

Access Control in Publicly Verifiable Outsourced Computation

James Alderman, Christian Janson, Carlos Cid, and Jason Crampton

Information Security Group, Royal Holloway, University of London
Egham, Surrey, TW20 0EX, United Kingdom
{James.Alderman.2011, Christian.Janson.2012}@live.rhul.ac.uk
{Carlos.Cid, Jason.Crampton}@rhul.ac.uk

Abstract

Publicly Verifiable Outsourced Computation (PVC) allows devices with restricted resources to delegate expensive computations to more powerful external servers, and to verify the correctness of results. Whilst highly beneficial in many situations, this increases the visibility and availability of potentially sensitive data, so we may wish to limit the sets of entities that can view input data and results. Additionally, it is highly unlikely that all users have identical and uncontrolled access to all functionality within an organization. Thus there is a need for access control mechanisms in PVC environments.

In this work, we define a new framework for Publicly Verifiable Outsourced Computation with Access Control (PVC-AC). We formally define algorithms to provide different PVC functionality for each entity within a large outsourced computation environment, and discuss the forms of access control policies that are applicable, and necessary, in such environments, as well as formally modelling the resulting security properties. Finally, we give an example instantiation that (in a black-box and generic fashion) combines existing PVC schemes with symmetric Key Assignment Schemes to cryptographically enforce the policies of interest.

1 Introduction

The increasing use of mobile devices as general computing devices, “big data” and cloud computing results in a growing discrepancy between the resources of end-users and those required for computational tasks. As a result, client devices need to be able to efficiently delegate computations to a server and verify the correctness of results. Consider, for example, a company operating a “bring your own device” policy for smartphones and tablets that may be unable to perform complex computations locally. Instead, the company subscribes to an external (untrusted) company providing software-as-a-service (SaaS), but needs to ensure that results remain correct. Alternatively, in the context of battlefield communications, a squadron of soldiers may be deployed with light-weight devices to gather data from their surroundings and send it to regional servers for analysis before receiving tactical commands. Those servers may not be fully trusted, e.g. if part of a coalition network, so soldiers must be assured that the command has been computed correctly. This setting is known as *Verifiable Outsourced Computation* (VC) [16, 20, 10].

Recent work has proposed extensions to VC where there exist large pools of delegators, known as Publicly Verifiable Outsourced Computation (PVC) [20], and servers [2]. As with any multi-user setting, we may wish to control access to resources. In this paper, we show that not only is this setting of multi-user VC well-suited to the cryptographic enforcement of access control policies, but that such policies fulfil a natural and vital role in protecting outsourced computations. Specifically in the setting of multi-user VC, we may wish to (i) restrict the

computations that may be outsourced by delegators; and (ii) restrict the computations a server may perform. The first need stems from separation of duties and the observation that, within an organization, it is extremely unlikely that all users have equal, uncontrolled access to all functionality. We may restrict the set of delegators that may outsource a computation to those that are authorized internally to compute it (if given sufficient resources). The second requirement arises, for example, from servers being selected from the pool for a particular job without the delegator necessarily having prior knowledge. Thus delegators may not authenticate the server beforehand (in contrast with prior schemes where a single server was chosen with which to set up a VC system) and have less control over which servers may operate on their data. The sensitivity of the data or other requirements, such as the physical location or resources of the server, may limit the servers that should be permitted to perform the computation.

Some VC settings distinguish between delegators and verifiers [2]. Delegators (or distinguished verifiers) may learn the result of the computation, whereas standard verifiers may only confirm that the result was computed correctly. Again, in a multi-user VC setting, we may wish to restrict the users that: (i) verify the result; and (ii) learn the result. When operating on sensitive data, this second restriction ensures that read access to the generated results is limited to those that satisfy the access control policy e.g. only entities that may read the input data may read the output.

A third motivation for access control in the VC setting is that computational services may be charged for (e.g. in subscription-based utility computing [17, 21]) and that service providers may offer different levels of service to different clients (e.g. different levels may provide access to different functions or computational resources). We must ensure that only valid subscribers may access each tier of service.

In many multi-user settings for access control to stored data [7, 22], servers enforce access control policies by authenticating users and granting or denying access based on *access control lists* or *capability lists*. This reference monitor is not appropriate in the multi-user VC setting since the servers are assumed to be untrusted and may have a vested interest in violating the policies. We instead use a cryptographic enforcement mechanism for access control policies where cryptographic keys are used to protect objects and restrict access — the access control mechanism reduces to the appropriate distribution of keys to authorized entities. We use the trusted Key Distribution Center (KDC) introduced in the Revocable Publicly Verifiable Outsourced Computation model (RPVC) [2] and extend its duties to instantiate the access control mechanism. In RPVC, the KDC issued keys that enable the delegation and evaluation of particular functions; we additionally require it to issue keys appropriate to the access control policy that enables read and write access to certain components of the PVC system. For example, input data for particular functions may be protected such that only authorized servers may read the data and hence perform the computation.

Note that in an RPVC scheme, the KDC implicitly provides some access control in that servers are certified to perform specific functions through the generation of evaluation keys. However, no access control is applied to delegators — *any* entity can outsource an evaluation of *any* function for which the KDC has published delegation information (essentially due to the use of asymmetric cryptographic primitives). In particular, a delegator may request a computation that the delegator itself is not authorized to perform.

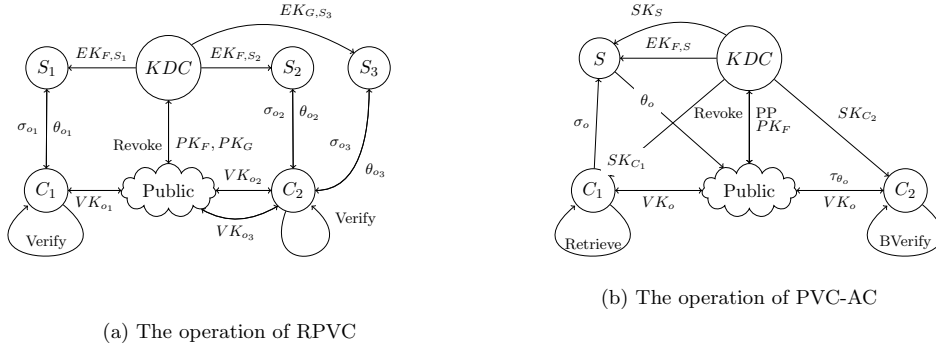
Cryptographic enforcement mechanisms are particularly appropriate when the objects and policies are relatively static (such that additional keys need not be generated and objects need not be re-encrypted). In the context of VC, we may assume that the set of functions that may be evaluated is fixed (a given VC construction can implement a specified family of functions) and that the input data to each function is also static (limited to the set of ‘valid’ inputs to that function). Thus, the set of objects (function evaluations in VC) is static, and policies will

primarily be specified in terms of these computations. Thus multi-user VC is a very natural setting in which to use cryptographic access control. However, VC leads to a somewhat novel application of these mechanisms as we will see in Section 3.

Related work Clear et al. [11] considered policies over delegators only and in a non-verifiable, multi-input outsourced computation setting using homomorphic Ciphertext-Policy ABE and fully homomorphic encryption. In independent and concurrent work, Xu et al. [23] also addressed the necessity for access control in the setting of verifiable computation, but limited their scope to non-public verifiable computation (i.e. not the full multi-user setting) enforcing access control on delegators only. We believe that the PVC setting also necessitates that delegators may specify access control requirements on those *servers* that may be selected for a given computation and, especially, limits must be placed on *verifiers* that may learn the output. Xu et al. discuss their notion purely in terms of using CP-ABE as the enforcement mechanism and did not discuss the form of the policies; in contrast, we discuss in detail the types of policies that may be of interest in these settings and present these in terms of generic graph-based access control policies. These may be enforced by a variety of enforcement mechanisms, including symmetric KASs as used here (which may well be more efficient than the pairing-based CP-ABE approach). We believe that we present a more generic treatment which, importantly, extends to the multi-user PVC setting and incorporates a more natural view of hierarchical access control, as well as additional security notions.

Structure of the paper We begin with some preliminary background material before, in Section 3, discussing example access control policies that are relevant to the setting of multi-user VC. We formulate policies in a generic fashion (irrespective of the cryptographic enforcement mechanism employed) in terms of authorization labels and graph-based policies over entities and computations. Then, in Section 4, we extend prior literature [16, 20, 2] to formally define a framework for PVC-AC. This results in novel security notions that we introduce in Section 5. Finally, in Section 6, we provide an example construction that provably meets the stated security goals. Our construction generically extends any RPVC scheme to use a symmetric Key Assignment Scheme (KAS). When instantiated with a PVC scheme built from Key-policy Attribute-based Encryption [2, 20], this provides a pragmatic blend of symmetric and asymmetric primitives – the symmetric KAS enables the efficient derivation of keys, each associated to a security label, while policies are specified in terms of these labels. A trusted authority on entities, the KDC, issues keys to each entity corresponding to the highest security label for which they are authorized, and they may derive “lower” keys as required. These keys may cryptographically protect input values and verification keys such that only entities satisfying the associated policies may operate on them. Unauthorised servers or eavesdroppers may not even learn the input values, thus exposure of data is limited to explicitly authorised (trusted) entities. Moreover, this restriction enforces access control policies on the part of the delegator device. To encrypt the inputs with a symmetric key associated with the policy, the delegator must be able to derive this key in the first place. Thus, the delegator’s security clearance must be at least the classification of the function in order to derive this key from that issued by the KDC.

We use the following notation. We write $y \leftarrow A(\cdot)$ to denote running a probabilistic algorithm A and assigning the result to y . We use PPT to denote probabilistic polynomial-time and say that $\text{negl}(\cdot)$ is a negligible function of its input. $\mathcal{A}^{\mathcal{O}}$ is used to denote the adversary \mathcal{A} being provided with oracle access.



(a) The operation of RPVC

(b) The operation of PVC-AC

Figure 1: The operation of RPVC and PVC-AC

2 Background

2.1 Revocable PVC

Non-interactive verifiable computation (VC), introduced by Gennaro et al. [16], is a protocol between two PPT parties: a *client* C and a *server* S . A successful protocol run results in the provably correct computation of $F(x)$ by the server for an input x given by the client. More specifically [16]:

1. C computes evaluation information EK_F given to S to enable the computation of F (pre-processing)
2. C sends the encoded input σ_o to S (input preparation)
3. S computes $y = F(x)$ using EK_F and σ_o and returns an encoded output θ_o (output computation)
4. C checks whether θ_o encodes $F(x)$ (verification)

KeyGen may be computationally expensive but the remaining operations should be efficient for the client. The cost of setup is amortized over multiple computations of F .

Parno et al. [20] introduced Publicly Verifiable Outsourced Computation (PVC) to allow multiple clients to delegate computations. Alderman et al. [2] extended PVC to consider a wider framework, Revocable PVC (RPVC). Here large pools of clients and servers exist and jobs may be submitted to the pool; an appropriate server is selected by a system-dependent mechanism. This model introduces a trusted entity known as a Key Distribution Center (KDC) that performs costly set-up operations and issues appropriate keys. RPVC comprises the algorithms Setup, FInit, Register, Certify, ProbGen, Compute, BVerif, Retrieve and Revoke illustrated in Figure 1a and corresponding to the following:

1. The KDC generates public parameters, issues personalised secret keys, and evaluation keys to servers and publishes function delegation information.
2. To outsource the evaluation of $F(x)$, a delegator C , sends an encoded input σ_o to a server S , and publishes a verification token and output retrieval key for the computation.
3. S uses σ_o and an evaluation key for F to produce an encoded output (sent to C , a manager, or published depending on the system architecture).
4. Any entity can use the verification token to verify correctness without learning the value of $F(x)$ unless in possession of a retrieval key (blind verification). If S cheated they may report S to the KDC for revocation.

5. If blind verification was successful, a party possessing the output retrieval key can recover $F(x)$.
6. The KDC may revoke a cheating server to prevent it from performing computations (and hence from receiving any reward for future work).

2.2 Cryptographic Enforcement of Access Control Policies

2.2.1 Graph-based Access Control

A *partially ordered set* (*poset*) is a set L equipped with a reflexive, anti-symmetric and transitive binary relation \leq . We may write $x < y$ if $x \leq y$ and $x \neq y$, and write $y \geq x$ if $x \leq y$. We say that x *covers* y , written $y \triangleleft x$, if $y < x$ and no z exists in L such that $y < z < x$. The *Hasse Diagram* of a poset (L, \leq) is the directed acyclic graph $H = (L, \triangleleft)$ where vertices are labelled by the elements of L and an edge connects vertex v to w if and only if $w \triangleleft v$. Let U be a set of entities, O be a set of resources to be protected, and (L, \leq) be a poset of security labels. Let $\lambda : U \cup O \rightarrow L$ be a labelling function assigning a security label to each entity and object. The tuple (L, \leq, U, O, λ) then denotes an *information flow policy* which can be represented by the H . Henceforth we refer to such policies as *graph-based access control policies*. The policy requires that information flow from objects to entities preserves the partial ordering relation — an entity $u \in U$ may read an object $o \in O$ if and only if $\lambda(u) \geq \lambda(o)$ i.e. there exists a directed path from $\lambda(u)$ to $\lambda(o)$ in H . Note that this statement is the *simple security property* of the Bell-LaPadula security model [4]. Thus an entity assigned clearance label x is prevented from accessing objects classified with label y if $y > x$.

