

Predicate Encryption Supporting Disjunctions, Polynomial Equations, and Inner Products

Jonathan Katz
jkatz@cs.umd.edu

Amit Sahai
sahai@cs.ucla.edu

Brent Waters*
bwaters@cs1.sri.com

Abstract

Predicate encryption is a new paradigm generalizing, among other things, identity-based encryption. In a predicate encryption scheme, secret keys correspond to predicates and ciphertexts are associated with attributes; the secret key SK_f corresponding to a predicate f can be used to decrypt a ciphertext associated with attribute I if and only if $f(I) = 1$. Constructions of such schemes are currently known for relatively few classes of predicates.

We construct such a scheme for predicates corresponding to the evaluation of *inner products* over \mathbb{Z}_N (for some large integer N). This, in turn, enables constructions in which predicates correspond to the evaluation of disjunctions, polynomials, CNF/DNF formulae, or threshold predicates (among others). Besides serving as a significant step forward in the theory of predicate encryption, our results lead to a number of applications that are interesting in their own right.

1 Introduction

Traditional public-key encryption is rather coarse-grained: a sender encrypts a message M with respect to a given public key PK , and only the owner of the (unique) secret key associated with PK can decrypt the resulting ciphertext and recover the message. These straightforward semantics suffice for point-to-point communication, where encrypted data is intended for one particular user who is known to the sender in advance. In other settings, however, the sender may instead want to define some complex *policy* determining who is allowed to recover the encrypted data. For example, classified data might be associated with certain keywords; this data should be accessible to users who are allowed to read *all* classified information, as well as to users allowed to read information associated with the particular keywords in question. Or, in a health care application, a patient's records should perhaps be accessible only to a physician who has treated the patient in the past.

Applications such as those sketched above require new cryptographic mechanisms that provide more fine-grained control over access to encrypted data. *Predicate encryption* offers one such tool. At a high level (formal definitions are given in Section 2), secret keys in a predicate encryption scheme correspond to predicates in some class \mathcal{F} , and a sender associates a ciphertext with an attribute in a set Σ ; a ciphertext associated with the attribute $I \in \Sigma$ can be decrypted by a secret key SK_f corresponding to the predicate $f \in \mathcal{F}$ if and only if $f(I) = 1$.

*Supported by NSF CNS-0524252, CNS-0716199; the US Army Research Office under the CyberTA Grant No. W911NF-06-1-0316; and the U.S. Department of Homeland Security under Grant Award Number 2006-CS-001-000001. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security.

The “basic” level of security achieved by such schemes guarantees, informally, that a ciphertext associated with attribute I hides all information about the underlying message unless one is in the possession of a secret key giving the explicit ability to decrypt. I.e., if an adversary \mathcal{A} holds keys $SK_{f_1}, \dots, SK_{f_\ell}$, then \mathcal{A} learns nothing about the message if $f_1(I) = \dots = f_\ell(I) = 0$. We refer to this security notion as *payload hiding*. A stronger notion of security, that we call *attribute hiding*, requires further that a ciphertext hides all information about the associated attribute I except that which is explicitly leaked by the keys in one’s possession; i.e., an adversary holding secret keys as above learns only the values $f_1(I), \dots, f_\ell(I)$. Formal definitions are given in Section 2.

Much recent work aimed at constructing different types of fine-grained encryption schemes can be cast in the framework of predicate encryption. Identity-based encryption (IBE) [21, 9, 14, 4, 5, 23] can be viewed as predicate encryption for the class of equality tests; the standard notion of security for IBE [9, 13] corresponds to payload-hiding, while *anonymous* IBE [8, 12, 15] corresponds to the stronger notion of attribute hiding. Attribute-based encryption schemes [20, 16, 3, 19] can also be cast in the framework of predicate encryption, though in this case all the listed constructions achieve payload hiding only. Boneh and Waters [11] construct a predicate encryption scheme that handles conjunctions of, e.g., equality tests and range searches; their scheme satisfies the stronger notion of attribute hiding. Shi [22] et. al. looked at providing more efficient range searching, but under a weaker security model in which attribute hiding only holds if none of the attacker secret keys are associated with predicates that evaluate to 1 on the challenge ciphertext.

Other work introducing concepts related to the idea of predicate encryption includes [2, 1]. In contrast to the present work, however, the threat model in those works do not consider *collusion* among users holding different secret keys.

1.1 Our Results

An important research direction is to construct predicate encryption schemes for predicate classes \mathcal{F} that are as expressive as possible, with the ultimate goal being to handle all polynomial-time predicates. Most prior work, listed above, yields only payload-hiding schemes; existing techniques for obtaining attribute-hiding schemes seem limited to enforcing *conjunctions*. (Indeed, handling disjunctions was left as an open question in [11].) Getting slightly technical, this is because the underlying cryptographic mechanism used in the above schemes is to pair components of the secret key with corresponding components of the ciphertext and then multiply the intermediate results together; a “cancellation” occurs if everything “matches”, but a random group element results if there is any “mismatch”. Thus, the holder of a non-matching secret key learns only that there was a mismatch in *at least one* position, but does not learn the number of mismatches or their locations. Very different cryptographic techniques seem needed to support disjunctions, since a mismatch in a single position cannot result in a completely random group element but must still somehow result in a “cancellation” if there is a match in any other position. (We stress that what makes this difficult is that we must hide the position of a match and only reveal that there was a match in at least one position.)

The aim of our work is to construct attribute-hiding schemes handling disjunctions. As a stepping stone toward this goal, we first focus on predicates corresponding to the computation of inner products over \mathbb{Z}_N (for some large integer N). Formally, we take $\Sigma = \mathbb{Z}_N^n$ as our set of attributes, and take our class of predicates to be $\mathcal{F} = \{f_{\vec{x}} \mid \vec{x} \in \mathbb{Z}_N^n\}$ where $f_{\vec{x}}(\vec{y}) = 1$ iff $\langle \vec{x}, \vec{y} \rangle = 0$. (Here, $\langle \vec{x}, \vec{y} \rangle$ denotes the standard inner product $\sum_{i=1}^n x_i \cdot y_i \bmod N$ of two vectors \vec{x} and \vec{y} .) We construct a predicate encryption scheme for this \mathcal{F} without random oracles, based on two new

assumptions in composite-order groups equipped with a bilinear map. Our assumptions are non-interactive and of fixed size (i.e., not “ q -type”), and can be shown to hold in the generic group model. A pessimistic interpretation of our results would be that we prove security in the generic group model, but we believe it is of importance that we are able to distill our necessary assumptions to ones that are compact and falsifiable. Our construction uses new techniques, including the fact that we work in a bilinear group whose order is a product of *three* primes.

We view our main construction as a significant step toward increasing the expressiveness of predicate encryption in general. Moreover, we show that any predicate encryption scheme supporting “inner product” predicates as described above can be used as a building block to construct predicates of more general types:

- As an easy warm-up, we show that it implies (anonymous) identity-based encryption as well as hidden-vector encryption [11]. As a consequence, our work implies all the results of [11].
- We can also construct predicate encryption schemes supporting polynomial evaluation. Here, we take \mathbb{Z}_N as our set of attributes, and predicates correspond to polynomials over \mathbb{Z}_N of some bounded degree; a predicate evaluates to 1 iff the corresponding polynomial evaluates to 0 on the attribute in question. We can also extend this to include multi-variate polynomials (in some bounded number of variables). A “dual” of this construction allows the attributes to be polynomials, and the predicates to correspond to evaluation at a fixed point.
- Given the above, we can fairly easily support predicates that are *disjunctions* of other predicates (e.g., equality), thus achieving our main goal. In the context of identity-based encryption, this gives the ability to issue secret keys corresponding to a *set* of identities that enables decryption whenever a ciphertext is encrypted to any identity in this set (without leaking which identity was actually used to encrypt).
- We also show how to handle predicates corresponding to DNF and CNF formulas of some bounded size.
- Working directly with our “inner product” construction, we can derive a scheme supporting threshold queries of the following form: Attributes are subsets of $A = \{1, \dots, \ell\}$, and predicates take the form $\{f_{S,t} \mid S \subseteq A\}$ where $f_{S,t}(S') = 1$ iff $S \cap S' = t$. This is useful in the “fuzzy IBE” setting of Sahai and Waters [20], and improves on their work in that we achieve attribute hiding (rather than only payload hiding) and handle *exact* thresholds.

We defer further discussion regarding the above until Section 5.

2 Definitions

We define the syntax of predicate encryption and the security properties discussed informally in the Introduction. (Our definitions follow the general framework of those given in [11].) Throughout this section, we consider the general case where Σ denotes an arbitrary set of attributes and \mathcal{F} denotes an arbitrary set of predicates over Σ . Formally, both Σ and \mathcal{F} may depend on the security parameter and/or the master public parameters; for simplicity, we leave this implicit.

Definition 2.1. *A predicate encryption scheme for the class of predicates \mathcal{F} over the set of attributes Σ consists of four PPT algorithms Setup, GenKey, Enc, Dec such that:*

- **Setup** takes as input the security parameter 1^n and outputs a (master) public key PK and a (master) secret key SK .

- **GenKey** takes as input the master secret key SK and a (description of a) predicate $f \in \mathcal{F}$. It outputs a key SK_f .
- **Enc** takes as input the public key PK , an attribute $I \in \Sigma$, and a message M in some associated message space. It returns a ciphertext C . We write this as $C \leftarrow \text{Enc}_{PK}(I, M)$.
- **Dec** takes as input a secret key SK_f and a ciphertext C . It outputs either a message M or the distinguished symbol \perp .

For correctness, we require that for all n , all (PK, SK) generated by $\text{Setup}(1^n)$, all $f \in \mathcal{F}$, any key $SK_f \leftarrow \text{GenKey}_{SK}(f)$, and all $I \in \Sigma$:

- If $f(I) = 1$ then $\text{Dec}_{SK_f}(\text{Enc}_{PK}(I, M)) = M$.
- If $f(I) = 0$ then $\text{Dec}_{SK_f}(\text{Enc}_{PK}(I, M)) = \perp$ with all but negligible probability.

We will also consider a variant of the above that we call a *predicate-only scheme*. Here, **Enc** takes only an attribute I (and no message); the correctness requirement is that $\text{Dec}_{SK_f}(\text{Enc}_{PK}(I)) = f(I)$ and so all the receiver learns is whether the predicate is satisfied. A predicate-only scheme can serve as a useful building block toward a full-fledged predicate encryption scheme.

Our definition of attribute-hiding security corresponds to the notion described informally earlier. Here, an adversary may request keys corresponding to the predicates f_1, \dots, f_ℓ and is then given either $\text{Enc}_{PK}(I_0, M_0)$ or $\text{Enc}_{PK}(I_1, M_1)$ for attributes I_0, I_1 such that $f_i(I_0) = f_i(I_1)$ for all i . Furthermore, if $M_0 \neq M_1$ then it is required that $f_i(I_0) = f_i(I_1) = 0$ for all i . The goal of the adversary is to determine which attribute/message pair was encrypted, and the stated conditions ensure that this is not trivial. Our definition uses the “selective” notion of security introduced in [13]. Observe that when specialized to the case when \mathcal{F} consists of equality tests on strings, this notion corresponds to *anonymous* identity-based encryption (with selective-ID security).

Definition 2.2. A predicate encryption scheme with respect to \mathcal{F} and Σ is **attribute hiding** (or **simple secure**) if for all PPT adversaries \mathcal{A} , the advantage of \mathcal{A} in the following experiment is negligible in the security parameter n :

1. $\mathcal{A}(1^n)$ outputs $I_0, I_1 \in \Sigma$.
2. $\text{Setup}(1^n)$ is run to generate PK, SK , and the adversary is given PK .
3. \mathcal{A} may adaptively request keys for any predicates $f_1, \dots, f_\ell \in \mathcal{F}$ subject to the restriction that $f_i(I_0) = f_i(I_1)$ for all i . In response, \mathcal{A} is given the corresponding keys $SK_{f_i} \leftarrow \text{GenKey}_{SK}(f_i)$.
4. \mathcal{A} outputs two equal-length messages M_0, M_1 . If there is an i for which $f_i(I_0) = f_i(I_1) = 1$, then it is required that $M_0 = M_1$. A random bit ν is chosen, and \mathcal{A} is given the ciphertext $C \leftarrow \text{Enc}_{PK}(I_b, M_b)$.
5. The adversary may continue to request keys for additional predicates, subject to the same restrictions as before.
6. \mathcal{A} outputs a bit ν' , and succeeds if $\nu' = \nu$.

The advantage of \mathcal{A} is the absolute value of the difference between its success probability and $1/2$.

Payload hiding, a strictly weaker notion, is defined by forcing $I_0 = I_1 = I$ in the above (in which case \mathcal{A} has no possible advantage if it ever holds that $f_i(I) = 1$). For predicate-only encryption schemes we simply omit the messages in the above experiment. For convenience, we include in Appendix A a re-statement of the definition of security given above for the particular inner-product predicate we use in our main construction.

3 Background on Pairings and Complexity Assumptions

We review some general notions about bilinear groups of *composite order*, first used in cryptographic applications by [10]. In contrast to all prior work using composite-order bilinear groups, however, we use groups whose order N is a product of *three* (distinct) primes. This is for simplicity only, since a variant of our construction can be proven secure based on a “decisional linear”-type assumption [7] in a group of composite order N which is a product of two primes.¹

Let \mathcal{G} be an algorithm that takes as input a security parameter 1^n and outputs a tuple $(p, q, r, \mathbb{G}, \mathbb{G}_T, \hat{e})$ where p, q, r are distinct primes, \mathbb{G} and \mathbb{G}_T are two cyclic groups of order $N = pqr$, and $\hat{e} : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ is *bilinear* (i.e., $\forall u, v \in \mathbb{G}$ and $\forall a, b \in \mathbb{Z}$ we have $\hat{e}(u^a, v^b) = \hat{e}(u, v)^{ab}$) and *non-degenerate* (i.e., if g generates \mathbb{G} then $\hat{e}(g, g)$ generates \mathbb{G}_T). We assume that the group action in \mathbb{G} and \mathbb{G}_T as well as the bilinear map \hat{e} are all computable in time polynomial in n . Furthermore, we assume that the description of \mathbb{G} and \mathbb{G}_T includes generators of \mathbb{G} and \mathbb{G}_T respectively.

