CS 395T: Topics in Cryptography **Spring 2024**

Exercise Set 3

Due: February 28, 2024 at 11:59pm (Submit on Gradescope) **Instructor:** David Wu

Instructions. You **must** typeset your solution in LaTeX using the provided template:

`https://www.cs.utexas.edu/~dwu4/courses/sp24/static/homework.tex`

You must submit your problem set via Gradescope (accessible through Canvas).

Collaboration Policy. You may discuss your general approach with other students, but you may not share written documents. You should not search online for solutions to these problems. If you do consult external sources, you must cite them in your submission. You must include the names of all of your collaborators with your submission. Refer to the official course policies for the full details.

Problem 1: ABE from Public-Key Encryption [20 points]. We showed in class how to construct a key-policy ABE scheme for any policy family that can be described by a linear secret sharing scheme (which captures monotone Boolean formulas as a special case). The Goyal-Pandey-Sahai-Waters construction presented in class was (selectively) secure against an adversary that could request arbitrarily many secret keys of its choosing. We say that such schemes are "collusion-resistant:" given multiple keys that are not authorized to decrypt a particular ciphertext, the adversary still cannot break semantic security. In this problem, we will see that without collusion-resistance, ABE is much easier to construct.

Let $(\mathsf{KeyGen}', \mathsf{Encrypt}', \mathsf{Decrypt}')$ be a semantically-secure public-key encryption scheme. Using only this scheme, show how to construct a *ciphertext-policy* attribute-based encryption (ABE) scheme that supports any policy family that can be described by a linear secret sharing scheme, and where (selective) security holds against an adversary that makes most *one* key-generation query. In ciphertext-policy ABE, the encryption algorithm takes in a policy (i.e., the share-generation matrix for a linear secret sharing scheme) and the key-generation algorithm takes in a set of attributes S. A secret key for set S can be used to decrypt a ciphertext with policy P if $P(S) = 1$ (just like in key-policy ABE).

(a) Formally describe your ABE scheme $(\mathsf{Setup}, \mathsf{KeyGen}, \mathsf{Encrypt}, \mathsf{Decrypt})$.

(b) Prove the correctness of your scheme. Correctness should hold for all decryption policies and sets of attributes (that satisfy the decryption policy).

(c) Explain *informally* why your scheme is (selectively) secure against an adversary that can make at most one key-generation query. You may assume that the selective-security adversary commits to the key-generation query it makes at the beginning of the game (i.e., the set of attributes S it queries on). You do *not* need to give a formal reduction. Your argument should appeal to security of the linear secret sharing scheme (i.e., a set of unauthorized shares perfectly hides the secret) and semantic security of the public-key encryption scheme.

(d) What goes wrong if the adversary obtains more than one key in your scheme? Give a brief informal explanation.

Remark: Using a generalization of this construction, it is possible to construct a single-key-secure ABE scheme for arbitrary policies (not just those captured by linear secret sharing schemes). Using combinatoric techniques, it is also possible to support giving out a bounded number of secret keys in the security game. We say such ABE schemes are bounded-collusion resistant.

Optional Feedback. Please answer the following *optional* questions to help design future exercise sets. You do not need to answer these questions. However, we do encourage you to provide us feedback on how to improve the course experience.

(a) How long did you spend on this exercise set?

(b) Do you have any feedback for this exercise set?

(c) Do you have any feedback on the course so far?

(d) Are there specific topics that you are interested in seeing in this course?