2.2.2 Key Assignment Schemes

A *Key Assignment Scheme* (KAS) [1] provides a generic, cryptographic enforcement mechanism for such policies where a unique cryptographic key is associated to each node (representing a security label) in H . A KAS eases the problem of key distribution by allowing a trusted center to distribute a single key to each entity, who may combine this key with public information to derive any additional keys for which the user is authorized. More formally, for an information flow policy (L, \leq) [13]:

- $\text{MakeKeys}(1^\kappa, (L, \leq)) \rightarrow \kappa_L$: returns a labelled set of encryption keys $\{\kappa_x : x \in L\}$, denoted κ_L .
- $\text{MakeSecrets}(1^\kappa, (L, \leq)) \rightarrow \omega_L$: returns a labelled set of secret values $\{\omega_x : x \in L\}$, denoted ω_L .
- $\text{MakePublicData}(1^\kappa, (L, \leq)) \rightarrow \text{Pub}_{(L, \leq)}$: returns a set of data $\text{Pub}_{(L, \leq)}$ published by the trusted authority
- $\text{GetKey}(x, y, \omega_x, \text{Pub}_{(L, \leq)}) \rightarrow \kappa_y$: takes two nodes, $x, y \in L$, the secret information for x , ω_x , and the public information, $\text{Pub}_{(L, \leq)}$ and returns κ_y if $y \leq x$.

A well-known KAS construction (known as an *iterative key encrypting* (IKE) KAS [13]) publishes encrypted keys. In particular, for each $y < x$, such that no z exists with $y < z < x$, $\text{Encrypt}_{\kappa_x}(\kappa_y)$ is published. Then for any $x > y$, the key for each node on a path from x to y can be derived (in an iterative fashion) if κ_x is known. *Key indistinguishability* (KI) [3] is a crucial security property if derived keys are used in other protocols: given a set of keys $\kappa_{x_1}, \dots, \kappa_{x_n}$, an adversary should not be able to distinguish between the key for a challenge node x^* (not a descendent of any x_i) and a randomly chosen key. Recently, Freire et al. [15] suggested a new

Game 1 $\text{Exp}_{\mathcal{A}}^{\text{S-KI}} [1^\kappa, (L, \leq)]:$

- 1: $b \xleftarrow{\$} \{0, 1\}$
- 2: $\kappa_L \leftarrow \text{MakeKeys}(1^\kappa, (L, \leq))$
- 3: $\omega_L \leftarrow \text{MakeSecrets}(1^\kappa, (L, \leq))$
- 4: $\text{Pub}_{(L, \leq)} \leftarrow \text{MakePublicData}(1^\kappa, (L, \leq))$
- 5: $Q = \epsilon$
- 6: $v^* \leftarrow \mathcal{A}^{\mathcal{O}^{\text{Corrupt}(\cdot, \perp)}}(\text{Pub}_{(L, \leq)})$
- 7: **for** all $v_i \in Q$ **do**
- 8: **if** $v^* \leq v_i$ **then return** 0
- 9: **if** $b = 0$ **then** $k^* = k_{v^*}$
- 10: **else** $k^* \xleftarrow{\$} \mathcal{K}$
- 11: $b' \leftarrow \mathcal{A}^{\mathcal{O}^{\text{Corrupt}(\cdot, v^*)}}(k^*)$
- 12: **if** $b' = b$ **then return** 1
- 13: **else return** 0

Oracle Query 1 $\mathcal{O}^{\text{Corrupt}}(v_i, v^*):$

- 1: **if** $v^* \not\leq v_i$ **then**
- 2: $Q = Q \cup v_i$
- 3: **return** $(\kappa_{v_i}, \omega_{v_i})$
- 4: **else if** $v_i > v^*$ **then**
- 5: **return** κ_{v_i}
- 6: **else**
- 7: **return** \perp

KI security definition for hierarchial KASs called *Strong Key Indistinguishability* (S-KI) which strengthens the definition in [3] to capture a wider range of realistic attacks. S-KI requires that the adversary \mathcal{A} is not able to gain any information about a key which it should not have access to *even if* \mathcal{A} has access to keys associated to all other classes which are predecessors of the target class in the hierarchy. In Game 1 we formally capture the notion of S-KI for an adaptive adversary.

Definition 1. *The advantage of an adversary \mathcal{A} running in probabilistic polynomial time (PPT) is defined as:*

$$\text{Adv}_{\mathcal{A}, \mathcal{KAS}}^{\text{S-KI}}(1^\kappa, (L, \leq)) = \Pr[\text{S-KI}_{\mathcal{KAS}}^{\mathcal{A}}(1^\kappa, (L, \leq))] - \frac{1}{2}.$$

A KAS is secure in the sense of strong key indistinguishability against adaptive adversaries (S-KI) if for all PPT adversaries \mathcal{A} , $\text{Adv}_{\mathcal{A}, \mathcal{KAS}}^{\text{S-KI}}(1^\kappa, (L, \leq)) \leq \text{negl}(\kappa)$.

2.3 Authenticated Encryption

In this paper, we use two forms of encryption scheme: a public-key Key-Policy Attribute-based Encryption (KP-ABE) [18, 19] scheme is used as the foundation of the Publicly Verifiable Outsourced Computation functionality, as introduced by Parno et al. [20] – decryption of a ciphertext succeeds only if a function is satisfied by an input; the second form is a symmetric authenticated encryption scheme to preserve confidentiality of messages from unauthorized readers and to achieve data origin authentication. We refer the reader to the cited literature for an introduction to KP-ABE as we use this only in a black-box manner as a component of an existing RPVC scheme. In this section, we give some brief background on authenticated encryption in the symmetric setting and introduce the security notions we shall require.

Game 2 $\text{Exp}_A^{\text{IND-CPA}} [1^\kappa]$:

- 1: $b \xleftarrow{\$} \{0, 1\}$
 - 2: $SK^* \leftarrow \text{KeyGen}(1^\kappa)$
 - 3: $b' \leftarrow \mathcal{A}^{\mathcal{O}^{\text{LoR}}(\cdot, SK^*)}(1^\kappa)$
 - 4: **return** $b' = b$
-

Oracle Query 2 $\mathcal{O}^{\text{LoR}}(m_0, m_1, SK^*)$:

- 1: **return** $\text{Encrypt}(m_b, SK^*)$
-

Game 3 $\text{Exp}_A^{\text{IND-CCA}} [1^\kappa]$:

- 1: $b \xleftarrow{\$} \{0, 1\}$
 - 2: $L = \epsilon$
 - 3: $SK^* \leftarrow \text{KeyGen}(1^\kappa)$
 - 4: $b' \leftarrow \mathcal{A}^{\mathcal{O}^{\text{LoR}}(\cdot, SK^*), \text{Decrypt}(\cdot, SK^*)}(1^\kappa)$
 - 5: **return** $b' = b$
-

Oracle Query 3 $\mathcal{O}^{\text{LoR}}(m_0, m_1, SK^*)$:

- 1: $CT \leftarrow \text{Encrypt}(m_b, SK^*)$
 - 2: $L = L \cup CT$
 - 3: **return** CT
-

2.3.1 Symmetric Encryption

A symmetric encryption scheme [5] \mathcal{SE} comprises three algorithms: KeyGen , Encrypt and Decrypt . KeyGen is a randomized algorithm that takes a security parameter κ and returns a secret key SK ; Encrypt is a randomized¹ algorithm taking as input a message m and the secret key, and returning a ciphertext CT ; finally, the Decrypt algorithm is deterministic and takes a ciphertext and the secret key and returns the corresponding plaintext m or a failure symbol \perp . For correctness, we require that, for all $SK \leftarrow \text{KeyGen}(1^\kappa)$ and messages m in the messagespace, $m \leftarrow \text{Decrypt}(\text{Encrypt}(m, SK), SK)$. A symmetric authenticated encryption scheme is syntactically identical but considers integrity as well as privacy of messages, and \perp is returned if the ciphertext is deemed ‘inauthentic’.

Two standard notions of privacy for symmetric encryption schemes are *Indistinguishability under Chosen Plaintext Attack* (IND-CPA) and *Indistinguishability under Chosen Ciphertext Attack* (IND-CCA). These are given in Games 2 and 3 respectively. In each, the challenger chooses a random bit b and generates a key SK^* . The adversary is given the security parameter (so can generate other keys itself) and oracle access to LoR which takes two messages of equal length and returns the encryption of message m_b . The adversary must guess the value of b . In the IND-CCA notion, the adversary is also given a decryption oracle for the challenge key but may not query a ciphertext that was output from the LoR oracle to avoid a trivial win.

Definition 2. *The advantage of an adversary \mathcal{A} running in probabilistic polynomial time (PPT) for $\mathbf{X} \in \{\text{IND-CPA}, \text{IND-CCA}\}$ is defined as:*

$$\text{Adv}_{\mathcal{A}, \mathcal{SE}}^{\mathbf{X}}(1^\kappa) = \Pr[\mathbf{X}_{\mathcal{SE}}^{\mathcal{A}}(1^\kappa)] - \frac{1}{2}.$$

A symmetric encryption scheme is secure in the sense of \mathbf{X} if for all PPT adversaries \mathcal{A} , $\text{Adv}_{\mathcal{A}, \mathcal{SE}}^{\mathbf{X}}(1^\kappa) \leq \text{negl}(\kappa)$.

Bellare and Namprempre [5] considered two notions of integrity for symmetric encryption schemes: *Integrity of Plaintexts* (INT-PTXT) and *Integrity of Ciphertexts* (INT-CTXT) under chosen message attacks. INT-PTXT requires it to be hard to create a ciphertext which

¹Encrypt could also be a stateful algorithm that updates and maintains a state between calls.

Oracle Query 4 $\mathcal{O}^{\text{Decrypt}(CT, SK^*)}$:

- 1: **if** $CT \in L$ **then return** \perp
 - 2: **return** $\text{Decrypt}(CT, SK^*)$
-

Game 4 $\text{Exp}_{\mathcal{A}}^{\text{INT-PTXT}}[1^\kappa]$:

- 1: $L = \epsilon$
 - 2: $SK^* \leftarrow \text{KeyGen}(1^\kappa)$
 - 3: $CT^* \leftarrow \mathcal{A}^{\mathcal{O}^{\text{Encrypt}(\cdot, SK^*)}, \text{Ver}(\cdot)}(1^\kappa)$
 - 4: **return** $((\text{Ver}(CT^*) = 1) \wedge (\text{Decrypt}(CT^*, SK^*) \notin L))$
-

Oracle Query 5 $\mathcal{O}^{\text{Encrypt}(m, SK^*)}$:

- 1: $CT \leftarrow \text{Encrypt}(m, SK^*)$
 - 2: $L = L \cup m$
 - 3: **return** CT
-

Oracle Query 6 $\mathcal{O}^{\text{Ver}(CT)}$:

- 1: **if** $\perp \neq \text{Decrypt}(CT, SK^*)$ **then return** 1
 - 2: **else return** 0
-

decrypts to a message never encrypted by a legitimate sender, while INT-CTXT requires it to be hard to create a ciphertext that was not previously generated by a legitimate sender (regardless of the underlying plaintext). We will only require the first notion here, which is given in Game 4. The adversary is given both an encryption oracle and a verification oracle which returns 1 if the queried ciphertext is not deemed inauthentic. The adversary wins if it can produce a ciphertext that is deemed authentic and that decrypts to a message not previously queried to the encryption oracle.

Definition 3. *The advantage of an adversary \mathcal{A} running in probabilistic polynomial time (PPT) is defined as:*

$$\text{Adv}_{\mathcal{A}, \mathcal{SE}}^{\text{INT-PTXT}}(1^\kappa) = \Pr[\text{INT-PTXT}_{\mathcal{SE}}^{\mathcal{A}}(1^\kappa)].$$

A symmetric encryption scheme is secure in the sense of integrity of plaintexts if for all PPT adversaries \mathcal{A} , $\text{Adv}_{\mathcal{A}, \mathcal{SE}}^{\text{INT-PTXT}}(1^\kappa) \leq \text{negl}(\kappa)$.

In this work, we will require an authenticated symmetric encryption scheme that is secure in the sense of $\text{IND-CPA} \wedge \text{INT-PTXT}$. As noted by Bellare and Namprepre [5], such a scheme can be constructed from an IND-CPA symmetric encryption scheme and a weakly or strongly unforgeable MAC using the generic composition techniques of *MAC-then-Encrypt* or *Encrypt-then-MAC*. As this latter composition is shown to be secure for all security choices of the encryption and MAC schemes and is more standard, we will adopt this notion in this paper. Thus, let $\mathcal{SE} = (\text{KeyGen}, \text{Encrypt}, \text{Decrypt})$ be a symmetric authenticated encryption scheme constructed from a symmetric encryption scheme $\mathcal{SE}' = (\text{KeyGen}', \text{Encrypt}', \text{Decrypt}')$ and a MAC $\mathcal{MAC} = (\text{KeyGen}'', \text{Tag}, \text{Verify})$. KeyGen outputs two keys: SK_E corresponding to KeyGen' and SK_M relating to KeyGen''². The encryption operation is defined as $\text{Encrypt}(m, SK) = C \parallel \text{Tag}(C, SK_M)$ for $C \leftarrow \text{Encrypt}'(m, SK_E)$.