We use the notation $\mathbb{G}_p, \mathbb{G}_q, \mathbb{G}_r$ to denote the subgroups of \mathbb{G} having order p, q , and r , respectively. Observe that $\mathbb{G} = \mathbb{G}_p \times \mathbb{G}_q \times \mathbb{G}_r$. Note also that if g is a generator of \mathbb{G} , then the element g^{pq} is a generator of \mathbb{G}_r ; the element g^{pr} is a generator of \mathbb{G}_q ; and the element g^{qr} is a generator of \mathbb{G}_p . Furthermore, if, e.g., $h_p \in \mathbb{G}_p$ and $h_q \in \mathbb{G}_q$ then

$$\hat{e}(h_p, h_q) = \hat{e}((g^{qr})^{\alpha_1}, (g^{pr})^{\alpha_2}) = \hat{e}(g^{\alpha_1}, g^{r\alpha_2})^{pqr} = 1,$$

where $\alpha_1 = \log_{g^{qr}} h_p$ and $\alpha_2 = \log_{g^{pr}} h_q$. Similar rules hold whenever \hat{e} is applied to non-identity elements in distinct subgroups.

3.1 Our Assumptions

We now state the assumptions we use to prove security of our construction. As remarked earlier, these assumptions are new but we justify them in Appendix B by proving that they hold in the generic group model under the assumption that finding a non-trivial factor of N (the group order) is hard. At a minimum, then, our construction can be viewed as secure in the generic group model. Nevertheless, we state our assumptions explicitly and highlight that they are non-interactive and of fixed size.

Assumption 1. Let \mathcal{G} be as above. We say that \mathcal{G} satisfies Assumption 1 if the advantage of any PPT algorithm \mathcal{A} in the following experiment is negligible in the security parameter n :

1. $\mathcal{G}(1^n)$ is run to obtain $(p, q, r, \mathbb{G}, \mathbb{G}_T, \hat{e})$. Set $N = pqr$, and let g_p, g_q, g_r be generators of $\mathbb{G}_p, \mathbb{G}_q$, and \mathbb{G}_r , respectively.
2. Choose random $Q_1, Q_2, Q_3 \in \mathbb{G}_q$, random $R_1, R_2, R_3 \in \mathbb{G}_r$, random $a, b, s \in \mathbb{Z}_p$, and a random bit ν . Give to \mathcal{A} the values $(N, \mathbb{G}, \mathbb{G}_T, \hat{e})$ as well as

$$g_p, g_r, g_q R_1, g_p^b, g_p^{b^2}, g_p^a g_q, g_p^{ab} Q_1, g_p^s, g_p^{bs} Q_2 R_2.$$

If $\nu = 0$ give \mathcal{A} the value $T = g_p^{b^2 s} R_3$, while if $\nu = 1$ give \mathcal{A} the value $T = g_p^{b^2 s} Q_3 R_3$.

3. \mathcal{A} outputs a bit ν' , and succeeds if $\nu' = \nu$.

¹This is analogous to the “folklore” transformation (based on the decisional linear assumption) that converts any scheme using groups whose order N is a product of two primes, to a scheme that uses prime-order groups. (See, for example [17]). Typically, using composite order groups gives a simpler scheme; however, since the group sizes are larger group operations are less efficient.

The advantage of \mathcal{A} is the absolute value of the difference between its success probability and $1/2$.

Assumption 2. Let \mathcal{G} be as above. We say that \mathcal{G} satisfies Assumption 2 if the advantage of any PPT algorithm \mathcal{A} in the following experiment is negligible in the security parameter n :

1. $\mathcal{G}(1^n)$ is run to obtain $(p, q, r, \mathbb{G}, \mathbb{G}_T, \hat{e})$. Set $N = pqr$, and let g_p, g_q, g_r be generators of $\mathbb{G}_p, \mathbb{G}_q$, and \mathbb{G}_r , respectively.
2. Choose random $h \in \mathbb{G}_p$ and $Q_1, Q_2 \in \mathbb{G}_q$, random $s, \gamma \in \mathbb{Z}_q$, and a random bit ν . Give to \mathcal{A} the values $(N, \mathbb{G}, \mathbb{G}_T, \hat{e})$ as well as

$$g_p, g_q, g_r, h, g_p^s, h^s Q_1, g_p^\gamma Q_2, \hat{e}(g_p, h)^\gamma.$$

If $\nu = 0$ then give \mathcal{A} the value $\hat{e}(g_p, h)^{\gamma s}$, while if $\nu = 1$ then give \mathcal{A} a random element of \mathbb{G}_T .

3. \mathcal{A} outputs a bit ν' , and succeeds if $\nu' = \nu$.

The advantage of \mathcal{A} is the absolute value of the difference between its success probability and $1/2$.

Note that both the above assumptions imply the hardness of factoring N .

4 Our Main Construction

Our main construction is a predicate-only scheme where the set of attributes is $\Sigma = \mathbb{Z}_N^n$, and the class of predicates is $\mathcal{F} = \{f_{\vec{x}} \mid \vec{x} \in \mathbb{Z}_N^n\}$ with $f_{\vec{x}}(\vec{y}) = 1$ iff $\langle \vec{x}, \vec{y} \rangle = 0 \pmod N$. In this section we provide our predicate-only construction and give some intuition about our proof. For space considerations the details of the proof are presented in Appendix C. In Appendix D we show how our scheme can be extended to give a full-fledged predicate encryption scheme.

Intuition. In our construction, each ciphertext has associated with it a (secret) vector \vec{x} , and each secret key corresponds to a vector \vec{v} . The decryption procedure must check whether $\vec{x} \cdot \vec{v} = 0$, and reveal nothing about \vec{x} but whether this is true. To do this, we will make use of a bilinear group \mathbb{G} whose order N is the product of three primes p, q , and r . Let $\mathbb{G}_p, \mathbb{G}_q$, and \mathbb{G}_r denote the subgroups of \mathbb{G} having order p, q , and r , respectively. We will (informally) assume, as in [10], that a random element in any of these subgroups is indistinguishable from a random element of \mathbb{G} .² Thus, we can use random elements from one subgroup to “mask” elements from another subgroup.

At a high level, we will use these subgroups as follows. \mathbb{G}_q will be used to encode the vectors \vec{x} and \vec{v} in the ciphertext and secret keys, respectively. Computation of the inner product $\langle \vec{v}, \vec{x} \rangle$ will be done in \mathbb{G}_q (in the exponent), using the bilinear map. \mathbb{G}_p will be used to encode an equation (again in the exponent) that evaluates to zero when decryption is done properly. This subgroup is used to prevent an adversary from improperly “manipulating” the computation (by, e.g., changing the ordering of components of the ciphertext or secret key, raising these components to some power, etc.). On an intuitive level, if the adversary tries to manipulate the computation in any way, then the computation occurring in the \mathbb{G}_p subgroup will no longer yield the identity (i.e., will no longer yield 0 in the exponent), but will instead have the effect of “masking” the correct answer with a random element of \mathbb{G}_p (which will invalidate the entire computation). Elements in \mathbb{G}_r are used for “general masking” of terms in other subgroups; i.e., random elements of \mathbb{G}_r will be multiplied with various components of the ciphertext (and secret key) in order to “hide” information that might be present in the \mathbb{G}_p and \mathbb{G}_q subgroups.

We now proceed to the formal description of our scheme.

²This is only for intuition. Our actual computational assumption is given in Section 3.

4.1 The Construction

Setup(1^n) The setup algorithm first runs $\mathcal{G}(1^n)$ to obtain $(p, q, r, \mathbb{G}, \mathbb{G}_T, \hat{e})$ with $\mathbb{G} = \mathbb{G}_p \times \mathbb{G}_q \times \mathbb{G}_r$. Next, it computes g_p, g_q , and g_r as generators of $\mathbb{G}_p, \mathbb{G}_q$, and \mathbb{G}_r , respectively. It then chooses $R_{1,i}, R_{2,i} \in \mathbb{G}_r$ and $h_{1,i}, h_{2,i} \in \mathbb{G}_p$ uniformly at random for $i = 1$ to n , and $R_0 \in \mathbb{G}_r$ uniformly at random. The public parameters include $(N = pqr, \mathbb{G}, \mathbb{G}_T, \hat{e})$ along with:

$$PK = (g_p, g_r, Q = g_q \cdot R_0, \{H_{1,i} = h_{1,i} \cdot R_{1,i}, H_{2,i} = h_{2,i} \cdot R_{2,i}\}_{i=1}^n).$$

The master secret key SK is $(p, q, r, g_q, \{h_{1,i}, h_{2,i}\}_{i=1}^n)$.

Enc $_{PK}(\vec{x})$ Let $\vec{x} = (x_1, \dots, x_n)$ with $x_i \in \mathbb{Z}_N$. This algorithm chooses random $s, \alpha, \beta \in \mathbb{Z}_N$ and $R_{3,i}, R_{4,i} \in \mathbb{G}_r$ for $i = 1$ to n . (Note: a random element $R \in \mathbb{G}_r$ can be sampled by choosing random $\delta \in \mathbb{Z}_N$ and setting $R = g_r^\delta$.) It outputs the ciphertext

$$C = \left(C_0 = g_p^s, \left\{ C_{1,i} = H_{1,i}^s \cdot Q^{\alpha \cdot x_i} \cdot R_{3,i}, C_{2,i} = H_{2,i}^s \cdot Q^{\beta \cdot x_i} \cdot R_{4,i} \right\}_{i=1}^n \right).$$

GenKey $_{SK}(\vec{v})$ Let $\vec{v} = (v_1, \dots, v_n)$, and recall $SK = (p, q, r, g_q, \{h_{1,i}, h_{2,i}\}_{i=1}^n)$. This algorithm chooses random $r_{1,i}, r_{2,i} \in \mathbb{Z}_p$ for $i = 1$ to n , random $R_5 \in \mathbb{G}_r$, random $f_1, f_2 \in \mathbb{Z}_q$, and random $Q_6 \in \mathbb{G}_q$. It then outputs

$$SK_{\vec{v}} = \left(K = R_5 \cdot Q_6 \cdot \prod_{i=1}^n h_{1,i}^{-r_{1,i}} \cdot h_{2,i}^{-r_{2,i}}, \left\{ K_{1,i} = g_p^{r_{1,i}} \cdot g_q^{f_1 \cdot v_i}, K_{2,i} = g_p^{r_{2,i}} \cdot g_q^{f_2 \cdot v_i} \right\}_{i=1}^n \right).$$

Dec $_{SK_{\vec{v}}}(C)$ Let $C = (C_0, \{C_{1,i}, C_{2,i}\}_{i=1}^n)$ and $SK_{\vec{v}} = (K, \{K_{1,i}, K_{2,i}\}_{i=1}^n)$ be as above. The decryption algorithm outputs 1 iff

$$\hat{e}(C_0, K) \cdot \prod_{i=1}^n \hat{e}(C_{1,i}, K_{1,i}) \cdot \hat{e}(C_{2,i}, K_{2,i}) \stackrel{?}{=} 1.$$

Correctness. To see that correctness holds, let C and $SK_{\vec{v}}$ be as above. Then

$$\begin{aligned} & \hat{e}(C_0, K) \cdot \prod_{i=1}^n \hat{e}(C_{1,i}, K_{1,i}) \cdot \hat{e}(C_{2,i}, K_{2,i}) \\ &= \hat{e} \left(g_p^s, R_5 Q_6 \prod_{i=1}^n h_{1,i}^{-r_{1,i}} h_{2,i}^{-r_{2,i}} \right) \\ & \quad \cdot \prod_{i=1}^n \hat{e} \left(H_{1,i}^s Q^{\alpha \cdot x_i} R_{3,i}, g_p^{r_{1,i}} g_q^{f_1 \cdot v_i} \right) \cdot \hat{e} \left(H_{2,i}^s Q^{\beta \cdot x_i} R_{4,i}, g_p^{r_{2,i}} g_q^{f_2 \cdot v_i} \right) \\ &= \hat{e} \left(g_p^s, \prod_{i=1}^n h_{1,i}^{-r_{1,i}} h_{2,i}^{-r_{2,i}} \right) \cdot \prod_{i=1}^n \hat{e} \left(h_{1,i}^s \cdot g_q^{\alpha \cdot x_i}, g_p^{r_{1,i}} g_q^{f_1 \cdot v_i} \right) \cdot \hat{e} \left(h_{2,i}^s \cdot g_q^{\beta \cdot x_i}, g_p^{r_{2,i}} g_q^{f_2 \cdot v_i} \right) \\ &= \prod_{i=1}^n \hat{e}(g_q, g_q)^{(\alpha f_1 + \beta f_2) x_i v_i} = \hat{e}(g_q, g_q)^{(\alpha f_1 + \beta f_2 \bmod q) \langle \vec{x}, \vec{v} \rangle}, \end{aligned}$$

where α, β are random in \mathbb{Z}_N and f_1, f_2 are random in \mathbb{Z}_q . If $\langle \vec{x}, \vec{v} \rangle = 0 \pmod N$, then the above evaluates to 1. If $\langle \vec{x}, \vec{v} \rangle \neq 0 \pmod N$ there are two cases: if $\langle \vec{x}, \vec{v} \rangle \neq 0 \pmod q$ then with all but negligible probability (over choice of α, β, f_1, f_2) the above evaluates to an element other than the identity. The other possibility is that $\langle \vec{x}, \vec{v} \rangle = 0 \pmod q$, in which case the above would always evaluate to 1; however, this would reveal a non-trivial factor of N and so this occurs with only negligible probability (recall, our assumptions imply hardness of finding a non-trivial factor of N).

