies. POPL 2012.
- [4] Akinyele, Joseph A. and Garman, Christina and Miers, Ian and Pagano. "Charm: a framework for rapidly prototyping cryptosystems." Journal of Cryptographic Engineering 3.2 (2013): 111-128.
- [5] J. Bethencourt, A. Sahai, B. Waters. Ciphertext-Policy Attribute-Based Encryption.