3 Policies

We now discuss the type of access control policies that are relevant in a multi-user verifiable outsourced computation environment. We consider graph-based policies where “objects” to be

²Note that in our setting where keys are derived from a KAS, the KAS could simply output strings that can be split into SK_E and SK_M .

protected are not data files, as in traditional access control policies, but outsourced computations and their results. The “user” population comprises the sets of delegators \mathcal{C} , computational servers \mathcal{S} , and verifiers \mathcal{V} . We are interested in specifying and enforcing policies that restrict: (i) which functions a delegator may delegate; (ii) which functions a server may evaluate; (iii) which outputs a verifier may read. As noted in the introduction, we distinguish between entities that may learn a computational result and those that may only verify correctness (blind verify). In this work, we assume that *any* entity may perform the blind verification step but that only a restricted subset of verifiers should be able to retrieve the actual output value³.

There has been considerable research in recent years on the cryptographic enforcement of access control policies [13, 14, 1]. Informally, access is regulated in a distributed fashion by issuing appropriate cryptographic keys to authorized subjects, rather than a centralized reference monitor mediating all attempts to access protected resources. Cryptographic access control generally focuses on read access and is regarded as being particularly appropriate when the protected content is read often but written rarely (since this would require re-encryption). In the context of VC, we will use cryptographic access control in somewhat unusual ways. Rather than storing static encrypted data and ensuring that only authorized users hold the relevant decryption key, we will encrypt dynamic messages within a protocol execution. In particular, to enforce policies restricting the computations that may be outsourced, a delegator must use an appropriate key to encrypt input data. Without the appropriate encryption, the input will be discarded by the server. The enforcement of policies for performing computations is achieved by distributing keys to servers that can be used to decrypt encrypted inputs. Without decryption, the server will be unable to read the input data and evaluate the function. The enforcement of (read) policies on outputs uses cryptographic access control in a more conventional fashion; results are published and protected via encryption with an appropriate key.

Cryptographic access control has been particularly widely studied in the context of information flow and graph-based access control policies (see Section 2.2). We restrict our focus to graph-based policies as these have been shown to encompass many notions of access control that are desirable in practice including information flow policies, role-based access control and attribute-based access control [12]. Recall that we define “user” population as the sets of delegators \mathcal{C} , computational servers \mathcal{S} , and verifiers \mathcal{V} . We define a security function $\lambda : \mathcal{C} \cup \mathcal{S} \cup \mathcal{V} \cup \mathcal{O} \rightarrow L$ where \mathcal{O} is the family of computations that may be outsourced and (L, \leq) is a poset of security labels. This function assigns a label from L , representing the security classification, to each delegator, server and computation in the PVC-AC system.

We first consider, in Section 3.1, the form of policies that restrict the computations a delegator may outsource and that a server may compute. We begin by examining simple policies over the choice of functions, before considering more fine-grained policies over sets of input values and other security posets. Then, in Section 3.3, we discuss policies that restrict the computation results a verifier may learn. Although we consider a variety of policies capturing a range of conditions on entities and other contextual information, this amounts to altering the policy poset and the labels assigned to each entity and computation; the definition in Section 4 and the instantiation in Section 6 is generic and encompasses any of the following forms of policy. We differentiate between computation policies over delegators and servers, denoted $\lambda^C(\cdot)$, and verification policies over delegators and verifiers, denoted $\lambda^V(\cdot)$. For ease of notation, we use λ to denote λ^C in Section 3.1, and to denote λ^V in Section 3.3.

³The same techniques that we present could also be applied to protect the (blind) verification keys to restrict the entities that may validate correctness of a computation.




Figure 2: Example Posets of Security Labels

3.1 Delegation and Computation Policies

3.1.1 Policies over Functions.

We begin by considering a simple case where policies are formulated purely in terms of the functions being computed. In simple terms, we associate each delegator C and server S with a set of functions $\lambda(C) \subseteq \mathcal{F}$ and $\lambda(S) \subseteq \mathcal{F}$ respectively. We define a *correctness criterion* that states that C should be able to prepare inputs for all functions $F \in \lambda(C)$ (and similarly for S) i.e. entities should be able to perform all operations that they are authorized for. The *security criterion* requires $\lambda(C), \lambda(S) \geq \lambda(F)$ in order to delegate or compute for F respectively, and that a set of unauthorized entities cannot collude to perform an operation that any of them couldn't perform alone.

More formally, we define the set of security labels L to be $2^{\mathcal{F}}$ (the power set of all considered functions). Then, $\lambda(C) \subseteq \mathcal{F}$ defines the set of functions that a delegator C may outsource an evaluation of, $\lambda(S) \subseteq \mathcal{F}$ denotes the functions a server S may compute, and $\lambda(o) = \{F\}$ labels the computation of $F \in \mathcal{F}$. Then, for any $x, y \in L$ we define an order relation $<$ such that $x < y$ if and only if $x \in \mathcal{F}, y \subseteq \mathcal{F}$ and $x \in y$. The corresponding Hasse diagram with $\mathcal{F} = \{F, G, H\}$ is shown in Figure 2a.⁴ Any entity E authorized for $\lambda(E)$ is, by the correctness criterion, authorized to operate on all computations such that $\lambda(o) < \lambda(E)$. For example, in Figure 2a, an entity authorised for the set $\lambda(E) = \{F, G\}$ is authorized for the functions F and G as expected. Each label $l \in L$ will be associated with a key κ_l . To outsource a computation o of $F(x)$, C prepares the encoding of x using the key $\kappa_o = \{F\}$ which, by the security criterion, C knows if and only if $\lambda(C) \geq \lambda(o)$ i.e. if and only if $\{F\} \in \lambda(C)$. To compute $F(x)$, S uses the corresponding key κ_o . As before, S may do this if and only if $\{F\} \in \lambda(S)$.

3.1.2 Policies over Function Inputs.

As well as limiting the functions that may be outsourced, we may wish to implement a more fine-grained access control policy determined by input values to functions. For ease of exposition let us assume that all functions have the same domain – for all $F \in \mathcal{F}$, $\text{Dom}(F) = \{0, 1\}^n$ for a positive, non-zero integer n .⁵ We then redefine the security function such that “objects” are now considered to be pairs (F, x) where $F \in \mathcal{F}$ and $x \in \text{Dom}(F)$ – that is, $\lambda : \mathcal{C} \cup \mathcal{S} \cup (\mathcal{F} \times \{0, 1\}^n) \rightarrow L$.

We could, for example, let L be $2^{\mathcal{F}} \times \{0, 1\}^n$. To define the order relation on L we must first define an ordering on the input data. One choice is to define a co-ordinatewise ordering

⁴Nodes for empty sets are excluded from these figures.

⁵If this is not the case then we can define $n = \max_{F \in \mathcal{F}} |\text{Dom}(F)|$ and then redefine all functions F with $\text{Dom}(F) < n$ to satisfy the required property by, for example, adding fixed points $F(x) = x$ for all $x \in \{0, 1\}^n \setminus \text{Dom}(F)$.

on $\{0, 1\}^n$: that is $(x_1, \dots, x_n) \leq (y_1, \dots, y_n)$ if and only if $x_i \leq y_i$ for all i . Then for two labels (\mathcal{G}, x) and (\mathcal{G}', x') in L , $(\mathcal{G}, x) \leq (\mathcal{G}', x')$ if and only if $\mathcal{G} = \{F\}$ for some $F \in \mathcal{F}$, $F \in \mathcal{G}'$ and $x \leq x'$. This poset could be used to ensure delegators and servers operate over a limited range of data values. Specifically, for an entity E , we define $\lambda(E) = (\mathcal{G}, x)$, where $\mathcal{G} \subseteq \mathcal{F}$ and $x \in \{0, 1\}^n$, and $\lambda(o) = (\{F\}, 0^n)$. Then E is authorized to operate on $G(y)$ for all $G \in \mathcal{G}$ and all $y \leq x$. The Hasse diagram for this poset with $\mathcal{F} = \{F, G\}$ and $n = 2$ is shown in Figure 2b. Different choice of orderings over inputs lead to different restrictions e.g. one could consider bit strings as integers and use the natural ordering over integers for $x \leq y$.

An alternative would be to let L be the power set $2^{\mathcal{F} \times \{0, 1\}^n}$. Then each function may be associated with a different range of permissible inputs. An entity E , is labelled by a set of pairs, each comprising a function label and an associated input label – $\lambda(E) = \{(F_1, x_1), \dots, (F_m, x_m)\}$. A computation of $F(x)$ is labelled $\lambda(o) = (\{F\}, x)$. Then, for any two labels,

$$\{(F_1, x_1), \dots, (F_m, x_m)\} \leq \{(F'_1, x'_1), \dots, (F'_{m'}, x'_{m'})\}$$

if and only if $m = 1$ (the first label is a single pair), $F_1 = F'_i$ for some $i \in \{1, \dots, m'\}$ and $x_1 \leq x'_i$.

We remark that the first choice of poset presented in this section corresponds to the idea of assigning a possible data range to each delegator and then authorizing them to outsource a set of functions on that data. For example, a chinese wall policy may result in employees being provided with separate portions of a database. Then, depending on the employee’s role, they are permitted to evaluate a specific set of functions on that data partition. The second choice of poset is more akin to authorizing delegators to outsource a single, specified computation – that is, permitting the entity to evaluate $F(x)$ for only this choice of function and input data. For example, one could imagine an employee being given restricted access to some sensitive data to perform only a particular task, and they should not be able to use the data for other purposes. Finally, we observe that the choice of $L = 2^{\mathcal{F}} \times 2^{\{0, 1\}^n}$ extends this second case to a situation where each function can be associated with an arbitrary sets of input values.

3.1.3 Enforcing General Graph-based Policies.

We have seen how sets of functions and inputs can be mapped to a graph-based access control policy to restrict the functions a delegator may outsource and a server may compute. However, in practical applications, outsourced computation functionality will be required to integrate with existing workflows and existing access control policies. As an example, a company that operates Role-based Access Control (RBAC) on their local network and wishes to provide access to an external VC system must ensure that the same access control requirements are adhered to within this new environment. Additionally, it may not always be the case that a computation may be considered purely in terms of the function and input data. Indeed, although the evaluation will be uniquely determined by such factors, the level of protection required may depend on other contextual information. As a motivating example, consider a summation function over the integers. The semantic meaning of the integers in question may determine the overall classification of this computation – if the integers are city populations then this may not be classified at all, but troop deployments in different regions may be much more sensitive. We now briefly discuss how to formulate access control policies for multi-user VC settings that incorporate additional security labels relying on the environment or external access control policies.

We can use any poset of security labels L to classify each outsourced computation. For example, consider the total order $M = \{\text{Top-Secret}, \text{Secret}, \text{Classified}, \text{Unclassified}\}$ representing the Bell-LaPadula clearance levels [4] and let K to be a set of need-to-know categories.

Then we can enforce the security function $\lambda : \mathcal{C} \cup \mathcal{S} \cup \mathcal{O} \rightarrow L = M \times 2^K$ where K specifies the nature, and M specifies the sensitivity, of the computation $o \in \mathcal{O}$. To delegate or compute $F(x)$, we require that the entity E 's clearance level is at least the classification of the computation: $\lambda(E) \geq \lambda(o)$.

Similarly, it has been shown [12] that the set of roles for a Role-based Access Control Policy can be encoded as a poset. Thus, we could define L to be this role poset to integrate with existing RBAC policies within an organization.

We could also add additional criteria in the mapping to security labels. For example, we could formulate policies dependent on time by setting $\lambda : \mathcal{C} \cup \mathcal{S} \cup (\mathcal{F} \times \mathcal{T}) \rightarrow L$, where \mathcal{T} is a set of time periods. We could then set $L = 2^{\mathcal{F} \times \mathcal{T}}$ as in the previous section to allow entities to operate on specific functions only during specified time periods. On the other hand, by defining $L = M$ as above, certain functions can be more highly protected during certain times of the day. In place of a temporal poset \mathcal{T} we could also apply a geo-spatial poset to classify function evaluations differently depending on location data e.g. a function may be more sensitive if being outsourced from a battlefield as opposed to within a secured building. Finally, policies over function inputs can be extended to include characteristics of the input data (as in the averaging example above, contextual information often changes the level of protection required). Let Z be a set of labels describing data types that functions may be computed over, and define $L = 2^{\mathcal{F} \times (\{0,1\}^n \times Z)}$ for example. This enables the same input to the same function to be classified differently and hence to require different authorization from the delegator and the server.

As mentioned in the introduction, one interesting setting for multi-user VC is when users pay for computation-as-a-service [17, 21]. Users may pay a price per computation, in which case they could be issued a relevant key such as those discussed in Section 3.1.2, perhaps with a short time window in which they may perform the computation. Then, when outsourcing the computation, it is protected by this key to prove that the user has paid for this particular computation. An alternative model [17] would be for a server to allow subscriptions to different tiers of service. A user may pay more to subscribe to a higher tier, and then may submit multiple computations relating to this tier or lower. For example, consider a set $T = \{\text{Gold}, \text{Silver}, \text{Bronze}\}$ comprising tiers that users may subscribe to; **Gold** service may allow access to more computational resources, or access during busier periods, than **Silver**, and so on. For simplicity, suppose the resources to be considered are just RAM capacity and the number of processor cores available, and let R and C be sets comprising the available quantities of each. Then, the set of security labels L is set to be $L = (R \times C) \cup T$ — that is, the cartesian product of the available resources, with the addition of labels for each service tier. A user that pays for **Gold** membership can be assigned the label $\lambda(C) = \{\text{Gold}\}$. Each computation is labelled by the resources that it requires. The poset ordering over L allows each tier to access different subsets of resources (e.g. if a computation is particularly memory intensive then it may require a higher subscription fee to compute), as illustrated in Figure 2c.