There may appear to be some redundancy in our construction. For instance, the $C_{1,i}$ and $C_{2,i}$ components play almost identical roles. In fact we can view the encryption system as two parallel sub-systems linked via the C_0 component (and the corresponding component of the secret key). This two sub-system approach was first used by Boyen and Waters [12]; it can be viewed as a complex generalization of the Naor-Yung [18] “two-key” paradigm to the predicate encryption setting. A natural question is whether this redundancy can be eliminated to achieve better performance. While such a construction appears to be secure, our current proof (that utilizes a *non-interactive* assumption) relies in an essential way on having two parallel subsystems.

4.2 Proof Intuition

The most challenging aspect to providing a proof of our scheme naturally arises from the disjunctive capabilities of our system. In previous conjunctive systems (such as the one by Boneh and Waters [11]) the authors proved security by moving through a sequence of hybrid games, in which an encryption of a vector \vec{x} was changed component-by-component to the encryption of a vector \vec{y} . The adversary could only ask for queries that did not match either \vec{x} or \vec{y} , or queries that did not “look at” the components in which \vec{x} and \vec{y} differed. Thus, it was relatively straightforward to perform hybrid experiments over the components of \vec{x} and \vec{y} that differed, since the private keys given to the adversary did not “look at” these components.

In our proof an adversary will again try to determine whether a given ciphertext was encrypted with respect to \vec{x} or \vec{y} . However, in our case the adversary can legally request a secret key for a vector \vec{v} such that $\langle \vec{x}, \vec{v} \rangle = 0$ and $\langle \vec{y}, \vec{v} \rangle = 0$; i.e., it may obtain a key that should enable correct decryption in either case. This means that we cannot use the same proof strategy as in previous, conjunctive schemes. For example, if we change just one component at a time, then the “hybrid” vector used in an intermediate step will likely not be orthogonal to \vec{v} (and the adversary will be able to detect this because its secret key will no longer decrypt the given ciphertext). Therefore, we need to use a sequence of hybrid games in which entire vectors are changed in one step, instead of using a sequence of hybrid games where the vector is changed component-by-component.

To do this we take advantage of the fact that, as noted earlier, our encryption scheme contains two parallel sub-systems. In our proof we will use hybrid games where a challenge ciphertext will be encrypted with respect to one vector in the first sub-system and a *different* vector in the second sub-system. (Note that such a ciphertext is ill-formed, since any valid ciphertext will always use the same vector in each sub-system.) Let (\vec{a}, \vec{b}) denote a ciphertext encrypted using vector \vec{a} in the first sub-system and \vec{b} in the second sub-system. To prove indistinguishability when encrypting to \vec{x} (which corresponds to (\vec{x}, \vec{x})) and when encrypting to \vec{y} (which corresponds to (\vec{y}, \vec{y})), we will prove indistinguishability of the following sequence of hybrid games:

$$(\vec{x}, \vec{x}) \approx (\vec{x}, \vec{0}) \approx (\vec{x}, \vec{y}) \approx (\vec{0}, \vec{y}) \approx (\vec{y}, \vec{y}).$$

Using this structure in our proof allows us to use a simulator that will essentially be able to work in one sub-system without “knowing” what is happening in the other one. The simulator embeds

a “subgroup decision-like” assumption into the challenge ciphertext for each experiment. The structure of the challenge will determine whether a sub-system encrypts the given vector or the zero vector. Details of our proof and further discussion are given in Appendix C.

5 Applications of Our Main Construction

In this section we discuss some applications of predicate encryption schemes of the type constructed in this paper. Our treatment here is general and can be based on any predicate encryption scheme supporting “inner product” queries; we do not rely on any specific details of our construction.

Given a vector $\vec{x} \in \mathbb{Z}_N^\ell$, we denote by $f_{\vec{x}} : \mathbb{Z}_N^\ell \rightarrow \{0, 1\}$ the function such that $f_{\vec{x}}(\vec{y}) = 1$ iff $\langle \vec{x}, \vec{y} \rangle = 0$. We define $\mathcal{F}_\ell \stackrel{\text{def}}{=} \{f_{\vec{x}} \mid \vec{x} \in \mathbb{Z}_N^\ell\}$. An *inner product encryption scheme of dimension ℓ* is an attribute-hiding predicate encryption scheme for the class of predicates \mathcal{F}_ℓ .

5.1 Anonymous Identity-Based Encryption

As a warm-up, we show how anonymous identity-based encryption (IBE) can be recovered from any inner product encryption scheme of dimension 2. To generate the master public and secret keys for the IBE scheme, simply run the setup algorithm of the underlying inner product encryption scheme. To generate secret keys for the identity $I \in \mathbb{Z}_N$, set $\vec{I} := (1, I)$ and output the secret key for the predicate $f_{\vec{I}}$. To encrypt a message M for the identity $J \in \mathbb{Z}_N$, set $\vec{J} := (-J, 1)$ and encrypt the message using the encryption algorithm of the underlying inner product encryption scheme and the attribute \vec{J} . Since $\langle \vec{I}, \vec{J} \rangle = 0$ iff $I = J$, correctness and security follow.

5.2 Hidden-Vector Encryption

Given a set Σ , let $\Sigma_\star = \Sigma \cup \{\star\}$. Hidden-vector encryption (HVE) [11] corresponds to a predicate encryption scheme for the class of predicates $\Phi_\ell^{\text{hve}} = \{\phi_{(a_1, \dots, a_\ell)}^{\text{hve}} \mid a_1, \dots, a_\ell \in \Sigma_\star\}$, where

$$\phi_{(a_1, \dots, a_\ell)}^{\text{hve}}(x_1, \dots, x_\ell) = \begin{cases} 1 & \text{if, for all } i, \text{ either } a_i = x_i \text{ or } a_i = \star \\ 0 & \text{otherwise} \end{cases}.$$

A generalization of the ideas from the previous section can be used to realize hidden-vector encryption with $\Sigma = \mathbb{Z}_N$ from any inner product encryption scheme (Setup, GenKey, Enc, Dec) of dimension 2ℓ :

- The setup algorithm is unchanged.
- To generate a secret key corresponding to the predicate $\phi_{(a_1, \dots, a_\ell)}^{\text{hve}}$, first construct a vector $\vec{A} = (A_1, \dots, A_{2\ell})$ as follows:

$$\begin{aligned} \text{if } a_i \neq \star: & \quad A_{2i-1} := 1, \quad A_{2i} := a_i \\ \text{if } a_i = \star: & \quad A_{2i-1} := 0, \quad A_{2i} := 0. \end{aligned}$$

Then output the key obtained by running $\text{GenKey}_{SK}(f_{\vec{A}})$.

- To encrypt a message M for the attribute $x = (x_1, \dots, x_\ell)$, choose random $r_1, \dots, r_\ell \in \mathbb{Z}_N$ and construct a vector $\vec{X}_r = (X_1, \dots, X_{2\ell})$ as follows:

$$X_{2i-1} := -r_i \cdot x_i, \quad X_{2i} := r_i$$

(multiplication is done modulo N). Then output the ciphertext $C \leftarrow \text{Enc}_{PK}(\vec{X}_{\vec{r}}, M)$.

To see that correctness holds, let (a_1, \dots, a_ℓ) , \vec{A} , (x_1, \dots, x_ℓ) , \vec{r} , and $\vec{X}_{\vec{r}}$ be as above. Then:

$$\phi_{(a_1, \dots, a_\ell)}^{\text{hve}}(x_1, \dots, x_\ell) = 1 \Rightarrow \forall \vec{r}: \langle \vec{A}, \vec{X}_{\vec{r}} \rangle = 0 \Rightarrow \forall \vec{r}: f_{\vec{A}}(\vec{X}_{\vec{r}}) = 1.$$

Furthermore, assuming $\gcd(a_i - x_i, N) = 1$ for all i :

$$\phi_{(a_1, \dots, a_\ell)}^{\text{hve}}(x_1, \dots, x_\ell) = 0 \Rightarrow \Pr_{\vec{r}} \left[\langle \vec{A}, \vec{X}_{\vec{r}} \rangle = 0 \right] = 1/N \Rightarrow \Pr_{\vec{r}} \left[f_{\vec{A}}(\vec{X}_{\vec{r}}) = 1 \right] = 1/N,$$

which is negligible. Using this, one can prove security of the construction as well.

A straightforward modification of the above gives a scheme that is the “dual” of HVE, where the set of attributes is $(\Sigma_\star)^\ell$ and the class of predicates is $\bar{\Phi}_\ell^{\text{hve}} = \{\bar{\phi}_{(a_1, \dots, a_\ell)}^{\text{hve}} \mid a_1, \dots, a_\ell \in \Sigma\}$ with

$$\bar{\phi}_{(a_1, \dots, a_\ell)}^{\text{hve}}(x_1, \dots, x_\ell) = \begin{cases} 1 & \text{if, for all } i, \text{ either } a_i = x_i \text{ or } x_i = \star \\ 0 & \text{otherwise} \end{cases}.$$

5.3 Predicate Encryption Schemes Supporting Polynomial Evaluation

We can also construct predicate encryption schemes for classes of predicates corresponding to polynomial evaluation. Let $\Phi_{\leq d}^{\text{poly}} = \{f_p \mid p \in \mathbb{Z}_N[x], \deg(p) \leq d\}$, where

$$\phi_p(x) = \begin{cases} 1 & \text{if } p(x) = 0 \\ 0 & \text{otherwise} \end{cases}$$

for $x \in \mathbb{Z}_N$.³ Given an inner product encryption scheme (**Setup**, **GenKey**, **Enc**, **Dec**) of dimension $d+1$, we can construct a predicate encryption scheme for $\Phi_{\leq d}^{\text{poly}}$ as follows:

- The setup algorithm is unchanged.
- To generate a secret key corresponding to the polynomial $p = a_d x^d + \dots + a_0 x^0$, set $\vec{p} := (a_d, \dots, a_0)$ and output the key obtained by running $\text{GenKey}_{SK}(f_{\vec{p}})$.
- To encrypt a message M for the attribute $w \in \mathbb{Z}_N$, set $\vec{w} := (w^d \bmod N, \dots, w^0 \bmod N)$ and output the ciphertext $C \leftarrow \text{Enc}_{PK}(\vec{w}, M)$.

Since $p(w) = 0$ iff $\langle \vec{p}, \vec{w} \rangle = 0$, correctness and security follow.

The above shows that we can construct predicate encryption schemes where predicates correspond to univariate polynomials whose degree d is polynomial in the security parameter. This can be generalized to the case of polynomials in t variables, and degree at most d in each variable, as long as d^t is polynomial in the security parameter.

We can also construct schemes that are the “dual” of the above, in which attributes correspond to polynomials and predicates involve the evaluation of the input polynomial at some fixed point.

³Since this is a general treatment of inner products over \mathbb{Z}_N , we use p to denote a polynomial and not a factor of N in this context.

5.4 Disjunctions, Conjunctions, and Evaluating CNF and DNF Formulas

Given the polynomial-based constructions of the previous section, we can fairly easily build predicate encryption schemes for disjunctions of equality tests. For example, the predicate OR_{I_1, I_2} , where $\text{OR}_{I_1, I_2}(x) = 1$ iff either $x = I_1$ or $x = I_2$, can be encoded as the univariate polynomial

$$p(x) = (x - I_1) \cdot (x - I_2),$$

which evaluates to 0 iff the relevant predicate evaluates to 1. Similarly, the predicate $\overline{\text{OR}}_{I_1, I_2}$, where $\overline{\text{OR}}_{I_1, I_2}(x_1, x_2) = 1$ iff either $x_1 = I_1$ or $x_2 = I_2$, can be encoded as the bivariate polynomial

$$p'(x_1, x_2) = (x_1 - I_1) \cdot (x_2 - I_2).$$

Conjunctions can be handled in a similar fashion. Consider, for example, the predicate AND_{I_1, I_2} where $\text{AND}_{I_1, I_2}(x_1, x_2) = 1$ if both $x_1 = I_1$ and $x_2 = I_2$. Here, we determine the relevant secret key by choosing a random $r \in \mathbb{Z}_N$ and letting the secret key correspond to the polynomial

$$p''(x_1, x_2) = r \cdot (x_1 - I_1) + (x_2 - I_2).$$

Note that if $\text{AND}_{I_1, I_2}(x_1, x_2) = 1$ then $p''(x_1, x_2) = 0$, whereas if $\text{AND}_{I_1, I_2}(x_1, x_2) = 0$ then, with all but negligible probability over choice of r , it will hold⁴ that $p''(x_1, x_2) \neq 0$.

The above ideas extend to more complex combinations of disjunctions and conjunctions, and for boolean variables this means we can handle arbitrary CNF or DNF formulas. (For non-boolean variables we do not know how to directly handle negation.) As pointed out in the previous section, the complexity of the resulting scheme depends polynomially on d^t , where t is the number of variables and d is the maximum degree (of the resulting polynomial) in each variable.

5.5 Exact Thresholds

We conclude with an application that relies directly on inner product encryption. Here, we consider the setting of “fuzzy IBE” [20], which can be mapped to the predicate encryption framework as follows: fix a set $A = \{1, \dots, \ell\}$ and let the set of attributes be all subsets of A . Predicates take the form $\Phi = \{\phi_S \mid S \subseteq A\}$ where $\phi_S(S') = 1$ iff $|S \cap S'| \geq t$, i.e., S and S' overlap in *at least* t positions. Sahai and Waters [20] show a construction of a payload-hiding predicate encryption scheme for this class of predicates.