3.2 Verification Policies

Thus far, we have discussed policies restricting the entities that may delegate and evaluate given computations. In the multi-user VC setting, it is equally important to apply access control to the act of verifying the results of a computation and, in particular, learning the output. Clearly, when considering computations over confidential data, it is not always appropriate to publish the results. As a trivial example, outsourcing the identity function could “legitimately” leak confidential data. The scheme of Alderman et al. [2] distinguished between blind verification (to ensure correctness) and output retrieval (to learn results). Here, we assume that all entities

are able to blind verify a result, but the set of entities that may learn the actual output should be restricted⁶. Thus, the set of verifiers we refer to in the following are just those distinguished verifiers that are able to retrieve computation results, and we wish to restrict the results that each may read.

The notion of “no write down” is an important property in many traditional access control policies. In the VC setting, this amounts to ensuring that any verifier should at least have the access rights of the delegating or computing entity, that is $\lambda(C), \lambda(S) \leq \lambda(V)$ e.g. a verifier may not read a result arising from input data that he is not also authorized to read. In some cases, however, the KDC may decide that the results of some computations are not as highly classified as the input data, or indeed the act of computing it. For example, statistics of company spending across all departments may be less sensitive than the spending of the research department alone. Alternatively, the results of a computation over classified data may be included in a public document, despite the input data remaining classified. The encoded output from the computation may be published with the verification key such that recipients can verify the legitimacy of the published values. It may also be reasonable to expect that only authorised and trusted reviewers of the document should be able to verify, or that some (lower) form of clearance is still required to access the result.

3.2.1 Published Verification Token.

In this second case, where verification tokens are published according to some policy, we can use similar posets and security functions to those in Section 3.1. The user population is the set of delegators \mathcal{C} and verifiers \mathcal{V} , and objects are encoded outputs. The security function is defined to be $\lambda^V : \mathcal{V} \cup \mathcal{C} \cup \mathcal{O} \rightarrow L$, where \mathcal{O} is the set of outsourced computations and (L, \leq) is a poset of security labels. As part of the VC functionality, when encoding an input, C also generates an output retrieval key RK_o that enables the result of a computation to be retrieved given the encoded output and the verification key. We will protect this retrieval key and require the verifier to satisfy the verification policy in order to use it – that is $\lambda(V) \geq \lambda(o)$ for $o \in \mathcal{O}$. Since the delegator generates the retrieval key, and hence may also use it to retrieve the output, we also assume $\lambda(C) \geq \lambda(o)$. This is reasonable given that the delegator must have the right to compute any function that he outsources and hence (resources permitting) could certainly learn the result. As before we can extend the definition of λ and of L to accommodate more fine-grained policies over outputs.

3.2.2 Enforcement of No Write Down Policies.

As mentioned, “No Write Down” is an important requirement of many access control policies. It ensures that an entity C may not write (encrypt) an object to a lower classification level $\lambda(o) < \lambda(C)$ as this could constitute a leak of classified data. It is, of course, possible for a delegator to write the result at his (maximal) clearance level and then any verifiers with higher access rights, by the correctness criterion of the poset, will be able to read the result. However, we may want to protect a result at a higher classification level (write up) e.g. when preparing a report that should only be read by management.

To do so, we define two posets: a computation poset $\mathcal{P}_C = (L, \leq)$ (as per Section 3.1) and a verification poset \mathcal{P}_V . \mathcal{P}_V is constructed by inverting the order relation on \mathcal{P}_C where $\mathcal{P}_V = (L, \geq)$ i.e. $x \leq y$ in \mathcal{P}_V if and only if $y \leq x$ in \mathcal{P}_C . Then, delegation and computation of $o \in \mathcal{O}$ are performed using a key associated to $\lambda^C(o) \in \mathcal{P}_C$. Retrieving the result of a

⁶Note, we could use the same techniques to restrict blind verification if desired.

computation requires a key related to $\lambda^V(o) \in \mathcal{P}_V$. A verifier V may then retrieve the output if and only if $\lambda^V(V) \geq \lambda^V(o)$.

3.3 Verification Policies

Thus far, we have discussed policies restricting the entities that may delegate and evaluate given computations. In the multi-user VC setting, it is equally important to apply access control to the act of verifying the results of a computation and, in particular, learning the output. Clearly, when considering computations over confidential data, it is not always appropriate to publish the results. As a trivial example, outsourcing the identity function could “legitimately” leak confidential data. The scheme of Alderman et al. [2] distinguished between blind verification (to ensure correctness) and output retrieval (to learn results). Here, we assume that all entities are able to blind verify a result, but the set of entities that may learn the actual output should be restricted⁷. Thus, the set of verifiers we refer to in the following are just those distinguished verifiers that are able to retrieve computation results, and we wish to restrict the results that each may read.

The notion of “no write down” is an important property in many traditional access control policies. In the VC setting, this amounts to ensuring that any verifier should at least have the access rights of the delegating or computing entity, that is $\lambda(C), \lambda(S) \leq \lambda(V)$ e.g. a verifier may not read a result arising from input data that he is not also authorized to read. In some cases, however, the KDC may decide that the results of some computations are not as highly classified as the input data, or indeed the act of computing it. For example, statistics of company spending across all departments may be less sensitive than the spending of the research department alone. Alternatively, the results of a computation over classified data may be included in a public document, despite the input data remaining classified. The encoded output from the computation may be published with the verification key such that recipients can verify the legitimacy of the published values. It may also be reasonable to expect that only authorised and trusted reviewers of the document should be able to verify, or that some (lower) form of clearance is still required to access the result.

3.3.1 Published Verification Token.

In this second case, where verification tokens are published according to some policy, we can use similar posets and security functions to those in Section 3.1. The user population is the set of delegators \mathcal{C} and verifiers \mathcal{V} , and objects are encoded outputs. The security function is defined to be $\lambda^V : \mathcal{V} \cup \mathcal{C} \cup \mathcal{O} \rightarrow L$, where \mathcal{O} is the set of outsourced computations and (L, \leq) is a poset of security labels. As part of the VC functionality, when encoding an input, C also generates an output retrieval key RK_o that enables the result of a computation to be retrieved given the encoded output and the verification key. We will protect this retrieval key and require the verifier to satisfy the verification policy in order to use it – that is $\lambda(V) \geq \lambda(o)$ for $o \in \mathcal{O}$. Since the delegator generates the retrieval key, and hence may also use it to retrieve the output, we also assume $\lambda(C) \geq \lambda(o)$. This is reasonable given that the delegator must have the right to compute any function that he outsources and hence (resources permitting) could certainly learn the result. As before we can extend the definition of λ and of L to accommodate more fine-grained policies over outputs.

⁷Note, we could use the same techniques to restrict blind verification if desired.

3.3.2 Enforcement of No Write Down Policies.

As mentioned, “No Write Down” is an important requirement of many access control policies. It ensures that an entity C may not write (encrypt) an object to a lower classification level $\lambda(o) < \lambda(C)$ as this could constitute a leak of classified data. It is, of course, possible for a delegator to write the result at his (maximal) clearance level and then any verifiers with higher access rights, by the correctness criterion of the poset, will be able to read the result. However, we may want to protect a result at a higher classification level (write up) e.g. when preparing a report that should only be read by management.

To do so, we define two posets: a computation poset $\mathcal{P}_C = (L, \leq)$ (as per Section 3.1) and a verification poset \mathcal{P}_V . \mathcal{P}_V is constructed by inverting the order relation on \mathcal{P}_C where $\mathcal{P}_V = (L, \geq)$ i.e. $x \leq y$ in \mathcal{P}_V if and only if $y \leq x$ in \mathcal{P}_C . Then, delegation and computation of $o \in \mathcal{O}$ are performed using a key associated to $\lambda^C(o) \in \mathcal{P}_C$. Retrieving the result of a computation requires a key related to $\lambda^V(o) \in \mathcal{P}_V$. A verifier V may then retrieve the output if and only if $\lambda^V(V) \geq \lambda^V(o)$.

4 PVC with Access Control

We now formally define the notion of PVC-AC to enforce graph-based access control policies over delegators, servers and verifiers. Generally, as discussed above, we wish to impose restrictions on three activities: first, we wish to specify a (write) policy determining the computations a client is authorized to delegate – this means restricting the inputs a client may correctly prepare; second, we specify a (read) policy that determines the computations a server may compute – that is, encoded inputs he may access; lastly, we specify a (read) policy dictating the outputs a verifier may read. We differentiate between computation policies over delegators and servers, denoted $\lambda^C(\cdot)$, and verification policies over delegators and verifiers, denoted $\lambda^V(\cdot)$. We may also use $\lambda(\cdot)$ to denote both labels where appropriate. Finally, \mathcal{P}_C and \mathcal{P}_V denote posets encoding the computation and verification policies respectively. Recall that the set of outsourced computations is denoted by \mathcal{O} . A specific element $o \in \mathcal{O}$ consists of the function F , some input x as well as some auxiliary input (such as a label or additional contextual information).

We extend the functionality of the KDC from [2] to grant access control credentials to delegators, servers and verifiers. We may split the responsibilities between two KDCs: a *computation* KDC (CKDC) that generates function keys for the VC functionality, and an *authorization* KDC (AKDC) that manages access control policies. The AKDC could be a trusted authority on entities to grant relevant permissions, and may be used by multiple systems. For clarity, here we use only one KDC that performs both duties.

Definition 4. A Publicly Verifiable Outsourced Computation Scheme with Access Control (PVC-AC) comprises the following algorithms:

- $\text{Setup}(1^\kappa, \mathcal{P}_C, \mathcal{P}_V) \rightarrow (\text{PP}, \text{MK})$: Run by the KDC to generate public parameters for the PVC-AC system as well as secret information used to generate keys.
- $\text{Fnlit}(F, \text{MK}, \text{PP}) \rightarrow PK_F$: Run by the KDC to generate a delegation key, PK_F , for a function F .
- $\text{Register}(\text{ID}, \lambda(\text{ID}), \text{MK}, \text{PP}) \rightarrow SK_{\text{ID}}$: Run by the KDC to create a personalised key SK_{ID} for an entity with identifier ID ,⁸ which grants rights for the label $\lambda(\text{ID})$ according to the computation or verification policy.

⁸In future algorithms this will be S , C or V to denote a server, delegator or verifier respectively.

- $\text{Certify}(S, F, \text{MK}, \text{PP}) \rightarrow EK_{F,S}$: The KDC creates an evaluation key $EK_{F,S}$ for a function F and server S .
- $\text{ProbGen}(x, SK_C, PK_F, \lambda(C), \lambda(o), \text{PP}) \rightarrow (\sigma_o, VK_o, RK_o)$: Run by a delegator C to outsource the computation of $F(x)$. C may do so only if it satisfies the computation policy – that is $\lambda^C(C) \geq \lambda^C(o)$. It outputs an encoded input σ_o , a public verification key VK_o to verify correctness, and an output retrieval key RK_o which verifiers satisfying $\lambda^V(V) \geq \lambda^V(o)$ may use to read $F(x)$.
- $\text{Compute}(\sigma_o, EK_{F,S}, SK_S, \lambda^C(S), \lambda^C(o), \text{PP}) \rightarrow \theta_o$: Run by a server S holding an evaluation key $EK_{F,S}$, SK_S and an encoded input σ_o to evaluate $F(x)$ and output an encoding, θ_o , of the result. This succeeds if and only if $\lambda^C(S) \geq \lambda^C(o)$, that is S satisfies the computation policy for this evaluation.
- $\text{Verify}(\theta_o, SK_V, VK_o, \lambda^V(V), \lambda^V(o), \text{PP}) \rightarrow (\tilde{y}, \tau_{\theta_o})$: Verification comprises two steps:
 - $\text{BVerif}(\text{PP}, \theta_o, VK_o) \rightarrow (RT_o, \tau_{\theta_o})$: Run by *any* verifier to produce a retrieval token RT_o and a token $\tau_{\theta_o} = (\text{accept}, S)$ for a correct result, or $\tau_{\theta_o} = (\text{reject}, S)$ to signify a cheating server S .
 - $\text{Retrieve}(SK_V, RT_o, \tau_{\theta_o}, VK_o, RK_o, \lambda^V(V), \lambda^V(o), \text{PP}) \rightarrow \tilde{y}$: Run by verifiers V in possession of the output retrieval key RK_o and the retrieval token RT_o from BVerif . If $\lambda^V(V) \geq \lambda^V(o)$ then V should be able to read the actual result $\tilde{y} = F(x)$ or \perp .
- $\text{Revoke}(\tau_{\theta_o}, F, \text{MK}, \text{PP}) \rightarrow \{EK_{F,S'}\}$ or \perp : If a verifier reports a misbehaving server (i.e. Verify returned $\tau_{\theta_o} = (\text{reject}, S)$), the KDC issues updated evaluation keys $EK_{F,S'}$ to all servers, preventing S from performing further computations. If $\tau_{\theta_o} = (\text{accept}, S)$ then this algorithm should output \perp .

We say that a PVC-AC scheme is *correct* if for all functions $F \in \mathcal{F}$, inputs x , honestly generated parameters, and honestly registered entities C , S and V (delegator, certified server and verifier respectively) such that $\lambda^C(C), \lambda^C(S) \geq \lambda^C(o)$ and $\lambda^V(C), \lambda^V(V) \geq \lambda^V(o)$, if S honestly runs Compute for F on an encoding of x generated by C , then V running Verify on the output from S will almost certainly output accept and $F(x)$. That is, if all algorithms are run honestly by authorised parties then the verifier should almost certainly accept. More formally we can write

Definition 5 (Correctness). *A Publicly Verifiable Outsourced Computation Scheme with Access Control (PVC-AC) is correct for a family of functions \mathcal{F} if for all functions $F \in \mathcal{F}$ and inputs*

x , where $\text{negl}(\cdot)$ is a negligible function of its input:

$$\begin{aligned}
& \Pr[\text{Setup}(1^\lambda) \rightarrow (\text{PP}, \text{MK}), \\
& \quad \text{FnInit}(F, \text{MK}, \text{PP}) \rightarrow PK_F, \\
& \quad \text{Register}(C, \lambda(C), \text{MK}, \text{PP}) \rightarrow SK_C, \\
& \quad \text{Register}(S, \lambda(S), \text{MK}, \text{PP}) \rightarrow SK_S, \\
& \quad \text{Register}(V, \lambda(V), \text{MK}, \text{PP}) \rightarrow SK_V, \\
& \quad \text{Certify}(S, F, \text{MK}, \text{PP}) \rightarrow EK_{F,S}, \\
& \quad \text{ProbGen}(x, SK_C, PK_F, \lambda^C(C), \lambda^C(o), \text{PP}) \rightarrow (\sigma_x, VK_{F,x}, RK_{F,x}), \\
& \quad \text{Verify}(\text{Compute}(\sigma_x, EK_{F,S}, SK_S, \lambda^C(S), \lambda^C(o), PP), SK_V, \\
& \quad \quad VK_o, \lambda^V(V), \lambda^V(o), \text{PP}) \rightarrow (F(x), (\text{accept}, S))] \\
& = 1 - \text{negl}(\lambda),
\end{aligned}$$

where $\lambda^C(C) \geq \lambda^C(o)$, $\lambda^V(C) \geq \lambda^V(o)$, $\lambda^C(S) \geq \lambda^C(o)$ and $\lambda^V(V) \geq \lambda^V(o)$ holds.

5 Security Definitions

We now introduce several security models capturing requirements of PVC-AC. These are important both to ensure that the enforcement mechanism for the access control policies correctly captures the required properties, and that the access control mechanism and the PVC implementation (if built from separate primitives) interact securely. As noted by Ferrara et al. [14], it is not always immediate that (probabilistic) cryptographic enforcement mechanisms can safely implement policies that are generally specified in absolute terms; hence formally defining and proving the required security goals is necessary. Notions of Public Verifiability, Revocation and Vindictive Servers follow naturally from [2] with additional inputs for policy declarations. As these do not rely on the access control properties, the proofs follow naturally. We let o denote an outsourced computation (including descriptors such as the function, input and environmental factors) such that $\lambda(o)$ encompasses the required clearance to operate on this computation. \mathcal{A}^O denotes \mathcal{A} having oracle access to the following functions: $\text{FnInit}(\cdot, \text{MK}, \text{PP})$, $\text{Register}(\cdot, \cdot, \text{MK}, \text{PP})$, $\text{Certify}(\cdot, \cdot, \text{MK}, \text{PP})$, $\text{ProbGen}(\cdot, SK_{(\cdot)}, \cdot, \cdot, \cdot, \text{PP})$, $\text{Compute}(\cdot, \cdot, SK_{(\cdot)}, \cdot, \cdot, PP)$ and $\text{Revoke}(\cdot, \cdot, \text{MK}, \text{PP})$. Each oracle, aside from Register , simply runs the relevant algorithm, with any restrictions on permissible queries stated below. The Register oracle specified in Oracle Query 7 is used in all games except Authorized Verification which instead uses Oracle Query 8 which uses the verification policy instead of the computation policy. We define the *advantage* and *security* of \mathcal{A} in each of the following games as:

Definition 6. *The advantage of a PPT adversary \mathcal{A} making a polynomial number of queries q is defined as follows:*

- For $\mathbf{X} \in \{AO, AC\}$:

$$\text{Adv}_{\mathcal{A}}^{\mathbf{X}}(\mathcal{PVC}_{AC}, F, 1^\kappa, q) = \Pr[\mathbf{Exp}_{\mathcal{A}}^{\mathbf{X}}[\mathcal{PVC}_{AC}, F, 1^\kappa] = 1]$$

- For $\mathbf{X} \in \{AV, wIP\}$:

$$\text{Adv}_{\mathcal{A}}^{\mathbf{X}}(\mathcal{PVC}_{AC}, F, 1^\kappa, q) = |\Pr[\mathbf{Exp}_{\mathcal{A}}^{\mathbf{X}}[\mathcal{PVC}_{AC}, F, 1^\kappa] = 1] - \frac{1}{2}|$$

A PVC-AC is secure against Game \mathbf{X} for a function F , if for all PPT adversaries \mathcal{A} , $\text{Adv}_{\mathcal{A}}^{\mathbf{X}}(\mathcal{PVC}_{AC}, F, 1^\kappa, q) \leq \text{negl}(\kappa)$.

Game 5 $\text{Exp}_{\mathcal{A}}^{\text{AO}}[\mathcal{PVC}_{\text{AC}}, F, 1^\kappa]$:

```

1:  $L = \epsilon, o = \perp$ 
2:  $(\text{PP}, \text{MK}) \leftarrow \text{Setup}(1^\kappa, \mathcal{P}_C, \mathcal{P}_V)$ ;
3:  $\lambda^C(\mathcal{A}) \leftarrow \mathcal{A}^{\text{O}}(\text{PP})$ ;
4:  $\text{SK}_{\mathcal{A}} \leftarrow \text{Register}(\mathcal{A}, \lambda^C(\mathcal{A}), \text{MK}, \text{PP})$ ;
5:  $L = L \cup \lambda^C(\mathcal{A})$ ;
6:  $o = (F, x, \text{aux}) \leftarrow \mathcal{A}^{\text{O}}(\lambda^C(\mathcal{A}), \text{SK}_{\mathcal{A}}, \text{PP})$ ;
7: for all  $\lambda^C(v_i) \in L$  do
8:     if  $\lambda^C(o) \leq \lambda^C(v_i)$  then
9:         return 0
10:  $\text{SK}_S \leftarrow \text{Register}(S, \lambda^C(S), \text{MK}, \text{PP})$  s.t.  $\lambda^C(S) \geq \lambda^C(o)$ ;
11:  $\text{SK}_V \leftarrow \text{Register}(V, \lambda^V(V), \text{MK}, \text{PP})$  s.t.  $\lambda^V(V) \geq \lambda^V(o)$ ;
12:  $\text{PK}_F \leftarrow \text{FnInit}(F, \text{MK}, \text{PP})$ ;
13:  $\text{EK}_{F,S} \leftarrow \text{Certify}(S, F, \text{MK}, \text{PP})$ ;
14:  $(\sigma_o, \text{VK}_o, \text{RK}_o) \leftarrow \mathcal{A}^{\text{O}}(o, \lambda(o), \lambda(\mathcal{A}), \text{PK}_F, \text{PP})$ ;
15:  $\theta_o \leftarrow \text{Compute}(\sigma_o, \text{EK}_{F,S}, \text{SK}_S, \lambda^C(S), \lambda^C(o), \text{PP})$ ;
16:  $(\tilde{y}, \tau_{\theta_o}) \leftarrow \text{Verify}(\theta_o, \text{SK}_V, \text{VK}_o, \lambda^V(V), \lambda^V(o), \text{PP})$ ;
17: if  $((\tilde{y}, \tau_{\theta_o}) \neq (\perp, (\text{reject}, \mathcal{A})))$  then return 1
18: else return 0

```

5.1 Authorized Outsourcing

In Game 5 we formalize the notion that a delegator may not outsource any computation for which he is not authorized – that is, any computation o where $\lambda^C(o) \not\leq \lambda^C(C)$. Clearly, within our framework we cannot make any guarantees about what an entity may do externally. For instance, we cannot enforce that a client does not circumvent the access control policies using an external server not subject to these controls, however we do enforce that to use the provided functionality (i.e. to contract an available server registered in this system) they must prove authorization. This assumption seems reasonable, taking into consideration organizational boundaries – external control should perhaps be enforced by limiting external communication channels (e.g. using a strict firewall). Similarly, we cannot enforce that an entity does not share key content, but we do ensure that such collusion does not enable access that either entity alone could not access. Additionally, due to the revocation functionality, it may be undesirable for a server to share key material as he must trust the additional server not to cheat.

The game proceeds as follows. First the challenger runs the setup procedures for the scheme and registers the adversary for the adversary’s choice of security label. The adversary is given the public information, his secret key and oracle access, and outputs a choice of outsourced computation o to attack, including the function and input, F and x . We require that no security label that \mathcal{A} has queried to the Register oracle may allow the derivation of a key for $\lambda^C(o)$ as this would be a trivial win, and we return 0 since the adversary has not found a valid attack target. We also require that any ID that \mathcal{A} queries to its oracle must previously have been queried to the Register oracle (which is in keeping with realistic operation). The ProbGen oracle may not be queried on the challenge computation o . Note that \mathcal{A} may register entities of his choice for any security label that is not an ancestor of the challenge label, and hence does not need oracle access for these labels. BVerif relies only on public information and \mathcal{A} may Register any verifier he likes to learn $\kappa_{\lambda^V(\cdot)}$ to run Retrieve, so these oracles are not required. \mathcal{C} sets up keys for F and simulates registering two entities that are authorized to perform the computation and verification for o respectively. The adversary is then challenged, given all information that a real attacker may learn and oracle access, to output an encoded input that the Compute and Verify algorithms will accept – that is, an unauthorized adversary must produce an input that is accepted and computed on by honest entities.

Oracle Query 7 $\mathcal{O}^{\text{Register}}(ID, \lambda(ID), \text{MK}, \text{PP})$:

```

1: if  $(o \neq \perp) \wedge \lambda(ID) \geq \lambda^C(o)$  then
2:   return  $\perp$ 
3:  $L = L \cup \lambda(ID)$ ;
4: return  $SK_{ID} \leftarrow \text{Register}(ID, \lambda(ID), \text{MK}, \text{PP})$ 

```

5.2 Authorized Computation

Authorized Computation, in Game 6, proceeds similarly to Game 5 and captures the notion that a computational result should only be considered valid if generated by an authorized party. The challenger sets up the system and registers the adversary for its choice of security label. \mathcal{A} then chooses a target computation that it is not authorized to compute by *any* of the keys it holds. The challenger simulates two entities: a delegator and a verifier that are authorized for this computation and certifies \mathcal{A} as a server for F . The challenger creates an encoded input by running ProbGen on the adversary's target computation and gives this to \mathcal{A} who must output an encoded output that is accepted by the verifier. Note that although the adversary chooses the input to the computation, and therefore can certainly compute $F(x)$, it still should not be able to convince the verifier to accept. The oracle queries are the same as in the previous section except that ProbGen may be queried for any input (even the challenge computation) and the Compute oracle cannot be run on the challenge input.

Game 6 $\text{Exp}_{\mathcal{A}}^{AC}[\mathcal{PVC}_{AC}, F, 1^\kappa]$:

```

1:  $L = \epsilon, o = \perp$ 
2:  $(\text{PP}, \text{MK}) \leftarrow \text{Setup}(1^\kappa, \mathcal{P}_C, \mathcal{P}_V)$ ;
3:  $\lambda^C(\mathcal{A}) \leftarrow \mathcal{A}^{\mathcal{O}}(\text{PP})$ ;
4:  $SK_{\mathcal{A}} \leftarrow \text{Register}(\mathcal{A}, \lambda^C(\mathcal{A}), \text{MK}, \text{PP})$ ;
5:  $L = L \cup \lambda^C(\mathcal{A})$ ;
6:  $o = (F, x, \text{aux}) \leftarrow \mathcal{A}^{\mathcal{O}}(\lambda^C(\mathcal{A}), SK_{\mathcal{A}}, \text{PP})$ ;
7: for all  $\lambda^C(v_i) \in L$  do
8:   if  $\lambda^C(o) \leq \lambda^C(v_i)$  then
9:     return 0
10:  $SK_C \leftarrow \text{Register}(C, \lambda^C(C), \text{MK}, \text{PP})$  s.t.  $\lambda^C(C) \geq \lambda^C(o)$ ;
11:  $SK_V \leftarrow \text{Register}(V, \lambda^V(V), \text{MK}, \text{PP})$  s.t.  $\lambda^V(V) \geq \lambda^V(o)$ ;
12:  $PK_F \leftarrow \text{FnInit}(F, \text{MK}, \text{PP})$ ;
13:  $EK_{F, \mathcal{A}} \leftarrow \text{Certify}(\mathcal{A}, F, \text{MK}, \text{PP})$ ;
14:  $(\sigma_o, VK_o, RK_o) \leftarrow \text{ProbGen}(x, SK_C, PK_F, \lambda(C), \lambda(o), \text{PP})$ ;
15:  $\theta_o \leftarrow \mathcal{A}^{\mathcal{O}}(\sigma_o, VK_o, o, \lambda^C(o), EK_{F, \mathcal{A}}, PK_F, RK_o, \text{PP})$ ;
16:  $(\tilde{y}, \tau_{\theta_o}) \leftarrow \text{Verify}(\theta_o, SK_V, VK_o, \lambda^V(V), \lambda^V(o), RK_o, \text{PP})$ ;
17: if  $((\tilde{y}, \tau_{\theta_o}) \neq (\perp, (\text{reject}, \mathcal{A})))$  then return 1
18: else return 0