We can construct a scheme where the attribute space is the same as before, but the class of predicates corresponds to overlap in *exactly* t positions. (Our scheme will also be attribute hiding.) Namely, set $\Phi' = \{\phi'_S \mid S \subseteq A\}$ with $\phi'_S(S') = 1$ iff $|S \cap S'| = t$. Then, given any inner product encryption scheme of dimension $\ell + 1$:

- The setup algorithm is unchanged.
- To generate a secret key for the predicate ϕ'_S , first define a vector $\vec{v} \in \mathbb{Z}_N^{\ell+1}$ as follows:

$$\begin{aligned} \text{for } 1 \leq i \leq \ell: \quad v_i &= 1 \text{ iff } i \in S \\ v_{\ell+1} &= 1. \end{aligned}$$

Then output the key obtained by running $\text{GenKey}_{SK}(f_{\vec{v}})$.

⁴In general, the secret key may leak the value of r in which case the adversary will be able to find x_1, x_2 such that $\text{AND}_{I_1, I_2}(x_1, x_2) \neq 1$ yet $p''(x_1, x_2) = 0$. Since, however, we consider the “selective” notion of security (where the adversary must commit to x_1, x_2 at the outset of the experiment), this is not a problem in our setting.

- To encrypt a message M for the attribute $S' \subseteq A$, define a vector \vec{v}' as follows:

$$\begin{aligned} \text{for } 1 \leq i \leq \ell: \quad & v_i = 1 \text{ iff } i \in S' \\ & v_{\ell+1} = -t \bmod N. \end{aligned}$$

Then output the ciphertext $C \leftarrow \text{Enc}_{PK}(\vec{v}', M)$.

Since $|S \cap S'| = t$ exactly when $\langle \vec{v}, \vec{v}' \rangle = 0$, correctness and security follow.

References

- [1] S. Al-Riyami, J. Malone-Lee, and N. Smart. Escrow-free encryption supporting cryptographic workflow. *Intl. J. Information Security*, 5(4):217–229, 2006.
- [2] W. Bagga and R. Molva. Policy-based cryptography and applications. In *Financial Cryptography*, 2005.
- [3] J. Bethencourt, A. Sahai, and B. Waters. Ciphertext-policy attribute-based encryption. In *IEEE Symposium on Security and Privacy*, 2007.
- [4] D. Boneh and X. Boyen. Efficient selective-ID identity based encryption without random oracles. In *Advances in Cryptology — Eurocrypt 2004*.
- [5] D. Boneh and X. Boyen. Secure identity based encryption without random oracles. In *Advances in Cryptology — Crypto 2004*.
- [6] D. Boneh, X. Boyen, and E. Goh. Hierarchical identity based encryption with constant size ciphertexts. In *Proceedings of Eurocrypt '05*, 2005.
- [7] D. Boneh, X. Boyen, and H. Shacham. Short group signatures. In *Advances in Cryptology — Crypto 2004*.
- [8] D. Boneh, G. Di Crescenzo, R. Ostrovsky, and G. Persiano. Public-key encryption with keyword search. In *Advances in Cryptology — Eurocrypt 2004*.
- [9] D. Boneh and M. Franklin. Identity-based encryption from the Weil pairing. *SIAM J. Computing*, 32(3):586–615, 2003.
- [10] D. Boneh, E.-J. Goh, and K. Nissim. Evaluating 2-DNF formulas on ciphertexts. In *Theory of Cryptography Conference*, 2005.
- [11] D. Boneh and B. Waters. Conjunctive, subset, and range queries on encrypted data. In *Theory of Cryptography Conference*, 2007.
- [12] X. Boyen and B. Waters. Anonymous hierarchical identity-based encryption (without random oracles). In *Advances in Cryptology — Crypto 2006*.
- [13] R. Canetti, S. Halevi, and J. Katz. A forward-secure public-key encryption scheme. In *Advances in Cryptology — Eurocrypt 2003*.
- [14] C. Cocks. An identity based encryption scheme based on quadratic residues. In *Proc. IMA Intl. Conf. on Cryptography and Coding*, 2001.

- [15] C. Gentry. Practical identity-based encryption without random oracles. In *Advances in Cryptology — Eurocrypt 2006*.
- [16] V. Goyal, O. Pandey, A. Sahai, and B. Waters. Attribute-based encryption for fine-grained access control of encrypted data. In *ACM CCCS*, 2006.
- [17] J. Groth and A. Sahai. Efficient non-interactive proof systems for bilinear groups. 2008.
- [18] M. Naor and M. Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *STOC*, pages 427–437, 1990.
- [19] R. Ostrovsky, A. Sahai, and B. Waters. Attribute-based encryption with non-monotonic access structures. In *ACM CCCS*, 2007.
- [20] A. Sahai and B. Waters. Fuzzy identity-based encryption. In *Advances in Cryptology — Eurocrypt 2005*.
- [21] A. Shamir. Identity-based cryptosystems and signature schemes. In *Advances in Cryptology — Crypto '84*.
- [22] E. Shi, J. Bethencourt, H. T.-H. Chan, D. X. Song, and A. Perrig. Multi-dimensional range queries over encrypted data. In *IEEE Symposium on Security and Privacy*, 2007.
- [23] B. Waters. Efficient identity-based encryption without random oracles. In *Advances in Cryptology — Eurocrypt 2005*.

A Security Definition for Inner-Product Encryption

Here, we re-state Definition 2.2 in the particular setting of our main construction, which is a predicate-only scheme where the set of attributes⁵ is $\Sigma = \mathbb{Z}_N^n$ and the class of predicates is $\mathcal{F} = \{f_{\vec{x}} \mid \vec{x} \in \mathbb{Z}_N^n\}$ such that $f_{\vec{x}}(\vec{y}) = 1 \Leftrightarrow \langle \vec{x}, \vec{y} \rangle = 0$.

Definition A.1. *A predicate-only encryption scheme for Σ, \mathcal{F} as above is attribute-hiding if for all PPT adversaries \mathcal{A} , the advantage of \mathcal{A} in the following experiment is negligible in the security parameter n :*

1. **Setup**(1^n) is run to generate keys PK, SK . This defines a value N which is given to \mathcal{A} .
2. \mathcal{A} outputs $\vec{x}, \vec{y} \in \mathbb{Z}_N^n$, and is then given PK .
3. \mathcal{A} may adaptively request keys corresponding to the vectors $\vec{v}_1, \dots, \vec{v}_\ell \in \mathbb{Z}_N^n$, subject to the restriction that, for all i , $\langle \vec{v}_i, \vec{x} \rangle = 0$ if and only if $\langle \vec{v}_i, \vec{y} \rangle = 0$. In response, \mathcal{A} is given the corresponding keys $SK_{\vec{v}_i} \leftarrow \text{GenKey}_{SK}(f_{\vec{v}_i})$.
4. A random bit ν is chosen. If $\nu = 0$ then \mathcal{A} is given $C \leftarrow \text{Enc}_{PK}(\vec{x})$, and if $\nu = 1$ then \mathcal{A} is given $C \leftarrow \text{Enc}_{PK}(\vec{y})$.
5. The adversary may continue to request keys for additional vectors, subject to the same restriction as before.

⁵Technically speaking, both Σ and \mathcal{F} depend on the public parameters (since N is generated as part of PK), but we ignore this technicality. We remark also that we consider vectors of length n , the security parameter, for convenience only.

6. \mathcal{A} outputs a bit ν' , and succeeds if $\nu' = \nu$.

The advantage of \mathcal{A} is the absolute value of the difference between its success probability and $1/2$.

B Justifying our Assumptions in the Generic Group Model

We justify our assumptions by showing that they hold in generic bilinear groups of composite order N , as long as finding a non-trivial factor of N is hard. In doing so, we first sketch a “master theorem” for proving that certain assumptions hold in such groups. Our theorem generalizes the result by Boneh, Boyen, and Goh [6]; in addition to handling groups of composite order, our framework can be used for proving security of assumptions where the target element is in the bilinear group \mathbb{G} (instead of the target group \mathbb{G}_T). Thus, it also applies to assumptions such as the linear assumption introduced by Boneh, Boyen, and Shacham [7] and the subgroup decision assumption introduced by Boneh, Goh, and Nissim [10].

B.1 A “Master Theorem” for Hardness in Composite Order Bilinear Groups

Let $(N, \mathbb{G}, \mathbb{G}_T, \hat{e})$ represent a bilinear group of order N , where N is the product of m distinct, equal-length primes p_1, \dots, p_m . Every element $h \in \mathbb{G}$ can be written as $h = g_{p_1}^{a_1} g_{p_2}^{a_2} \dots g_{p_m}^{a_m}$, where $a_i \in \mathbb{Z}_{p_i}$ and g_{p_i} denotes a generator of the subgroup of order p_i . We can therefore represent each element $h \in \mathbb{G}$ as an m -tuple (a_1, \dots, a_m) . We can do the same with elements in \mathbb{G}_T (with respect to the generators $\hat{e}(p_i, p_i)$), and will represent elements in \mathbb{G}_T as bracketed tuples $[a_1, \dots, a_m]$. We will work with tuples of this sort for notational (and mathematical) convenience, but stress that when, e.g., we say “the adversary is given (a_1, \dots, a_m) ” we mean that it is given the element $h \in \mathbb{G}$ whose representation is (a_1, \dots, a_m) .

Note also that if $h \in \mathbb{G}$ corresponds to (a_1, \dots, a_m) and $h' \in \mathbb{G}$ corresponds to (a'_1, \dots, a'_m) , then $\hat{e}(h, h')$ corresponds to $[a_1 a'_1, \dots, a_m a'_m]$ (where multiplication in the i th component is performed modulo p_i). We write $(a_1, \dots, a_m) + (b_1, \dots, b_m) = (a_1 + b_1, \dots, a_m + b_m)$ and, for $\gamma \in \mathbb{Z}_N$, we write $\gamma \cdot (a_1, \dots, a_m) = (\gamma a_1, \dots, \gamma a_m)$.

We will also refer to random variables in \mathbb{G} and \mathbb{G}_T , and denote these by (X_1, \dots, X_m) and $[X_1, \dots, X_m]$, respectively, where X_i takes values in \mathbb{Z}_{p_i} . Given such a tuple $X = (X_1, \dots, X_m)$ and a set $B = \{B_i\}$ where $B_i \stackrel{\text{def}}{=} (X_1^i, \dots, X_m^i)$, we say that X is *dependent* on B if there exist $\gamma_i \in \mathbb{Z}_N$ such that $\Pr[X = \sum_i \gamma_i \cdot B_i]$ is not negligible. Otherwise, X is *independent* of B . We define dependence and independence in \mathbb{G}_T analogously.

Theorem B.1. *Let $\{A_i\}$ be a set of random variables in \mathbb{G} , where $A_i = (X_1^i, \dots, X_m^i)$, and let $\{B_i\}$ be a set of random variables in \mathbb{G}_T , where $B_i = [Y_1^i, \dots, Y_m^i]$. Let T_0 and T_1 be random variables in \mathbb{G} . (We stress that all these random variables are in the same space, and may therefore be correlated arbitrarily.) Let $S \stackrel{\text{def}}{=} \{i \mid \hat{e}(T_0, A_i) \neq \hat{e}(T_1, A_i)\}$, and assume that for all $k \in S$ it holds that $\hat{e}(T_0, A_k)$ is independent of $\{B_i\} \cup \{\hat{e}(A_i, A_j)\} \cup \{\hat{e}(A_i, T_0)\}_{i \neq k}$ (and similarly for $\hat{e}(T_1, A_k)$).*

Consider the following experiment:

An adversary is given $(N, \mathbb{G}, \mathbb{G}_T, \hat{e})$ (generated using $\mathcal{G}(1^n)$) and $\{a_i = (x_1^i, \dots, x_m^i)\}$, $\{b_i = [y_1^i, \dots, y_m^i]\}$ distributed appropriately. A random bit ν is chosen, and the adversary is also given T_ν . The adversary makes queries to its generic group oracles, outputs a guess for ν , and succeeds if this guess is correct. The adversary’s advantage is the absolute value of the difference between its success probability and $1/2$.

Then assuming it is hard to find a non-trivial factor of N , the advantage of any polynomial-time adversary as above is negligible.

Theorem B.2. Let $\{A_i\}$, $\{B_i\}$ be as in the previous theorem, and let T_0 , and T_1 be random variables in \mathbb{G}_T . Assume that T_0 and T_1 are each independent of $\{B_i\} \cup \{\hat{e}(A_i, A_j)\}$. Consider the following experiment:

An adversary is given $(N, \mathbb{G}, \mathbb{G}_T, \hat{e})$ (generated using $\mathcal{G}(1^n)$) and $\{a_i = (x_1^i, \dots, x_m^i)\}$, $\{b_i = [y_1^i, \dots, y_m^i]\}$ distributed appropriately. A random bit ν is chosen, and the adversary is also given T_ν . The adversary makes queries to its generic group oracles, and outputs a guess for ν ; its advantage is defined as before.

Then assuming it is hard to find a non-trivial factor of N , the advantage of any polynomial-time adversary as above is negligible.

Proof Sketch. We now sketch a proof of the above theorems.

- Two encoding lists E and E_T are kept in the generic group model (one for the bilinear group and one for the target group). The encoding lists contain pairs of symbolic polynomials and encodings of length ν , where ν is chosen to be sufficiently large to avoid collisions for random encodings. Every group operation requested by the adversary will result in the polynomials two elements from the list being either component-wise added or subtracted (depending on whether the requested operation was multiplication or division). Note again that polynomials in the i -th component are added in \mathbb{Z}_{p_i} . The generic group simulator checks to see whether the resulting polynomial is already on the list. If it is, it gives out the previous encoding; otherwise it picks a new random encoding and adds the polynomial m -tuple and new encoding to the encoding list. The bilinear map operation is modeled by multiplying the polynomials of two elements from the list E and checking the result against the encodings in E_T , adding a new encoding if necessary.
- To prove the generic security of our assumption we define a set of three hybrid security games (in the generic group) model.