```

5.3 Authorized Verification

In Game 7 we capture the notion that an unauthorized verifier should not be able to learn the output of a computation. This is formulated in an indistinguishability game, where the adversary chooses two computations to give to the challenger. To avoid a trivial win, we require that neither of the associated labels in the verification poset are less than or equal to a label queried to the Register oracle. The challenger chooses *one* of these at random and simulates outsourcing and computing this computation. The adversary is provided with the encoded output from this computation, along with the verification keys and other public information, and must guess which computation was chosen. As the adversary chose the computations, it can work out the results; however, no information about the result should leak from the encoded output unless the verifier is authorized, and hence the adversary should be unable to tell which result the output corresponds to. The oracle queries are the same as in the previous section

Game 7 $\text{Exp}_{\mathcal{A}}^{AV}[\mathcal{PVC}_{AC}, F, 1^\kappa]$:

- 1: $L = \epsilon, o_0 = 0_1 = \perp$
 - 2: $(PP, MK) \leftarrow \text{Setup}(1^\kappa, \mathcal{P}_C, \mathcal{P}_V)$;
 - 3: $\lambda^V(\mathcal{A}) \leftarrow \mathcal{A}^{\mathcal{O}}(PP)$;
 - 4: $SK_{\mathcal{A}} \leftarrow \text{Register}(\mathcal{A}, \lambda^V(\mathcal{A}), MK, PP)$;
 - 5: $L = L \cup \lambda^V(\mathcal{A})$;
 - 6: $(o_0, o_1) \xleftarrow{\$} \mathcal{A}^{\mathcal{O}}(SK_{\mathcal{A}}, PP)$ s.t. $\lambda^V(o_0), \lambda^V(o_1) > \lambda^V(v_i) \forall \lambda^V(v_i) \in L$;
 - 7: $b \xleftarrow{\$} \{0, 1\}$;
 - 8: $SK_C \leftarrow \text{Register}(C, \lambda^C(C), MK, PP)$ s.t. $\lambda^C(C) \geq \lambda^C(o_b)$;
 - 9: $SK_S \leftarrow \text{Register}(V, \lambda^C(S), MK, PP)$ s.t. $\lambda^C(S) \geq \lambda^C(o_b)$;
 - 10: $PK_F \leftarrow \text{FnInit}(F, MK, PP)$;
 - 11: $EK_{F,S} \leftarrow \text{Certify}(S, F, MK, PP)$;
 - 12: $(\sigma_{o_b}, VK_{o_b}, RK_{o_b}) \leftarrow \text{ProbGen}(x, SK_C, PK_F, \lambda(C), \lambda(o_b), PP)$;
 - 13: $\theta_{o_b} \leftarrow \text{Compute}(\sigma_{o_b}, EK_{F,S}, SK_S, \lambda^C(S), \lambda^C(o_b), PP)$;
 - 14: $b' \leftarrow \mathcal{A}^{\mathcal{O}}(\theta_{o_b}, VK_{o_b}, RK_{o_b}, \lambda(o_0), \lambda(o_1), SK_{\mathcal{A}}, PK_F, PP)$;
 - 15: **if** $b' = b$ **then return** 1
 - 16: **else return** 0
-

Oracle Query 8 $\mathcal{O}^{\text{Register}}(ID, \lambda(ID), MK, PP)$:

- 1: **if** $(o_b \neq \perp) \wedge \lambda(ID) \geq \lambda^V(o_b)$ **then**
 - 2: **return** \perp
 - 3: $L = L \cup \lambda(ID)$;
 - 4: **return** $SK_{ID} \leftarrow \text{Register}(ID, \lambda(ID), MK, PP)$
-

except that ProbGen and Compute may be queried for any input. The adversary is also given access to a Retrieve oracle (which cannot be queried for the challenge computation) in order to observe outputs for the challenge label.

5.4 Weak Input Privacy

The notion of weak input privacy, captured in Game 8, is not as strong as the input privacy often considered in PVC settings where computational servers learn nothing about the data they are working on. Instead, we are interested in ensuring that servers (or other entities) that are not authorized to access (compute on) the input data may not learn x . If full input privacy is required then the underlying PVC scheme, such as the KP-ABE based one used as a black box in the construction, could be replaced by one built from a predicate encryption scheme for the same class of functions.

The game begins like the Authorized Verification game with the adversary selecting two computations that he is not authorized for by *any* key that he has queried to the Register oracle. The Compute oracle can be queried for any input except the challenge inputs. The challenger then simulates the outsourcing of one of these computations and gives the adversary the resulting encoded input. The adversary must guess which computation the input corresponds to i.e. which input data is encoded in the input.

6 Instantiation

We now provide an example provably secure instantiation of PVC-AC. The approach we take makes generic, black-box use of existing RPVC schemes such as the Key-policy Attribute-based Encryption approaches [2, 20], and introduces the use of a symmetric Key Assignment Scheme (KAS) to restrict the behaviour of entities.

Game 8 Exp $_{\mathcal{A}}^{wIP}[\mathcal{PVC}_{AC}, F, 1^\kappa]$:

- 1: $L = \epsilon, o = \perp$
 - 2: $(PP, MK) \leftarrow \text{Setup}(1^\kappa, \mathcal{P}_C, \mathcal{P}_V)$;
 - 3: $\lambda^C(\mathcal{A}) \leftarrow \mathcal{A}^O(PP)$;
 - 4: $SK_{\mathcal{A}} \leftarrow \text{Register}(\mathcal{A}, \lambda^C(\mathcal{A}), MK, PP)$;
 - 5: $L = L \cup \lambda^C(\mathcal{A})$;
 - 6: $(o_0, o_1) \xleftarrow{\$} \mathcal{A}^O(SK_{\mathcal{A}}, PP)$ s.t. $\lambda^C(o_0), \lambda^C(o_1) > \lambda^C(v_i) \forall \lambda^C(v_i) \in L$;
 - 7: $b \xleftarrow{\$} \{0, 1\}$;
 - 8: $SK_C \leftarrow \text{Register}(C, \lambda^C(C), MK, PP)$ s.t. $\lambda^C(C) \geq \lambda^C(o_b)$;
 - 9: $SK_V \leftarrow \text{Register}(V, \lambda^V(V), MK, PP)$ s.t. $\lambda^V(V) \geq \lambda^V(o_b)$;
 - 10: $PK_F \leftarrow \text{FnInit}(F, MK, PP)$;
 - 11: $EK_{F, \mathcal{A}} \leftarrow \text{Certify}(\mathcal{A}, F, MK, PP)$;
 - 12: $(\sigma_{o_b}, VK_{o_b}, RK_{o_b}) \leftarrow \text{ProbGen}(x, SK_C, PK_F, \lambda(C), \lambda(o_b), PP)$;
 - 13: $b' \leftarrow \mathcal{A}^O(\sigma_{o_b}, VK_{o_b}, \lambda(o_0), \lambda(o_1), EK_{F, \mathcal{A}}, SK_{\mathcal{A}}, PK_F, PP)$;
 - 14: **if** $b' = b$ **then return** 1
 - 15: **else return** 0
-

6.1 Informal Overview

The graph-based policies discussed in Section 3 assign labels to each entity and outsourced computation. The ordering relation and the correctness criterion ensures that entities are authorized for all appropriate computations (those that are descendants of their label in the Hasse diagram of the poset). A KAS is designed to enforce such policies, and assigns a key to each label. Each entity is provided with a key corresponding to their label, and they may derive keys for all descendants. As per the security criterion, users may not collude to derive keys for which they are not authorized.

In our setting, we restrict the computations a delegator may outsource, the computations a server may perform, and the results a verifier may learn. We use two independent KASs instantiated over the computation and verification posets, \mathcal{P}_C and \mathcal{P}_V respectively. Delegators and servers are issued a key according to their label in the computation poset, while delegators and verifiers are given a key from the verification KAS. Appropriate keys from the computation KAS are used to encrypt the encoded input for a computation. To encode an input in a manner that will be accepted by a server, delegators must encrypt the encoded input using the key, $\kappa_{\lambda^C(o)}$, associated to the computation within the computation poset. To do so, the delegator must be able to derive $\kappa_{\lambda^C(o)}$ and hence must hold a key at that level or higher i.e. $\lambda^C(C) \geq \lambda^C(o)$ – delegators must be authorized by the KDC to outsource the computation. Similarly, only authorized servers can derive the decryption key to access the encoded input and perform the computation. By the IND-CPA security of the symmetric encryption scheme used, no information about the encoded input is learnt by an unauthorized entity. To enforce verification policies, delegators generate a retrieval key during delegation, and encrypt this using an appropriate key from the verification KAS. Then, only verifiers that can derive this key may decrypt the ciphertext to recover the retrieval key and learn the output.

6.2 Formal Details

Let $\mathcal{RPVC} = (\text{RPVC.Setup}, \text{RPVC.FnInit}, \text{RPVC.Register}, \text{RPVC.Certify}, \text{RPVC.ProbGen}, \text{RPVC.Compute}, \text{RPVC.BVerif}, \text{RPVC.Retrieve}, \text{RPVC.Revoke})$ be an RPVC scheme, as defined by Alderman et al. [2], for a class of functions \mathcal{F} . Let $\mathcal{SE} = (\text{SE.KeyGen}, \text{SE.Encrypt}, \text{SE.Decrypt})$ be an authenticated symmetric encryption scheme secure in the sense of $\text{IND-CPA} \wedge \text{INT-PTXT}$, and let $\mathcal{KAS} = (\text{KAS.MakeKeys}, \text{KAS.MakeSecrets}, \text{KAS.MakePublicData}, \text{KAS.GetKey})$ be a Key Assignment Scheme secure against Strong-Key Indistinguishability⁹ whose keys are

⁹It has been shown that S-KI is equivalent to KI [9]. We make use of the additional queries in the S-KI game, but due to the equivalence this is not a strengthening of the assumptions. It is interesting to note that the form

compatible with \mathcal{SE} . Finally, let \mathcal{P}_C denote the poset encoding computation policies (e.g. (L, \leq)), and similarly let \mathcal{P}_V denote the poset encoding verification policies (e.g. (L, \geq))¹⁰. Then, for the same class of functions \mathcal{F} , there is a PVC-AC scheme $\mathcal{PVC}_{AC} = (\text{PVC}_{AC}.\text{Setup}, \text{PVC}_{AC}.\text{FnInit}, \text{PVC}_{AC}.\text{Register}, \text{PVC}_{AC}.\text{Certify}, \text{PVC}_{AC}.\text{ProbGen}, \text{PVC}_{AC}.\text{Compute}, \text{PVC}_{AC}.\text{Verify}, \text{PVC}_{AC}.\text{Revoke})$ defined in Algorithms 1-9.

Alg. 1 $\text{Setup}(1^\kappa, \mathcal{P}_C, \mathcal{P}_V) \rightarrow (\text{PP}, \text{MK})$

- 1: $(\text{PP}', \text{MK}') \leftarrow \text{RPVC}.\text{Setup}(1^\kappa)$
 - 2: $\kappa_{\mathcal{P}_C} \leftarrow \text{KAS}.\text{MakeKeys}(1^\kappa, \mathcal{P}_C)$
 - 3: $\omega_{\mathcal{P}_C} \leftarrow \text{KAS}.\text{MakeSecrets}(1^\kappa, \mathcal{P}_C)$
 - 4: $\text{Pub}_{\mathcal{P}_C} \leftarrow \text{KAS}.\text{MakePublicData}(1^\kappa, \mathcal{P}_C)$
 - 5: $\kappa_{\mathcal{P}_V} \leftarrow \text{KAS}.\text{MakeKeys}(1^\kappa, \mathcal{P}_V)$
 - 6: $\omega_{\mathcal{P}_V} \leftarrow \text{KAS}.\text{MakeSecrets}(1^\kappa, \mathcal{P}_V)$
 - 7: $\text{Pub}_{\mathcal{P}_V} \leftarrow \text{KAS}.\text{MakePublicData}(1^\kappa, \mathcal{P}_V)$
 - 8: Set $\text{PP} = (\text{PP}', \text{Pub}_{\mathcal{P}_C}, \text{Pub}_{\mathcal{P}_V})$
 - 9: Set $\text{MK} = (\text{MK}', \kappa_{\mathcal{P}_C}, \omega_{\mathcal{P}_C}, \kappa_{\mathcal{P}_V}, \omega_{\mathcal{P}_V})$
-

Alg. 2 $\text{FnInit}(F, \text{MK}, \text{PP}) \rightarrow PK_F$

- 1: $PK_F \leftarrow \text{RPVC}.\text{FnInit}(F, \text{MK}', \text{PP}')$
-

Alg. 3 $\text{Register}(\text{ID}, \lambda(\text{ID}), \text{MK}, \text{PP}) \rightarrow SK_{\text{ID}}$