The first game is identical to the real (generic model) security game. In this game the challenger controls the generic group interface. It chooses random values for the all variables in the assumption. E.g. variables for the i -th component are chosen randomly in \mathbb{Z}_{p_i} . Recall that $T_0, T_1 \in \mathbb{G}$ are the two m -component tuples that the adversary is attempting to distinguish between. The challenger flips a coin ν and give the adversary the target, which is based on the coin flip $T = T_\nu$. The adversary then attempts to guess ν with non-negligible advantage.

- In the second game the generic group will output encodings such that two encoding will only be the same if they are *symbolically* equivalent. A standard argument using the Schwartz-Zippel lemma shows that with all but negligible probability the adversary's view in these two experiment will be the same. (Note the concrete parameters depend upon the number of queries made and degrees of the polynomials in use.)
- In the next game we define the challenger will still interpret elements as m -tuples, however, for the i -th component it will keep track of the coefficients in \mathbb{Z}_N instead of \mathbb{Z}_{p_i} . Two different elements will have a different representation if any of their coefficients are different in \mathbb{Z}_N .

The only way the adversary will notice the difference between this game and the previous game is requests the encoding of a m -tuple polynomial such that in the previous game it was symbolically equivalent to a previous requested encoding, but results in a different encoding in this game. In this case, the coefficients for a particular polynomial in two elements in the list are the same mod \mathbb{Z}_{p_i} for some i , but are different mod \mathbb{Z}_N . If this condition happens with non-negligible probability then we can use the generic adversary to find non-trivial factors of N .

Intuitively, this final game puts us in a situation where the adversary needs to attack a set of polynomials where each component is defined over \mathbb{Z}_N . This last hybrid step effectively deals with the complexity that arises from using composite order groups.

- Suppose that the adversary is given elements represented by the polynomial tuples A_1, \dots, A_s in the bilinear group and $B_1, \dots, B_{s'}$ in the target group \mathbb{G}_T . In addition, the adversary is given the target T .

Consider a m -tuple of polynomials with coefficients in \mathbb{Z}_n formed as:

$$\sum_{i=1, \dots, s} t_i \cdot T \cdot A_i + \sum_{i,j=1, \dots, s} a_{i,j} \cdot A_{i,j} \cdot A_j + \sum_{i=1, \dots, s'} b_i B_i$$

where $a_{i,j}, t_i, b_i$ values are in \mathbb{Z}_N . An adversary can win the game iff he can find coefficients such that such a symbolic expression evaluates to 0 *conditionally* on the bit ν .

We now define a sufficient condition, which when met will ensure that this will not occur. Let S be the set of all k such that the polynomial tuple $T \cdot A_k$ is dependent upon ν . The condition we must show is that for each k in S that $A \cdot T$ is linearly independent of the polynomials $B_i, A_i \cdot A_j$ for all i, j and $A_i T$ for $i \neq k$.

If our condition holds for all elements in S then any expression of the above form must not include any non-zero coefficient for the terms $T \cdot A_k$ for $k \in S$. However, by definition all terms of the form $T \cdot A_k$ for $k \notin S$ are independent of ν therefore, the adversary will have no advantage in guessing ν .

- For the case when the target of the assumption is in the group \mathbb{G}_T the argument is even simpler as we only need to consider whether T is linearly dependent of $B_i, A_i \cdot A_j$ for all i, j .

□

B.2 Applying the Master Theorem to Our Assumptions

For completeness, we show how to apply the theorems of the previous section to prove that our assumptions hold in the generic group model.

Assumption 1. Recall that in defining this assumption, we choose random $Q_1, Q_2, Q_3 \in \mathbb{G}_q$, random $R_1, R_2, R_3 \in \mathbb{G}_r$, random $a, b, s \in \mathbb{Z}_p$, and a random bit ν ; the adversary \mathcal{A} is given $(N, \mathbb{G}, \mathbb{G}_T, \hat{e})$ as well as

$$g_p, g_r, g_q R_1, g_p^b, g_p^{b^2}, g_p^a g_q, g_p^{ab} Q_1, g_p^s, g_p^{bs} Q_2 R_2.$$

Furthermore, if $\nu = 0$ then \mathcal{A} is given $T_0 = g_p^{b^2 s} R_3$, while if $\nu = 1$ then \mathcal{A} is given $T_1 = g_p^{b^2 s} Q_3 R_3$.

Write $R_i = g_r^{c_i}$ and $Q_i = g_q^{d_i}$. Thus, the values in the assumption correspond to the following tuples in $\mathbb{G}_p \times \mathbb{G}_q \times \mathbb{G}_r$:

$$\begin{aligned} A_1 &= (1, 0, 0), & A_2 &= (0, 0, 1), & A_3 &= (0, 1, c_1), \\ A_4 &= (b, 0, 0), & A_5 &= (b^2, 0, 0), & A_6 &= (a, 1, 0), \\ A_7 &= (ab, d_1, 0), & A_8 &= (s, 0, 0), & A_9 &= (bs, d_2, c_2), \end{aligned}$$

and the two possible target values are: $T_0 = (b^2s, 0, c_3)$ and $T_1 = (b^2s, d_3, c_3)$.

Since $T_0, T_1 \in \mathbb{G}$, we are in the setting of Theorem B.1. Note that only A_3, A_6, A_7, A_9 give different results when paired with T_0 or T_1 . Considering T_0 first, we obtain the following tuples:

$$\begin{aligned} \hat{e}(A_3, T_0) &= C_1 \stackrel{\text{def}}{=} [0, 0, c_1c_3] & \hat{e}(A_6, T_0) &= C_2 \stackrel{\text{def}}{=} [ab^2s, 0, 0] \\ \hat{e}(A_7, T_0) &= C_3 \stackrel{\text{def}}{=} [ab^3s, 0, 0] & \hat{e}(A_9, T_0) &= C_4 \stackrel{\text{def}}{=} [b^3s^2, 0, c_2c_3]. \end{aligned}$$

Assumption 2. We now move on to Assumption 2. Recall that in defining this assumption, we choose random $h \in \mathbb{G}_p$ and $Q_1, Q_2 \in \mathbb{G}_q$, random $s, \gamma \in \mathbb{Z}_q$, and a random bit ν ; the adversary \mathcal{A} is given $(N, \mathbb{G}, \mathbb{G}_T, \hat{e})$ as well as

$$g_p, g_q, g_r, h, g_p^s, h^sQ_1, g_p^\gamma Q_2, \hat{e}(g_p, h)^\gamma.$$

If $\nu = 0$ then \mathcal{A} is given $T_0 = \hat{e}(g_p, h)^{\gamma s}$, while if $\nu = 1$ then \mathcal{A} is given a random element $T_1 \in \mathbb{G}_T$.

Write $Q_i = g_q^{d_i}$ and $h = g_p^a$. Thus, the given values in the assumption correspond to the following tuples:

$$\begin{aligned} A_1 &= (1, 0, 0), & A_2 &= (0, 1, 0), & A_3 &= (0, 0, 1), \\ A_4 &= (a, 0, 0), & A_5 &= (s, 0, 0), & A_6 &= (as, d_1, 0), & A_7 &= (\gamma, d_2, 0), \\ & & & & B_1 &= [a\gamma, 0, 0] \end{aligned}$$

The two possible target tuples are $T_0 = [a\gamma s, 0, 0]$ or $T_1 = [e, f, g]$.

C Proof of Security

This section is devoted to a proof of the following theorem:

Theorem C.1. *If \mathcal{G} satisfies Assumption 1 then the scheme described in Section 4 is an attribute-hiding, predicate-only encryption scheme.*

Throughout, we will refer to the experiment as described in Definition A.1. We establish the theorem using a sequence of games, defined as follows:

Game₁: The challenge ciphertext is generated as a proper encryption using \vec{x} . (Recall from Definition A.1 that we let \vec{x}, \vec{y} denote the two vectors output by the adversary.) That is, we choose random $s, \alpha, \beta \in \mathbb{Z}_N$ and random $\{R_{3,i}, R_{4,i}\} \in \mathbb{G}_r$ and compute the ciphertext as

$$C = \left(C_0 = g_p^s, \quad \left\{ C_{1,i} = H_{1,i}^s Q^{\alpha x_i} R_{3,i}, \quad C_{2,i} = H_{2,i}^s Q^{\beta x_i} R_{4,i} \right\}_{i=1}^n \right).$$

Game₂: We now generate the $\{C_{2,i}\}$ components as if encryption were done using $\vec{0}$. That is, we choose random $s, \alpha, \beta \in \mathbb{Z}_N$ and random $\{R_{3,i}, R_{4,i}\} \in \mathbb{G}_r$ and compute the ciphertext as

$$C = \left(C_0 = g_p^s, \quad \{ C_{1,i} = H_{1,i}^s Q^{\alpha x_i} R_{3,i}, \quad C_{2,i} = H_{2,i}^s R_{4,i} \}_{i=1}^n \right).$$

Game₃: We now generate the $\{C_{2,i}\}$ components using vector \vec{y} . That is, we choose random $s, \alpha, \beta \in \mathbb{Z}_N$ and random $\{R_{3,i}, R_{4,i}\} \in \mathbb{G}_r$ and compute the ciphertext as

$$C = \left(C_0 = g_p^s, \quad \{ C_{1,i} = H_{1,i}^s Q^{\alpha x_i} R_{3,i}, \quad C_{2,i} = H_{2,i}^s Q^{\beta y_i} R_{4,i} \}_{i=1}^n \right).$$

Game₄ and **Game₅**: These games are defined symmetrically to **Game₂** and **Game₃**: In **Game₄** the $\{C_{i,1}\}$ components are generated using $\vec{0}$. That is, we choose random $s, \alpha, \beta \in \mathbb{Z}_N$ and random $\{R_{3,i}, R_{4,i}\} \in \mathbb{G}_r$ and compute the ciphertext as

$$C = \left(C_0 = g_p^s, \quad \{ C_{1,i} = H_{1,i}^s R_{3,i}, \quad C_{2,i} = H_{2,i}^s Q^{\beta y_i} R_{4,i} \}_{i=1}^n \right).$$

In **Game₅**, the $\{C_{i,1}\}$ components are generated using \vec{y} . I.e., we choose random $s, \alpha, \beta \in \mathbb{Z}_N$ and random $\{R_{3,i}, R_{4,i}\} \in \mathbb{G}_r$ and compute the ciphertext as

$$C = \left(C_0 = g_p^s, \quad \{ C_{1,i} = H_{1,i}^s Q^{\alpha y_i} R_{3,i}, \quad C_{2,i} = H_{2,i}^s Q^{\beta y_i} R_{4,i} \}_{i=1}^n \right).$$

In **Game₅** the challenge ciphertext is a proper encryption with respect to the vector \vec{y} . So, the proof of the theorem is concluded once we show that the adversary cannot distinguish between **Game_i** and **Game_{i+1}** for each i .

As discussed in Section 4.2, it is difficult to proceed directly from a game in which the challenge ciphertext is generated as a proper encryption using \vec{x} , to a game in which the challenge ciphertext is generated as a proper encryption using \vec{y} . (Indeed, this is the reason our construction uses two “sub-systems” to begin with.) That is why our proof proceeds via the intermediate **Game₃** where half of the challenge ciphertext corresponds to an encryption using \vec{x} and the other half corresponds to an encryption using \vec{y} . Intermediate games **Game₂** and **Game₄** are used to simplify the proof; informally speaking, it helps when part of the ciphertext corresponds to an encryption using $\vec{0}$ since this vector is orthogonal to everything.

The main difficulty in our proofs will be to answer queries for decryption keys. In considering the indistinguishability of **Game₁** and **Game₂** (and, symmetrically, **Game₄** and **Game₅**), we will actually be able to construct *all* decryption keys (i.e., even keys that would allow the adversary to distinguish an encryption relative to \vec{x} from an encryption relative to \vec{y}). In essence, we will be showing that even such keys cannot be used to distinguish a well-formed encryption of \vec{x} (or \vec{y}) from a badly-formed one.

On the other hand, in considering the indistinguishability of **Game₂** and **Game₃** (and, symmetrically, **Game₃** and **Game₄**) we will not be able to construct all decryption keys. Instead, we will deal separately with the problems of (1) providing keys for vectors \vec{v} with $\langle \vec{v}, \vec{x} \rangle = 0 = \langle \vec{v}, \vec{y} \rangle$ and (2) providing keys for vectors \vec{v} with $\langle \vec{v}, \vec{x} \rangle \neq 0 \neq \langle \vec{v}, \vec{y} \rangle$.

C.1 Indistinguishability of Game₁ and Game₂

Fix an adversary \mathcal{A} . We describe a simulator who is given $(N = pqr, \mathbb{G}, \mathbb{G}_T, \hat{e})$ along with the elements $g_p, g_r, g_q R_1, h_p = g_p^b, k_p = g_p^{b^2}, g_p^a g_q, g_p^{ab} Q_1, g_p^s, g_p^{bs} Q_2 R_2$, and an element $T = g_p^{b^2 s} g_q^\beta R_3$ where β is either 0 or uniform in \mathbb{Z}_q (cf. Assumption 1).

Before describing the simulation in detail, we observe that the simulator can sample a random element $R \in \mathbb{G}_r$ by choosing random $\delta \in \mathbb{Z}_N$ and setting $R = g_r^\delta$. Although there does not appear to be any way for the simulator to sample a random element of \mathbb{G}_q (since g_q is not provided to the simulator), it is possible for the simulator to choose a random element $QR \in \mathbb{G}_{qr} \stackrel{\text{def}}{=} \mathbb{G}_q \times \mathbb{G}_r$: this can be done by choosing random $\delta_1, \delta_2 \in \mathbb{Z}_N$ and setting $QR = (g_q R_1)^{\delta_1} \cdot g_r^{\delta_2}$. Henceforth, we simply describe the simulator as sampling uniformly from \mathbb{G}_r and \mathbb{G}_{qr} with the understanding that such sampling is done in this way.

Public parameters. The simulator begins by giving N to \mathcal{A} , who outputs vectors \vec{x}, \vec{y} . The simulator chooses random $\{w_{1,i}, w_{2,i}\} \in \mathbb{Z}_N$ and random $\{R_{1,i}, R_{2,i}\} \in \mathbb{G}_r$, includes $(N, \mathbb{G}, \mathbb{G}_T, \hat{e})$ in the public parameters, and sets the remaining values as follows:

$$PK = (g_p, g_r, g_q R_1, \{H_{1,i} = (h_p)^{x_i} g_p^{w_{1,i}} R_{1,i}, H_{2,i} = (k_p)^{x_i} g_p^{w_{2,i}} R_{2,i}\}).$$

By doing so, the simulator is implicitly setting $h_{1,i} = h_p^{x_i} g_p^{w_{1,i}}$ and $h_{2,i} = k_p^{x_i} g_p^{w_{2,i}}$. Note that PK has the appropriate distribution.

Key derivation. We now describe how the simulator prepares the secret key corresponding to the vector $\vec{v} = (v_1, \dots, v_n)$. We stress that although Definition A.1 restricts the vectors \vec{v} for which the adversary is allowed to request secret keys, we do not rely on this restriction here. This is because the purpose of this hybrid proof is to show that the adversary cannot distinguish between properly formed encryptions of \vec{x} and improperly formed encryptions (a combination of an encryption of \vec{x} and $\vec{0}$).

We begin with some intuition: We must construct the $K_{1,i}$ and $K_{2,i}$ components of the key. Note that we do not have access to g_q , but we do have $g_q g_p^a$. We will make use of this element from the assumption here. This will give rise to terms containing a in the exponent of g_p . Note, however, that we will later have to construct the K component of the key, whose purpose is to cancel out terms in the G_p subgroup. If $\langle \vec{v}, \vec{x} \rangle \neq 0$, then additional terms involving ab and ab^2 will have to appear in K . However, we do not have access to $g_p^{ab^2}$; indeed if we did, the assumption would be false and we could easily distinguish between Game₁ and Game₂. We deal with this problem by adding a term (using the $g_p^{ab} g_q^d$ term given in the assumption) to the $K_{1,i}$ components that will allow us to cancel out the ab^2 terms that will appear in K due to the $K_{2,i}$ components.

The simulator begins by choosing random $f'_1, f'_2, \{r'_{1,i}\}, \{r'_{2,i}\} \in \mathbb{Z}_N$. In constructing the key, the simulator will be implicitly setting:

$$r_{1,i} = r'_{1,i} + v_i \cdot (a f'_1 - a b f'_2) \quad (1)$$

$$r_{2,i} = r'_{2,i} + a f'_2 v_i, \quad (2)$$

as well as $f_1 = f'_1 - d f'_2$ and $f_2 = f'_2$, where we set $d = \log_{g_q} Q_1$. Note that these values are each independently and uniformly distributed in \mathbb{Z}_N , just as they would be in actual secret key components.

Next, for all i it computes:

$$\begin{aligned} K_{1,i} &= (g_p^a g_q)^{f_1' v_i} \cdot (g_p^{ab} Q_1)^{-f_2' v_i} \cdot g_p^{r_{1,i}'} \\ &= g_p^{(af_1' - abf_2') \cdot v_i + r_{1,i}'} \cdot g_q^{(f_1' - df_2') \cdot v_i} \end{aligned}$$

and

$$\begin{aligned} K_{2,i} &= (g_p^a g_q)^{f_2' v_i} \cdot g_p^{r_{2,i}'} \\ &= g_p^{af_2' v_i + r_{2,i}'} \cdot g_q^{f_2' v_i}. \end{aligned}$$

Now, to construct the K element for the decryption key. Recall that $h_{1,i} = (g_p)^{bx_i} g_p^{w_{1,i}}$. Therefore, the exponents in K will contain a term of the form $\sum_i r_{1,i} bx_i$. But because of how we chose $r_{1,i}$, we have that $\sum_i r_{1,i} bx_i = k(abf_1' - ab^2 f_2) + \sum_i r_{1,i}' x_i$ where $k = \langle \vec{v}, \vec{x} \rangle$. A similar equation holds for the terms arising out of the $h_{2,i}$ parts of K , and allows us to cancel out all the ab^2 terms that arise in K . Thus, we can compute K as follows:

Let $k = \langle \vec{v}, \vec{x} \rangle$. Finally, the simulator chooses random $QR \in \mathbb{G}_{qr}$ and computes

$$\begin{aligned} K &= QR \cdot (g_p^{ab} Q_1)^{-k \cdot f_1'} \\ &\quad \cdot \prod_i (g_p^a g_q)^{-f_1' v_i w_{1,i} - f_2' v_i w_{2,i}} \cdot (g_p^{ab} Q_1)^{f_2' v_i w_{1,i}} \cdot g_p^{-w_{1,i} r_{1,i}' - w_{2,i} r_{2,i}'} \cdot h_p^{-x_i r_{1,i}'} \cdot k_p^{-x_i r_{2,i}'} \end{aligned}$$

The simulator then hands the adversary $SK_{\vec{v}} = (K, \{K_{1,i}, K_{2,i}\}_{i=1}^n)$ as the key.

To see formally that the K component has the correct distribution, let $K_p, K_q,$ and K_r denote the projections of K in $\mathbb{G}_p, \mathbb{G}_q,$ and $\mathbb{G}_r,$ respectively. It is easy to see that K_q and K_r are independently and uniformly distributed, as required. Furthermore,

$$\begin{aligned} K_p &= g_p^{-abk f_1'} \cdot \prod_i g_p^{-af_1' v_i w_{1,i} - af_2' v_i w_{2,i}} g_p^{abf_2' v_i w_{1,i}} g_p^{-w_{1,i} r_{1,i}' - w_{2,i} r_{2,i}'} h_p^{-x_i r_{1,i}'} k_p^{-x_i r_{2,i}'} \\ &= h_p^{-ak f_1'} \prod_i \left(h_p^{-x_i r_{1,i}'} g_p^{-w_{1,i} r_{1,i}'} g_p^{-w_{1,i} v_i (af_1' - abf_2')} \right) \cdot \left(k_p^{-x_i r_{2,i}'} g_p^{-w_{2,i} r_{2,i}'} g_p^{-w_{2,i} af_2' v_i} \right) \\ &= \prod_i h_p^{-ax_i v_i f_1'} \cdot \left(h_p^{-x_i r_{1,i}'} g_p^{-w_{1,i} r_{1,i}'} g_p^{-w_{1,i} v_i (af_1' - abf_2')} \right) \cdot \left(h_p^{abx_i v_i f_2'} \cdot h_p^{-abx_i v_i f_2'} \right) \\ &\quad \cdot \left(k_p^{-x_i r_{2,i}'} g_p^{-w_{2,i} r_{2,i}'} g_p^{-w_{2,i} af_2' v_i} \right), \end{aligned}$$

using the fact that $k = \langle \vec{x}, \vec{v} \rangle = \sum_i x_i v_i$. Using simple (but tedious) algebra, we obtain

$$\begin{aligned} K_p &= \prod_i \left(h_p^{-x_i r_{1,i}'} g_p^{-w_{1,i} r_{1,i}'} h_p^{-x_i v_i (af_1' - abf_2')} g_p^{-w_{1,i} v_i (af_1' - abf_2')} \right) \cdot \left(k_p^{-x_i r_{2,i}'} g_p^{-w_{2,i} r_{2,i}'} k_p^{-x_i af_2' v_i} g_p^{-w_{2,i} af_2' v_i} \right) \\ &= \prod_i (h_p^{x_i} g_p^{w_{1,i}})^{-r_{1,i}} (k_p^{x_i} g_p^{w_{2,i}})^{-r_{2,i}} = \prod_i h_{1,i}^{-r_{1,i}} h_{2,i}^{-r_{2,i}} \end{aligned}$$

(using Eqs. (1) and (2)), and thus K_p has the correct distribution.

The challenge ciphertext. The challenge ciphertext is generated in a straightforward way, as follows. The simulator chooses $\{R_{7,i}, R_{8,i}\} \in \mathbb{G}_r$ at random, sets C_0 equal to g_p^s , and computes:

$$\begin{aligned} C_{1,i} &= \left(g_p^{bs} Q_2 R_2\right)^{x_i} \cdot (g_p^s)^{w_{1,i}} \cdot R_{7,i} \\ &= h_p^{x_i s} g_p^{w_{1,i} s} Q_2^{x_i} R'_{7,i} \\ &= (h_{1,i})^s Q_2^{x_i} R'_{7,i} \\ C_{2,i} &= T^{x_i} \cdot (g_p^s)^{w_{2,i}} \cdot R_{8,i} \\ &= (h_{2,i})^s \left(g_q^\beta\right)^{x_i} R'_{8,i}, \end{aligned}$$

where $\{R'_{7,i}, R'_{8,i}\}$ refer to elements of \mathbb{G}_r whose exact values are unimportant.

Analysis. By examining the projections of the components of the challenge ciphertext in the groups \mathbb{G}_p , \mathbb{G}_q , and \mathbb{G}_r , it can be verified that when β is random the challenge ciphertext is distributed exactly as in Game_1 , whereas if $\beta = 0$ the challenge ciphertext is distributed exactly as in Game_2 . We conclude that, under Assumption 1, these two games are indistinguishable. ■

C.2 Indistinguishability of Game_2 and Game_3

Fix again some adversary \mathcal{A} . We describe a simulator who is given $(N = pqr, \mathbb{G}, \mathbb{G}_T, \hat{e})$ along with the elements $g_p, g_r, g_q R_1, h_p = g_p^b, k_p = g_p^{b^2}, g_p^a g_q, g_p^{ab} Q_1, g_p^s, g_p^{bs} Q_2 R_2$, and an element $T = g_p^{b^2 s} g_q^\beta R_3$ where β is either 0 or uniform in \mathbb{Z}_q . Recall that sampling uniform elements from \mathbb{G}_r or \mathbb{G}_{qr} can be done efficiently. The simulator interacts with \mathcal{A} as we now describe.

Public parameters. The simulator begins by giving N to \mathcal{A} , who outputs vectors \vec{x}, \vec{y} . The simulator chooses random $\{w_{1,i}, w_{2,i}\} \in \mathbb{Z}_N$ and random $\{R_{1,i}, R_{2,i}\} \in \mathbb{G}_r$, includes $(N, \mathbb{G}, \mathbb{G}_T, \hat{e})$ in the public parameters, and sets the public parameters as follows:

$$PK = \left(g_p, g_r, g_q R_1, \left\{H_{1,i} = (h_p)^{x_i} g_p^{w_{1,i}} R_{1,i} \quad H_{2,i} = (k_p)^{y_i} g_p^{w_{2,i}} R_{2,i}\right\}\right).$$

By doing so, the simulator is implicitly setting $h_{1,i} = h_p^{x_i} g_p^{w_{1,i}}$ and $h_{2,i} = k_p^{y_i} g_p^{w_{2,i}}$. Note that PK has the appropriate distribution.

Key derivation. The adversary \mathcal{A} may request secret keys corresponding to different vectors, and we now describe how the simulator prepares the secret key corresponding to the vector $\vec{v} = (v_1, \dots, v_n)$. Here, the simulator will only be able to produce the appropriate secret key when the vector \vec{v} satisfies the restriction imposed by Definition A.1. We distinguish two cases, depending on whether $\langle \vec{v}, \vec{x} \rangle$ and $\langle \vec{v}, \vec{y} \rangle$ are both 0 or whether they are both non-zero.

Case 1. We first consider the case where $\langle \vec{v}, \vec{x} \rangle = 0 = \langle \vec{v}, \vec{y} \rangle$. The simulator begins by choosing random $f_1, f_2, \{r'_{1,1}\}, \{r'_{2,1}\} \in \mathbb{Z}_N$. Then for all i it computes:

$$\begin{aligned} K_{1,i} &= \left(g_p^a g_q\right)^{f_1 v_i} \cdot (g_p)^{r'_{1,i}} \\ &= g_p^{a f_1 v_i + r'_{1,i}} \cdot g_q^{f_1 v_i} \\ K_{2,i} &= \left(g_p^a g_q\right)^{f_2 v_i} \cdot (g_p)^{r'_{2,i}} \\ &= g_p^{a f_2 v_i + r'_{2,i}} \cdot g_q^{f_2 v_i}. \end{aligned}$$

Finally, the simulator chooses random $\mathcal{QR} \in \mathbb{G}_{qr}$ and computes

$$K = \mathcal{QR} \cdot \prod_i (g_p^a g_q)^{-f_1 v_i w_{1,i} - f_2 v_i w_{2,i}} \cdot g_p^{-w_{1,i} \cdot r'_{1,i} - w_{2,i} \cdot r'_{2,i}} \cdot h_p^{-x_i \cdot r'_{1,i}} \cdot k_p^{-y_i \cdot r'_{2,i}}.$$

The simulator then hands the adversary $SK_{\vec{v}} = (K, \{K_{1,i}, K_{2,i}\})$ as the key.