- 1: **if** ID is a server **then**
 - 2: $SK'_{\text{ID}} \leftarrow \text{RPVC}.\text{Register}(\text{ID}, \text{MK}', \text{PP}')$
 - 3: $SK_{\text{ID}} = (SK'_{\text{ID}}, \kappa_{\lambda^C(\text{ID})}, \omega_{\lambda^C(\text{ID})})$
 - 4: **else if** ID is a delegator **then**
 - 5: $SK_{\text{ID}} = (\kappa_{\lambda^C(\text{ID})}, \omega_{\lambda^C(\text{ID})}, \kappa_{\lambda^V(\text{ID})}, \omega_{\lambda^V(\text{ID})})$
 - 6: **else**
 - 7: $SK_{\text{ID}} = (\kappa_{\lambda^V(\text{ID})}, \omega_{\lambda^V(\text{ID})})$
-

Alg. 4 $\text{Certify}(S, F, \text{MK}, \text{PP}) \rightarrow EK_{F,S}$

- 1: $EK_{F,S} \leftarrow \text{RPVC}.\text{Certify}(S, F, \text{MK}', \text{PP}')$
-

Alg. 5 $\text{ProbGen}(x, SK_C, PK_F, \lambda(C), \lambda(o), \text{PP}) \rightarrow (\sigma_o, VK_o, RK_o)$

- 1: $(\sigma'_o, VK_o, RK'_o) \leftarrow \text{RPVC}.\text{ProbGen}(x, PK_F, \text{PP}')$
 - 2: $\kappa_{\lambda^C(o)} \leftarrow \text{KAS}.\text{GetKey}(\lambda^C(C), \lambda^C(o), \omega_{\lambda^C(C)}, \text{PP})$
 - 3: $\kappa_{\lambda^V(o)} \leftarrow \text{KAS}.\text{GetKey}(\lambda^V(C), \lambda^V(o), \omega_{\lambda^V(C)}, \text{PP})$
 - 4: **if** $(\kappa_{\lambda^C(o)} \neq \perp)$ and $(\kappa_{\lambda^V(o)} \neq \perp)$ **then**
 - 5: $\sigma_o = (\lambda^C(o), \text{SE}.\text{Encrypt}(\sigma'_o, \kappa_{\lambda^C(o)}))$
 - 6: $RK_o = (\lambda^V(o), \text{SE}.\text{Encrypt}(RK'_o, \kappa_{\lambda^V(o)}))$
-

Alg. 6 $\text{Compute}(\sigma_o, EK_{F,S}, SK_S, \lambda^C(S), \lambda^C(o), \text{PP}) \rightarrow \theta_o$

- 1: Parse σ_o as $(\lambda^C(o), c)$
 - 2: $\kappa_{\lambda^C(o)} \leftarrow \text{KAS}.\text{GetKey}(\lambda^C(S), \lambda^C(o), \omega_{\lambda^C(S)}, \text{PP})$
 - 3: **if** $\kappa_{\lambda^C(o)} = \perp$ **then**
 - 4: **return** \perp
 - 5: **else**
 - 6: $\sigma'_o \leftarrow \text{SE}.\text{Decrypt}(c, \kappa_{\lambda^C(o)})$
 - 7: **if** $\sigma'_o = \perp$ **then return** \perp
 - 8: **else** $\theta_o \leftarrow \text{RPVC}.\text{Compute}(\sigma'_o, EK_{F,S}, SK_S, \text{PP}')$
-

Alg. 7 $\text{BVerify}(\theta_o, VK_o, \text{PP}) \rightarrow (RT_o, \tau_{\theta_o})$:

- 1: $(RT_o, \tau_{\theta_o}) \leftarrow \text{RPVC}.\text{BVerify}(\theta_o, VK_o, \text{PP}')$
-

of the S-KI game is useful as a proof technique besides the original motivation to reflect realistic attacks.

¹⁰We use $\kappa_{\mathcal{P}_C}$ to denote the set of all keys generated by the KAS for the computation poset \mathcal{P}_C and, for a label $\lambda^C(\text{ID}) \in \mathcal{P}_C$, the associated key is $\kappa_{\lambda^C(\text{ID})} \in \kappa_{\mathcal{P}_C}$.

Alg. 8 $\text{Retrieve}(RT_o, \tau_{\theta_o}, VK_o, RK_o, \lambda^V(V), \lambda^V(o), PP) \rightarrow \tilde{y}$:

```

1: Parse  $RK_o$  as  $(\lambda^V(o), e)$ 
2:  $\kappa_{\lambda^V(o)} \leftarrow \text{KAS.GetKey}(\lambda^V(V), \lambda^V(o), \omega_{\lambda^V(V)}, PP)$ 
3: if  $\kappa_{\lambda^V(o)} = \perp$  then
4:   return  $\perp$ 
5: else
6:    $RK'_o \leftarrow \text{SE.Decrypt}(e, \kappa_{\lambda^V(o)})$ 
7:    $\tilde{y} \leftarrow \text{RPVC.Retrieve}(RT_o, \tau_{\theta_o}, VK_o, RK'_o, PP')$ 

```

Alg. 9 $\text{Revoke}(\tau_{\theta_o}, F, MK, PP) \rightarrow \{EK_{F,S'}\}$ or \perp

```

1: return  $\text{RPVC.Revoke}(\tau_{\theta_o}, F, MK, PP')$ 

```

We now give a theorem and proof that the construction presented above is secure against the games presented in Section 5.

Theorem 1. *Given any Strong-Key-Indistinguishability secure KAS, any RPVC scheme secure in the sense of Public Verifiability, Revocation, and Blind Verification, and an authenticated symmetric encryption scheme secure in the sense of $\text{IND-CPA} \wedge \text{INT-PTXT}$, let \mathcal{PVC}_{AC} be the PVC-AC scheme defined in Algorithms 1-9. Then \mathcal{PVC}_{AC} is secure in the sense of Public Verifiability, Revocation, Blind Verification, Authorized Outsourcing, Authorized Computation, Weak Input Privacy, and Authorized Verification.*

Informally, the security proofs follow from the security of the underlying RPVC scheme and the S-KI and IND-CPA security properties. We now present a formal proof for Theorem 1.

6.3 Security Proofs

Lemma 1. *Given a secure RPVC scheme, an IND-CPA secure symmetric encryption scheme and a KAS secure against Key-Indistinguishability, let \mathcal{PVC}_{AC} be the PVC-AC scheme defined in Algorithms 1-10. Then \mathcal{PVC}_{AC} is secure in the sense of Authorized Outsourcing (Game 5).*

Proof. Let \mathcal{A}_{VC} be an adversary with non-negligible advantage δ in the Authorized Outsourcing game and let \mathcal{A}_{SE} be an adversary playing the IND-CPA game with a challenger \mathcal{C} against \mathcal{SE} . We first transition to a slightly modified version of the Authorized Outsourcing game, showing a negligible difference between the two. We can then use an adversary against this modified game to break the IND-CPA security of the symmetric encryption scheme.

- **Game 0.** This is the Authorized Outsourcing game as defined in Section 5.1.
- **Game 1.** This is identical to **Game 0**, except that we replace the key $\kappa_{\lambda^C(o)}$ for the challenge computation o with a random key κ^* drawn uniformly from the keyspace.

Game 0 to Game 1 This game hop relies on the Strong Key Indistinguishability (S-KI) of the KAS. Suppose an adversary \mathcal{A}_{VC} exists that can distinguish **Game 0** from **Game 1** with non-negligible advantage δ . Then there exists an adversary \mathcal{A}_{SE} that breaks the S-KI of the KAS also with advantage δ . \mathcal{A}_{VC} will play either **Game 0** or **Game 1** with \mathcal{A}_{SE} acting as the challenger, and must guess correctly which game he is playing. \mathcal{A}_{SE} in turn will play the S-KI game with a challenger \mathcal{C} . This reduction is important to ensure that no additional information about the encryption key is leaked through the construction.

1. \mathcal{C} begins by selecting a bit $b \xleftarrow{\$} \{0, 1\}$ which will determine whether the challenge key is real or random, and hence whether \mathcal{A}_{VC} plays **Game 0** or **Game 1**. \mathcal{C} continues by instantiating the KAS in lines 1 to 5 and giving the public information to \mathcal{A}_{SE} .

2. \mathcal{A}_{SE} runs lines 1 and 2 of the Authorized Outsourcing game, and sends \mathcal{A}_{VC} the generated public parameters. During the **Setup** algorithm in line 2, \mathcal{A}_{SE} will not run lines 2 to 4 of Algorithm 1. Instead, it will set $Pub_{\mathcal{P}_C}$ to be the challenge public information from the KI challenger \mathcal{C} and will make use of oracle queries to \mathcal{C} for the remaining parameters.
3. \mathcal{A}_{VC} is then given oracle access. Queries to the **FnInit**, **Certify** and **Revoke** functions can be answered simply by running the corresponding algorithm. Queries to the remaining algorithms may need to make use of KAS derived keys. \mathcal{A}_{SE} can use its knowledge of the verification poset and the corresponding KAS which it owns where appropriate but it does not know anything other than the public information for the computational poset, because this is the challenge poset for the S-KI game. Thus, \mathcal{A}_{SE} will issue **Corrupt** oracle queries to retrieve the necessary keys and secret information for the queried identity label. Since \mathcal{A}_{VC} has not yet chosen a challenge computation o , and \mathcal{A}_{SE} has not yet chosen a challenge node v^* , the **Corrupt** oracle will always return a valid key and secret pair, which \mathcal{A}_{SE} can use to form a full, valid response for \mathcal{A}_{VC} .
4. \mathcal{A}_{VC} will select a challenge security label $\lambda^C(\mathcal{A})$ for itself, which \mathcal{A}_{SE} will register him for using another **Corrupt** oracle query to \mathcal{C} – that is, it will query $\mathcal{O}^{\text{Corrupt}}(\lambda^C(\mathcal{A}), \perp)$ in the KI game. \mathcal{C} will add $\lambda^C(\mathcal{A})$ to the list Q and return $(\kappa_{\lambda^C(\mathcal{A})}, \omega_{\lambda^C(\mathcal{A})})$. Additionally, \mathcal{A}_{SE} may run $\text{RPVC.Register}(\text{PP}_{VC}, \text{MK}_{VC}, \mathcal{A})$ to create SK'_A . Finally, \mathcal{A}_{SE} sets $SK_A = (SK'_A, \kappa_{\lambda^C(\mathcal{A})}, \omega_{\lambda^C(\mathcal{A})})$ as well as updating the list to $L = L \cup \lambda^C(\mathcal{A})$.
5. \mathcal{A}_{VC} is again given oracle access which \mathcal{A}_{SE} can respond to as above. Eventually, \mathcal{A}_{VC} will output a challenge computation o which \mathcal{A}_{SE} will send to \mathcal{C} as the challenge node v^* for the KI game. Note that this is a valid challenge in the S-KI game since, in the Authorized Outsourcing game, \mathcal{A}_{VC} may not choose a challenge computation o for which he is authorized – that is, any o such that $\lambda^C(o) \leq \lambda^C(v_i)$ for any $\lambda^C(v_i)$ queried to the **Register** oracle. This corresponds precisely to the condition on the choice of v^* in the S-KI game. \mathcal{C} will return a key κ^* which corresponds either to the real key $\kappa_{\lambda^C(o)}$ or a random key from the keyspace chosen according to the bit b chosen at the beginning of the game.
6. \mathcal{A}_{SE} must now simulate a server and a verifier. It does this by first registering two such entities, S and V respectively, that are authorized to compute and verify o respectively. It sets $\lambda^C(S) = \lambda^C(o)$ and hence $\kappa_{\lambda^C(S)} = \kappa^*$. Note that this is the only part of SK_S that is required, and in particular $\omega_{\lambda^C(S)}$ is not required (as this is only needed for deriving keys, and S will only need to use κ^*). To register S , \mathcal{A}_{SE} sets $SK_S = (\kappa^*, \perp, \kappa_{\lambda^V(S)}, \omega_{\lambda^V(S)})$. Furthermore, to register V , \mathcal{A}_{SE} sets $SK_V = (\kappa_{\lambda^V(V)}, \omega_{\lambda^V(V)})$.
7. \mathcal{A}_{SE} then runs lines 12 and 13 as written in the Authorized Outsourcing game. It can also run the **Compute** algorithm as it knows $\kappa_{\lambda^C(o)}$ from $\kappa_{\lambda^C(C)}$, and **Verify** as it owns the verification KAS. \mathcal{A}_{SE} can therefore return the result of the game to \mathcal{A}_{VC} as expected.

Now, \mathcal{A}_{VC} has been provided all information that an adversary against the Authorized Outsourcing game would be given and will guess which game he is playing with advantage δ . Notice that the distribution of the game generated by \mathcal{A}_{SE} is precisely that of **Game 0** if $b = 0$ (i.e. the real KAS key is used), and is precisely that of **Game 1** otherwise, if a random key was chosen. Thus, \mathcal{A}_{SE} can simply forward \mathcal{A}_{VC} 's guess to \mathcal{C} as its guess in the S-KI game. Thus we conclude that if \mathcal{A}_{VC} guesses correctly with non-negligible advantage δ , then \mathcal{A}_{SE} can break the Strong-Key Indistinguishability of the KAS also with non-negligible advantage δ . Since we assumed the KAS was S-KI secure, such an adversary may not exist and hence **Game 0** is indistinguishable from **Game 1** except with at most a negligible advantage $\epsilon \leq 1 - \delta$.