To see that this key has the correct distribution, note that by construction of the $\{K_{1,i}, K_{2,i}\}$ the values f_1, f_2 are random; furthermore, the simulator implicitly sets

$$\begin{aligned} r_{1,i} &= r'_{1,i} + a f_1 v_i \\ r_{2,i} &= r'_{2,i} + a f_2 v_i, \end{aligned}$$

which are uniformly distributed as well. Looking at K_p , the projection of K in \mathbb{G}_p (as in the proof in the previous section), we see that

$$\begin{aligned} K_p &= \prod_i g_p^{-a f_1 v_i w_{1,i} - a f_2 v_i w_{2,i}} \cdot g_p^{-w_{1,i} \cdot r'_{1,i} - w_{2,i} \cdot r'_{2,i}} \cdot h_p^{-x_i \cdot r'_{1,i}} \cdot k_p^{-y_i \cdot r'_{2,i}} \\ &= \prod_i h_p^{-a f_1 x_i v_i} \cdot k_p^{-a f_2 y_i v_i} \cdot g_p^{-a f_1 v_i w_{1,i} - a f_2 v_i w_{2,i}} \cdot g_p^{-w_{1,i} \cdot r'_{1,i} - w_{2,i} \cdot r'_{2,i}} \cdot h_p^{-x_i \cdot r'_{1,i}} \cdot k_p^{-y_i \cdot r'_{2,i}}, \end{aligned}$$

using the fact that $\prod_i h_p^{-a f_1 x_i v_i} = h_p^{-a f_1 \cdot \sum_i x_i v_i} = 1 = \prod_i k_p^{-a f_2 y_i v_i}$ (because $\langle \vec{v}, \vec{x} \rangle = 0 = \langle \vec{v}, \vec{y} \rangle$). Algebraic manipulation as in the previous section shows that K_p has the correct distribution.

Case 2. Here, we consider the case where $\langle \vec{v}, \vec{x} \rangle = c_x \neq 0$ and $\langle \vec{v}, \vec{y} \rangle = c_y \neq 0$. The simulator begins by choosing random $f'_1, f'_2, \{r'_{1,1}\}, \{r'_{2,1}\} \in \mathbb{Z}_N$. Next, for all i it computes

$$\begin{aligned} K_{1,i} &= (g_p^a g_q)^{f'_1 v_i} (g_p^{ab} Q_1)^{-c_y \cdot f'_2 v_i} \cdot (g_p)^{r'_{1,i}} \\ &= g_p^{(a f'_1 - a b c_y f'_2) \cdot v_i + r'_{1,i}} \cdot g_q^{(f'_1 - c_y d f'_2) \cdot v_i} \\ K_{2,i} &= (g_p^a g_q)^{c_x \cdot f'_2 v_i} \cdot (g_p)^{r'_{2,i}} \\ &= g_p^{a c_x f'_2 v_i + r'_{2,i}} \cdot g_q^{c_x \cdot f'_2 v_i}, \end{aligned}$$

where we set $d = \log_{g_q} Q_1$ as in the previous proof. Finally, the simulator chooses random $\mathcal{QR} \in \mathbb{G}_{qr}$ and computes

$$\begin{aligned} K &= \mathcal{QR} \cdot (g_p^{ab} Q_1)^{-c_x f'_1} \\ &\quad \cdot \prod_i (g_p^a g_q)^{-f'_1 v_i w_{1,i} - f'_2 c_x v_i w_{2,i}} \cdot (g_p^{ab} Q_1)^{f'_2 c_y v_i w_{1,i}} \cdot g_p^{-w_{1,i} \cdot r'_{1,i} - w_{2,i} \cdot r'_{2,i}} \cdot h_p^{-x_i \cdot r'_{1,i}} \cdot k_p^{-y_i \cdot r'_{2,i}}. \end{aligned}$$

The simulator then hands the key $SK_{\vec{v}} = (K, \{K_{1,i}, K_{2,i}\})$ to the adversary.

To see that this key has the correct distribution, note that by construction of the $\{K_{1,i}, K_{2,i}\}$ the simulator implicitly sets

$$\begin{aligned} r_{1,i} &= r'_{1,i} + (a f'_1 - c_y a b f'_2) \cdot v_i \\ r_{2,i} &= r'_{2,i} + a c_x f'_2 v_i, \end{aligned}$$

as well as $f_1 = f'_1 - c_y \cdot df'_2$ and $f_2 = c_x \cdot f'_2$. It is clear that f_1 and the $\{r_{1,i}, r_{2,i}\}$ are independently and uniformly distributed in \mathbb{Z}_N . The value f_2 is also uniformly distributed in \mathbb{Z}_N as long as $\gcd(c_x, N) = 1$. (If $\gcd(c_x, N) \neq 1$, then the adversary has found a non-trivial factor of N . This occurs with negligible probability under Assumption 1.)

As for element K of the secret key, it is once again easy to see that the projection of K in \mathbb{G}_{qr} is uniformly distributed. Looking at K_p , the projection of K in \mathbb{G}_p (as in the previous section), we see that

$$\begin{aligned}
K_p &= g_p^{-abc_x f'_1} \cdot \prod_i g_p^{-af'_1 v_i w_{1,i} - af'_2 c_x v_i w_{2,i}} \cdot g_p^{abf'_2 c_y v_i w_{1,i}} \cdot g_p^{-w_{1,i} \cdot r'_{1,i} - w_{2,i} \cdot r'_{2,i}} \cdot h_p^{-x_i \cdot r'_{1,i}} \cdot k_p^{-y_i \cdot r'_{2,i}} \\
&= \prod_i h_p^{-ax_i v_i f'_1} \cdot g_p^{-af'_1 v_i w_{1,i} - af'_2 c_x v_i w_{2,i}} \cdot g_p^{abf'_2 c_y v_i w_{1,i}} \cdot (h_{1,i})^{r'_{1,i}} (h_{2,i})^{r'_{2,i}} \\
&= h_p^{c_x c_y abf'_2} \cdot h_p^{-c_x c_y abf'_2} \prod_i g_p^{-af'_2 c_x v_i w_{2,i}} \cdot g_p^{abf'_2 c_y v_i w_{1,i}} \cdot (h_{1,i})^{-r'_{1,i} - av_i f'_1} (h_{2,i})^{-r'_{2,i}} \\
&= \prod_i h_p^{x_i v_i c_y abf'_2} \cdot k_p^{-c_x y_i v_i af'_2} \cdot g_p^{-af'_2 c_x v_i w_{2,i}} \cdot g_p^{abf'_2 c_y v_i w_{1,i}} \cdot (h_{1,i})^{-r'_{1,i} - av_i f'_1} (h_{2,i})^{-r'_{2,i}} \\
&= \prod_i (h_{1,i})^{-r'_{1,i} - av_i f'_1 + abf'_2 c_y v_i} (h_{2,i})^{-r'_{2,i} - ac_x v_i f'_2} = \prod_i (h_{1,i})^{-r_{1,i}} (h_{2,i})^{-r_{2,i}},
\end{aligned}$$

and so K_p has the right distribution.

The challenge ciphertext. The challenge ciphertext is generated in a straightforward way. The simulator chooses $\{R_{7,i}, R_{8,i}\} \in \mathbb{G}_r$ at random, sets $C_0 = g_p^s$, and computes:

$$\begin{aligned}
C_{1,i} &= \left(g_p^{bs} Q_2 R_2\right)^{x_i} \cdot (g_p^s)^{w_{1,i}} \cdot R_{7,i} \\
&= (h_{1,i})^s Q_2^{x_i} R'_{7,i} \\
C_{2,i} &= T^{y_i} (g_p^s)^{w_{2,i}} R_{8,i} \\
&= (h_{2,i})^s \left(g_q^\beta\right)^{y_i} R'_{8,i},
\end{aligned}$$

where $\{R'_{7,i}, R'_{8,i}\}$ again refer to elements of \mathbb{G}_r whose values are unimportant.

Analysis. By examining the projections of the components of the challenge ciphertext in the groups \mathbb{G}_p , \mathbb{G}_q , and \mathbb{G}_r , it can be verified that when β is random the challenge ciphertext is distributed exactly as in **Game₃**, whereas if $\beta = 0$ the challenge ciphertext is distributed exactly as in **Game₂**. We conclude that, under Assumption 1, these two games are indistinguishable. ■

C.3 Completing the Proof

Our scheme is symmetric with respect to the roles of $h_{1,i}$ and $h_{2,i}$. Thus, as mentioned earlier, the proof that **Game₃** and **Game₄** are indistinguishable exactly parallels the proof (given in the previous section) that **Game₂** and **Game₃** are indistinguishable, while the proof that **Game₄** and **Game₅** are indistinguishable exactly parallels the proof (given in Section C.1) that **Game₁** and **Game₂** are indistinguishable. This concludes the proof of our theorem.

D A Full-Fledged Predicate Encryption Scheme

In Section 4, we showed a construction of a *predicate-only* scheme. Here, we extend the previous scheme to obtain a full-fledged predicate encryption scheme in the sense of Definition 2.1. The construction follows.

Setup(1^n) The setup algorithm first runs $\mathcal{G}(1^n)$ to obtain $(p, q, r, \mathbb{G}, \mathbb{G}_T, \hat{e})$ with $\mathbb{G} = \mathbb{G}_p \times \mathbb{G}_q \times \mathbb{G}_r$. Next, it computes g_p, g_q , and g_r as generators of $\mathbb{G}_p, \mathbb{G}_q$, and \mathbb{G}_r , respectively. It then chooses $R_{1,i}, R_{2,i} \in \mathbb{G}_r$ and $h_{1,i}, h_{2,i} \in \mathbb{G}_p$ uniformly at random for $i = 1$ to n , and $R_0 \in \mathbb{G}_r$ uniformly at random. It also chooses random $\gamma \in \mathbb{Z}_p$ and $h \in \mathbb{G}_p$. The public parameters include $(N = pqr, \mathbb{G}, \mathbb{G}_T, \hat{e})$ along with:

$$PK = (g_p, g_r, Q = g_q \cdot R_0, P = \hat{e}(g_p, h)^\gamma, \{H_{1,i} = h_{1,i} \cdot R_{1,i}, H_{2,i} = h_{2,i} \cdot R_{2,i}\}_{i=1}^n).$$

The master secret key SK is $(p, q, r, g_q, h^{-\gamma}, \{h_{1,i}, h_{2,i}\}_{i=1}^n)$.

Enc $_{PK}(\vec{x}, M)$ View M as an element of \mathbb{G}_T , and let $\vec{x} = (x_1, \dots, x_n)$ with $x_i \in \mathbb{Z}_N$. This algorithm chooses random $s, \alpha, \beta \in \mathbb{Z}_N$ and $R_{3,i}, R_{4,i} \in \mathbb{G}_r$ for $i = 1$ to n . It outputs the ciphertext

$$C = \left(C' = M \cdot P^s, C_1 = g_p^s, \left\{ C_{1,i} = H_{1,i}^s \cdot Q^{\alpha \cdot x_i} \cdot R_{3,i}, C_{2,i} = H_{2,i}^s \cdot Q^{\beta \cdot x_i} \cdot R_{4,i} \right\}_{i=1}^n \right).$$

GenKey $_{SK}(\vec{v})$ Let $\vec{v} = (v_1, \dots, v_n)$. This algorithm chooses random $r_{1,i}, r_{2,i} \in \mathbb{Z}_p$ for $i = 1$ to n , random $f_1, f_2 \in \mathbb{Z}_q$, random $R_5 \in \mathbb{G}_r$, and random $Q_6 \in \mathbb{G}_q$. It then outputs

$$SK_{\vec{v}} = \left(K = R_5 \cdot Q_6 \cdot h^{-\gamma} \cdot \prod_{i=1}^n h_{1,i}^{-r_{1,i}} \cdot h_{2,i}^{-r_{2,i}}, \left\{ K_{1,i} = g_p^{r_{1,i}} \cdot g_q^{f_1 \cdot v_i}, K_{2,i} = g_p^{r_{2,i}} \cdot g_q^{f_2 \cdot v_i} \right\}_{i=1}^n \right).$$

Dec $_{SK_{\vec{v}}}(C)$ Let C and $SK_{\vec{v}}$ be as above. The decryption algorithm outputs

$$C' \cdot \hat{e}(C_1, K) \cdot \prod_{i=1}^n \hat{e}(C_{1,i}, K_{1,i}) \cdot \hat{e}(C_{2,i}, K_{2,i}).$$

As we have described it, decryption never returns an error (i.e., even when $\langle \vec{v}, \vec{x} \rangle \neq 0$). We will show below that when $\langle \vec{v}, \vec{x} \rangle \neq 0$ then the output is essentially a random element in the order- q subgroup of \mathbb{G}_T . By restricting the message space to some efficiently-recognizable set of negligible density in this subgroup, we recover the desired semantics by returning an error if the recovered message does not lie in this space.

Correctness. Let C and $SK_{\vec{v}}$ be as above. Then

$$\begin{aligned}
& C' \cdot \hat{e}(C_1, K) \cdot \prod_{i=1}^n \hat{e}(C_{1,i}, K_{1,i}) \cdot \hat{e}(C_{2,i}, K_{2,i}) \\
&= M \cdot P^s \cdot \hat{e} \left(g_p^s, R_5 Q_6 h^{-\gamma} \prod_{i=1}^n h_{1,i}^{-r_{1,i}} h_{2,i}^{-r_{2,i}} \right) \\
&\quad \cdot \prod_{i=1}^n \hat{e} \left(H_{1,i}^s Q^{\alpha \cdot x_i} R_{3,i}, g_p^{r_{1,i}} g_q^{f_{1,i} \cdot v_i} \right) \cdot \hat{e} \left(H_{2,i}^s Q^{\beta \cdot x_i} R_{4,i}, g_p^{r_{2,i}} g_q^{f_{2,i} \cdot v_i} \right) \\
&= M \cdot P^s \cdot \hat{e} \left(g_p^s, h^{-\gamma} \prod_{i=1}^n h_{1,i}^{-r_{1,i}} h_{2,i}^{-r_{2,i}} \right) \cdot \prod_{i=1}^n \hat{e} \left(h_{1,i}^s g_q^{\alpha \cdot x_i}, g_p^{r_{1,i}} g_q^{f_{1,i} \cdot v_i} \right) \cdot \hat{e} \left(h_{2,i}^s g_q^{\beta \cdot x_i}, g_p^{r_{2,i}} g_q^{f_{2,i} \cdot v_i} \right) \\
&= M \cdot P^s \cdot \hat{e}(g_p, h)^{-\gamma s} \cdot \prod_{i=1}^n \hat{e}(g_q, g_q)^{(\alpha f_{1,i} + \beta f_{2,i}) x_i v_i} = M \cdot \hat{e}(g_q, g_q)^{(\alpha f_1 + \beta f_2) \langle \vec{x}, \vec{v} \rangle}.
\end{aligned}$$

If $\langle \vec{x}, \vec{v} \rangle = 0 \pmod N$, then the above evaluates to M . If $\langle \vec{x}, \vec{v} \rangle \neq 0 \pmod N$ there are two cases: if $\langle \vec{x}, \vec{v} \rangle \neq 0 \pmod q$ then the above evaluates to an element whose distribution is statistically close to uniform in the order- q subgroup of \mathbb{G}_T . (Recall that α, β are chosen at random.) It is possible that $\langle \vec{x}, \vec{v} \rangle = 0 \pmod q$, in which case the above always evaluates to M ; however, this reveals a non-trivial factor of N and so an adversary can cause this condition to occur with only negligible probability.

D.1 Proof of Security

Theorem D.1. *If \mathcal{G} satisfies Assumptions 1 and 2 then the scheme described in the previous section is an attribute-hiding predicate encryption scheme.*

We prove that the scheme described in the previous section satisfies Definition 2.2. In proving this, we distinguish two cases: when $M_0 = M_1$ and when $M_0 \neq M_1$. We show that the adversary's probability of success conditioned on the occurrence of each case is negligibly-close to $1/2$.

A proof for the case $M_0 = M_1$ follows *mutatis mutandis* from the proof given in Section 4. Specifically, if $M_0 = M_1 = M$ then the adversary gets no advantage from the extra term $M \cdot P^s$ included in the challenge ciphertext and so the only point to verify is that, throughout the proofs in Sections C.1 and C.2, the simulator can compute the value P^s (so that it can construct the additional element $C' = M \cdot P^s$). This is easy to do if the simulator computes P exactly as in the Setup algorithm, and stores $h^{-\gamma}$. We omit the straightforward details.

Given the above, we concentrate here on proving security under the assumption that $M_0 \neq M_1$. Since we are considering only this case, we will assume the adversary is restricted to requesting keys corresponding to vectors \vec{v} for which $\langle \vec{v}, \vec{x} \rangle \neq 0$ and $\langle \vec{v}, \vec{y} \rangle \neq 0$, where \vec{x}, \vec{y} are the vectors output by the adversary at the outset of the experiment (cf. Definition A.1). We establish the result in this case using a sequence of games, defined as follows.

Game₀: The challenge ciphertext is generated as a proper encryption of M_0 using \vec{x} . That is, we choose random $s, \alpha, \beta \in \mathbb{Z}_N$ and random $\{R_{3,i}, R_{4,i}\} \in \mathbb{G}_r$, and compute the ciphertext as

$$C = \left(C' = M_0 \cdot P^s, \quad C_1 = g_p^s, \quad \left\{ C_{1,i} = H_{1,i}^s Q^{\alpha x_i} R_{3,i}, \quad C_{2,i} = H_{2,i}^s Q^{\beta x_i} R_{4,i} \right\}_{i=1}^n \right).$$

Game₁: We now generate the challenge ciphertext as a proper encryption of a random element of \mathbb{G}_T , but still using \vec{x} . I.e., the ciphertext is formed as above except that C' is chosen uniformly from \mathbb{G}_T .

Game₂: We now generate the $\{C_{2,i}\}$ components as if encryption were done using $\vec{0}$. That is, we choose random $s, \alpha, \beta \in \mathbb{Z}_N$, random $\{R_{3,i}, R_{4,i}\} \in \mathbb{G}_r$, and random $C' \in \mathbb{G}_T$, and compute the ciphertext as

$$C = \left(C', \quad C_1 = g_p^s, \quad \{ C_{1,i} = H_{1,i}^s Q^{\alpha x_i} R_{3,i}, \quad C_{2,i} = H_{2,i}^s R_{4,i} \}_{i=1}^n \right).$$

Note that this exactly parallels **Game₂** in the proof of Theorem C.1.

Game₃: We now generate the $\{C_{2,i}\}$ components using vector \vec{y} . That is, we choose random $s, \alpha, \beta \in \mathbb{Z}_N$, random $\{R_{3,i}, R_{4,i}\} \in \mathbb{G}_r$, and random $C' \in \mathbb{G}_T$, and compute the ciphertext as

$$C = \left(C', \quad C_1 = g_p^s, \quad \{ C_{1,i} = H_{1,i}^s Q^{\alpha x_i} R_{3,i}, \quad C_{2,i} = H_{2,i}^s Q^{\beta y_i} R_{4,i} \}_{i=1}^n \right).$$

Note that this exactly parallels **Game₃** in the proof of Theorem C.1.

Game₄ and **Game₅**: These games are defined symmetrically to **Game₂** and **Game₃**, as in the proof of Theorem C.1. We continue to let C' be a random element of \mathbb{G}_T . Note that **Game₅** corresponds to a proper encryption of a random element of \mathbb{G}_T using \vec{y} .

Game₆: The challenge ciphertext is generated as a proper encryption of M_1 using \vec{y} .

In the next section we prove that, under Assumption 2, **Game₀** and **Game₁** are indistinguishable. Indistinguishability of **Game₁** and **Game₅** follows, as mentioned earlier, *mutatis mutandis* from the proofs in Sections C.1 and C.2. The proof that **Game₅** and **Game₆** are indistinguishable is symmetric to the proof that **Game₀** and **Game₁** are indistinguishable, and is therefore omitted.

D.1.1 Indistinguishability of **Game₀** and **Game₁**

Fix an adversary \mathcal{A} . We describe a simulator who is given $(N = pqr, \mathbb{G}, \mathbb{G}_T, \hat{e})$ along with the elements $g_p, g_q, g_r, h, g_p^s, h^s Q_1, g_p^\gamma Q_2, \hat{e}(g_p, h)^\gamma$, and an element T which is either equal to $\hat{e}(g_p, h)^\gamma$ or is uniformly distributed in \mathbb{G}_T . Note that the simulator is now able to sample uniformly from \mathbb{G}_q and \mathbb{G}_r using g_q and g_r , respectively. In particular, the simulator can sample uniformly from $\mathbb{G}_{qr} = \mathbb{G}_q \times \mathbb{G}_r$. The simulator interacts with \mathcal{A} as we now describe.

Public parameters. The simulator begins by giving N to \mathcal{A} , who outputs vectors \vec{x}, \vec{y} . The simulator chooses random $\{w_{1,i}, w_{2,i}\} \in \mathbb{Z}_N$ and random $\{R_{1,i}, R_{2,i}\}, R_0 \in \mathbb{G}_r$, includes $(N, \mathbb{G}, \mathbb{G}_T, \hat{e})$ in the public parameters, and sets the remainder of the parameters as follows:

$$PK = (g_p, g_r, Q = g_q R_0, P = \hat{e}(g_p, h)^\gamma, \quad \{ H_{1,i} = h^{x_i} g_p^{w_{1,i}} R_{1,i}, \quad H_{2,i} = h^{x_i} g_p^{w_{2,i}} R_{2,i} \}_{i=1}^n).$$

The simulator is implicitly setting $h_{1,i} = h^{x_i} g_p^{w_{1,i}}$ and $h_{2,i} = h^{x_i} g_p^{w_{2,i}}$. Note that PK has the appropriate distribution.

Key derivation. The adversary \mathcal{A} may request secret keys corresponding to different vectors \vec{v} , as long as $\langle \vec{v}, \vec{x} \rangle \neq 0$ (we do not use the fact that $\langle \vec{v}, \vec{y} \rangle \neq 0$ here). We now describe how the simulator prepares the secret key corresponding to any such vector.

Say the adversary requests the secret key for vector \vec{v} , and let $k = 1/2 \cdot \langle \vec{x}, \vec{v} \rangle \bmod N$. (If $\gcd(\langle \vec{x}, \vec{v} \rangle, N) \neq 1$) then the adversary has factored N ; this occurs with negligible probability.) The simulator first chooses random $f'_1, f'_2, \{r'_{1,i}, r'_{2,i}\} \in \mathbb{Z}_N$. Next, for all i it computes:

$$\begin{aligned} K_{1,i} &= (g_p^\gamma Q_2)^{-kv_i} \cdot g_q^{f'_1 v_i} \cdot g_p^{r'_{1,i}} \\ &= g_p^{-kv_i \gamma + r'_{1,i}} \cdot g_q^{(f'_1 - kc) \cdot v_i} \end{aligned}$$

(where we set $c = \log_{g_q} Q_2$), and

$$\begin{aligned} K_{2,i} &= (g_p^\gamma Q_2)^{-kv_i} \cdot g_q^{f'_2 v_i} \cdot g_p^{r'_{2,i}} \\ &= g_p^{-kv_i \gamma + r'_{2,i}} \cdot g_q^{(f'_2 - kc) \cdot v_i}. \end{aligned}$$

The simulator then chooses random $QR \in \mathbb{G}_{qr}$ and computes:

$$K = QR \cdot \prod_{i=1}^n \left((g_p^{w_{1,i}} h^{x_i})^{-r'_{1,i}} \cdot (g_p^\gamma Q_2)^{kv_i w_{1,i}} \right) \cdot \left((g_p^{w_{2,i}} h^{x_i})^{-r'_{2,i}} \cdot (g_p^\gamma Q_2)^{kv_i w_{2,i}} \right).$$

Finally, the simulator hands the adversary $SK_{\vec{v}} = (K, \{K_{1,i}, K_{2,i}\}_{i=1}^n)$ as the key.

To see that this key has the correct distribution, note that by construction of the $\{K_{1,i}, K_{2,i}\}$ the simulator is implicitly setting $f_1 = f'_1 - kc$ and, for all i , $r_{1,i} = -k\gamma v_i + r'_{1,i}$ (and analogously for f_2 and the $\{r_{2,i}\}$). These values are all uniformly and independently distributed in \mathbb{Z}_N . Next, note that

$$\begin{aligned} \prod_{i=1}^n (g_p^{w_{1,i}} h^{x_i})^{-r'_{1,i}} \cdot (g_p^\gamma)^{kv_i w_{1,i}} &= \prod_{i=1}^n g_p^{-w_{1,i} r'_{1,i} + k\gamma v_i w_{1,i}} \cdot h^{-x_i r'_{1,i}} \\ &= \prod_{i=1}^n g_p^{-w_{1,i} \cdot (r_{1,i} + k\gamma v_i) + k\gamma v_i w_{1,i}} \cdot h^{-x_i \cdot (r_{1,i} + k\gamma v_i)} \\ &= \prod_{i=1}^n (h^{x_i} g_p^{w_{1,i}})^{-r_{1,i}} \cdot h^{-\gamma kv_i x_i} = h^{-\gamma/2} \cdot \prod_{i=1}^n h_{1,i}^{-r_{1,i}}, \end{aligned}$$

using the fact that $\langle \vec{v}, \vec{x} \rangle = 1/2k \bmod N$. Thus, looking at K_p (the projection of K in \mathbb{G}_p) we see that

$$\begin{aligned} K_p &= \prod_{i=1}^n \left((g_p^{w_{1,i}} h^{x_i})^{-r'_{1,i}} \cdot (g_p^\gamma)^{kv_i w_{1,i}} \right) \cdot \left((g_p^{w_{2,i}} h^{x_i})^{-r'_{2,i}} \cdot (g_p^\gamma)^{kv_i w_{2,i}} \right) \\ &= h^{-\gamma} \cdot \prod_{i=1}^n h_{1,i}^{-r_{1,i}} \cdot h_{2,i}^{-r_{2,i}}, \end{aligned}$$

and so K_p (and hence K) is distributed appropriately.

The challenge ciphertext. The challenge ciphertext is generated as follows. The simulator chooses random $\{R_{7,i}, R_{8,i}\} \in \mathbb{G}_r$ and $Q'_1 \in \mathbb{G}_q$, sets $C' = M_0 \cdot T$, sets $C_1 = g_p^s$, and computes:

$$\begin{aligned} C_{1,i} &= (g_p^s)^{w_{1,i}} \cdot (h^s Q_1)^{x_i} \cdot R_{7,i} \\ &= (h^{x_i} g_p^{w_{1,i}})^s \cdot Q_1^{x_i} \cdot R_{7,i} \\ C_{2,i} &= (g_p^s)^{w_{2,i}} \cdot (h^s Q_1)^{x_i} \cdot (Q'_1)^{x_i} \cdot R_{8,i} \\ &= (h^{x_i} g_p^{w_{2,i}})^s \cdot (Q_1 Q'_1)^{x_i} \cdot R_{8,i}. \end{aligned}$$

Analysis. By examining the projections of the components of the challenge ciphertext in the groups \mathbb{G}_p , \mathbb{G}_q , and \mathbb{G}_r , it can be verified that when $T = \hat{e}(g_p, h)^{\gamma s}$ the challenge ciphertext is distributed exactly as in Game_0 , whereas if T is chosen uniformly from \mathbb{G}_T the challenge ciphertext is distributed exactly as in Game_1 . We conclude that, under Assumption 2, these two games are indistinguishable. ■