Reduction to INT-PTXT We have shown that, from the adversary’s point of view, **Game 1** is almost (with negligible distinguishing advantage) identical to **Game 0**. Thus, we may run the adversary against **Game 1** instead with at most an ϵ loss in the tightness of the reduction. In essence, we have now removed any information leakage from the KAS. Suppose \mathcal{A}_{VC} is an adversary with non-negligible advantage δ in **Game 1**. We now show that we can construct an adversary \mathcal{A}_{SE} that breaks the INT-PTXT security of the authenticated symmetric encryption scheme \mathcal{SE} using \mathcal{A}_{VC} as a subroutine. Let \mathcal{C} be the INT-PTXT challenger for \mathcal{A}_{SE} who in turn acts as the challenger in **Game 1** for \mathcal{A}_{VC} .

1. \mathcal{C} begins by initializing the list L and running $\text{SE.KeyGen}(1^\kappa)$ to generate a key κ^* . It sends the security parameter 1^κ to \mathcal{A}_{SE} .
2. \mathcal{A}_{SE} must now initialize Game 1 for \mathcal{A}_{VC} . Informally, it will set the KAS key for the label $\lambda^C(o)$ to be the random key κ^* chosen by \mathcal{C} . However, the challenge label is unknown until \mathcal{A}_{VC} chooses it, whilst the public parameters and oracle access must be provided before this choice. Thus, we require \mathcal{A}_{SE} to guess the challenge label during **Setup** so that the correct key can be implicitly set to be the INT-PTXT challenge key (and all encryptions under that key can be formed using oracle access to \mathcal{C}). If the number of labels in the poset is N , where N is polynomial in the security parameter (as the scheme must be efficiently instantiable), then \mathcal{A}_{SE} may guess $\lambda^C(o)$ with probability at least $\frac{1}{N}$. Assuming that the guess is correct, we proceed as follows.
3. \mathcal{A}_{SE} runs $\text{PVC}_{AC}.\text{Setup}$ as given in Algorithm 1 with the modification that the key for the guessed label in the computation poset, $\kappa_{\lambda^C(o)}$ is implicitly set to be the key generated by \mathcal{C} in the IND-CPA game and the KAS is made to be consistent with this choice. Of course, \mathcal{A}_{SE} does not have this key but will ensure that any operations using it will be performed using its oracles to \mathcal{C} . Since the challenge key, and hence any corresponding secret derivation information, is unknown, it is not trivial to construct a KAS incorporating this key. However, notice that the Authorized Outsourcing game (and by extension **Game 1**) does not permit the adversary to query any label that is an ancestor of the challenge label in the computation poset. Thus, KAS keys for the set of ancestors will not be needed, and a KAS can simply be instantiated over the remaining nodes (and the public information for the ancestor set simulated, as the keys cannot be derived, the public information need not be functionally correct). Remaining keys can simply be generated using the security parameter. From the adversarial point of view, this will be indistinguishable from the real games.
4. \mathcal{A}_{VC} is given the generated public parameters and oracle access which \mathcal{A}_{SE} may respond to as follows:
 - **FnInit**, **Certify**, and **Revoke** queries can be handled by simply calling the relevant algorithm as these have no dependence on the KAS.
 - If a **Register** query is made for a label $\lambda(ID) \geq \lambda^C(o)$ as guessed by \mathcal{A}_{SE} then \mathcal{A}_{SE} aborts the game since \mathcal{A}_{VC} would not then be able to choose $\lambda^C(o)$ as its challenge computation, and hence \mathcal{A}_{SE} ’s guess was incorrect. Otherwise, \mathcal{A}_{SE} holds the relevant KAS keys and may respond by running as specified in Oracle Query 7.
 - By the oracle restrictions specified for the Authorized Outsourcing game, \mathcal{A}_{VC} may not query the **ProbGen** oracle on the challenge computation $\lambda^C(o)$. Hence, if such a query is made at this point, then $\lambda^C(o)$ would not be chosen as the challenge computation and \mathcal{A}_{SE} ’s guess would be incorrect, and so the game is aborted. For

any other choice of computation queried to the **ProbGen** oracle, \mathcal{A}_{SE} holds the KAS key (or generated symmetric key) and may run Algorithm 5 as written.

- Observe that if a query is made to the **Compute** oracle for the challenge computation, then the adversary has either submitted a malformed encoded input, and \perp should be returned, or the adversary has submitted a correctly formed encoded input, and hence has already won the game. For all other queries, \mathcal{A}_{SE} can use one of the keys it holds to respond to the challenge by running Algorithm 6.
5. \mathcal{A}_{VC} eventually outputs a choice of label which \mathcal{A}_{SE} can register him for, as in the oracle query above, and \mathcal{A}_{VC} is again given oracle access which is handled as before. It eventually outputs a challenge computation, and if this is not the computation chosen by \mathcal{A}_{SE} at the beginning of the game, then the game is aborted. Otherwise, the choice of computation label is valid since the game has not already been aborted during the oracle queries.
 6. \mathcal{A}_{SE} should now register two entities: a computational server S and a verifier V . However, these will not be required in the following, so \mathcal{A}_{SE} can simulate registering them with labels $\lambda^C(o)$ and $\lambda^V(o)$ respectively and update the public parameters accordingly without actually requiring the correct keys for these entities (as the adversary will not see any output from these entities other than their presence in the lists in the public parameters). \mathcal{A}_{SE} can also run the **FnInit** and **Certify** algorithms as written.
 7. \mathcal{A}_{VC} is then provided with all relevant information and is given oracle access again. Queries can be handled as above, but **Register** queries now return \perp if the queried label is an ancestor of the computation label in the poset i.e. exactly when \mathcal{A}_{SE} does not hold a KAS key for the queried label.
 8. \mathcal{A}_{VC} finally outputs a forged encoded input σ_o .

Now, observe that from the point of view of \mathcal{A}_{VC} , the game has been simulated correctly up to this point and that \mathcal{A}_{SE} has not made any queries to its **Encrypt** oracle in the INT-PTXT game. Thus, if the ciphertext c within σ_o decrypts successfully (which it must for \mathcal{A}_{VC} to win the Authorized Outsourcing game), then the verification oracle $\text{Ver}(c)$ will output 1 *and* the decrypted message has certainly not been queried to the **Encrypt** oracle, and thus \mathcal{A}_{SE} can forward c as its answer in the INT-PTXT game and win with exactly the advantage of \mathcal{A}_{VC} in the Authorized Outsourcing game. This is assuming that \mathcal{A}_{SE} correctly guessed the challenge computation label so that the game did not abort. Thus, for a poset of polynomial size N , the advantage of \mathcal{A}_{SE} is $\frac{\delta}{N}$ which is non-negligible if \mathcal{A}_{VC} has non-negligible advantage δ in the Authorized Outsourcing game. However, since we assumed the authenticated symmetric encryption scheme to be secure, such an adversary with non-negligible advantage against the Authorized Outsourcing game cannot exist.

Thus, we conclude, the overall advantage against the Authorized Outsourcing game is the sum of the distinguishing advantage between **Game 0** and **Game 1**, and the advantage in the reduction to INT-PTXT, both of which we have shown to be negligible. Therefore, the overall advantage against the Authorized Outsourcing game is negligible. □

We remark that in the above proof we required \mathcal{A}_{SE} to correctly guess the label that \mathcal{A}_{VC} would select ahead of time. This is very similar to the notions of security commonly found in functional encryption primitives, particularly Identity-based Encryption [8] and Attribute-based Encryption [18, 6]. In these settings, it is common to refer to the method of guessing the correct label as *complexity leveraging* which results in a polynomial loss in the tightness of the

reduction. An alternative, which could equally be taken here, is to formulate a weaker *selective* notion of security in which the adversary must select the challenge label at the beginning of the game before seeing the public parameters.

We also note that the restriction on queries to the **ProbGen** oracle (i.e. that the challenge computation cannot be queried) is stronger than required for our particular construction due to the **Encrypt** oracle available in the INT-PTXT game. An alternative would be to allow queries to **ProbGen** for the challenge computation but require that the final adversarial output is distinct from those given in response to these queries.

Lemma 2. *Given a secure RPVC scheme, an authenticated symmetric encryption scheme secure in the sense of IND-CPA \wedge INT-PTXT and a KAS secure against Strong Key-Indistinguishability, let \mathcal{PVC}_{AC} be the PVC-AC scheme defined in Algorithms 1-10. Then \mathcal{PVC}_{AC} is secure in the sense of Authorized Computation (Game 6).*

Proof. Let \mathcal{A}_{VC} be an adversary with non-negligible advantage δ in the Authorized Computation game and let \mathcal{A}_{SE} be an adversary playing the IND-CPA game with a challenger \mathcal{C} against \mathcal{SE} . We first transition to a slightly modified version of the Authorized Computation game, showing a negligible difference between the two. We can then use an adversary against this modified game to break the IND-CPA security of the symmetric encryption scheme.

- **Game 0.** This is the Authorized Computation game as defined in Section 5.2.
- **Game 1.** This is identical to **Game 0**, except that we replace the key $\kappa_{\lambda^{\mathcal{C}}(o)}$ for the challenge computation o with a random key κ^* drawn uniformly from the keyspace.

Game 0 to Game 1 This game hop relies on the Strong Key-Indistinguishability of the KAS and proceeds very similarly to the corresponding hop in the proof of Lemma 1. As such, we do not state the full proof here but leave it as an easy exercise. We conclude that if \mathcal{A} distinguishes **Game 0** from **Game 1** with non-negligible advantage δ , then an adversary can use \mathcal{A} as a subroutine to break the Strong Key-Indistinguishability of the KAS also with non-negligible advantage δ . Since the KAS is assumed S-KI secure, such an adversary may not exist and hence **Game 0** is indistinguishable from **Game 1** except with at most a negligible advantage $\epsilon \leq 1 - \delta$.

Reduction to IND-CPA We have shown that, from the adversary’s point of view, **Game 1** is almost (with negligible distinguishing advantage) identical to **Game 0**. Thus, we may run the adversary against **Game 1** instead to remove any information leakage from the KAS. Suppose \mathcal{A}_{VC} is an adversary with non-negligible advantage against **Game 1**. We show that we can construct an adversary \mathcal{A}_{SE} that breaks the IND-CPA security of the symmetric encryption scheme \mathcal{SE} using \mathcal{A}_{VC} as a subroutine. Let \mathcal{C} be the IND-CPA challenger for \mathcal{A}_{SE} who in turn acts as the challenger in **Game 1** for \mathcal{A}_{VC} .

1. \mathcal{C} begins by choosing a random bit b and running $\text{SE.KeyGen}(1^\kappa)$ to generate a key κ^* . It sends the security parameter 1^κ to \mathcal{A}_{SE} .
2. \mathcal{A}_{SE} must now initialize **Game 1** for \mathcal{A}_{VC} . Informally, it will set the KAS key for the label $\lambda^{\mathcal{C}}(o)$ to be the random key κ^* chosen by \mathcal{C} . However, the challenge label is unknown until \mathcal{A}_{VC} chooses it, whilst the public parameters and oracle access must be provided before this choice. Thus, we require \mathcal{A}_{SE} to guess the challenge label during **Setup** so that the correct key can be implicitly set to be the IND-CPA challenge key (and all encryptions under that key can be formed using \mathcal{C}). If the number of labels in the poset

is N , where N is polynomial in the security parameter (as the scheme must be efficiently instantiable), then \mathcal{A}_{SE} may guess $\lambda^C(o)$ with probability at least $\frac{1}{N}$. Assuming that the guess is correct, we proceed as follows.

3. \mathcal{A}_{SE} runs $\text{PVC}_{AC}.\text{Setup}$ as given in Algorithm 1 except that the key for the guessed label in the computation poset, $\kappa_{\lambda^C(o)}$ is implicitly set to be κ^* and the KAS is made consistent with this choice. Note that the Authorized Computation game (and by extension **Game 1**) does not permit the adversary to query any label that is an ancestor of the challenge label in the computation poset. Thus a KAS can be instantiated over the remaining nodes (and the public information for the ancestor set simulated – as the keys cannot be derived, the public information need not be functionally correct). Remaining keys can simply be generated using the security parameter. From the adversarial point of view, this will be indistinguishable from the real games.
4. \mathcal{A}_{VC} is given the generated public parameters and oracle access which \mathcal{A}_{SE} responds to as follows:
 - FnInit , Certify , and Revoke queries can be handled by simply calling the relevant algorithm.
 - If a Register query is made for a label $\lambda(ID) \geq \lambda^C(o)$ then \mathcal{A}_{SE} aborts the game since \mathcal{A}_{VC} would not then be able to choose $\lambda^C(o)$ as its challenge computation, and hence \mathcal{A}_{SE} 's guess was incorrect. Otherwise, \mathcal{A}_{SE} holds the relevant KAS keys and may respond by running Oracle Query 7.
 - ProbGen queries for a computation labelled $\lambda(o)$ can be handled by running Algorithm 5 with the exception that lines 2 and 5 are simulated as follows. Line 2 is not run at all as \mathcal{A}_{SE} may not hold the challenge key, and line 5 is run using an oracle query to the IND-CPA LoR oracle for the choice of messages $m_0 = m_1 = \sigma'_o$. For any other queried computation, \mathcal{A}_{SE} holds the KAS key (or generated symmetric key) and may run Algorithm 5 as written.
 - Observe that by the INT-PTXT property of the authenticated symmetric encryption scheme, \mathcal{A}_{VC} cannot form a valid ciphertext (one that will decrypt to anything other than \perp) without holding the encryption key. The encryption keys that \mathcal{A}_{VC} may hold are precisely those revealed through Register queries, and since queries may not be made for labels $\lambda(ID) \geq \lambda^C(o)$, each query will be for a label belonging to the KAS instantiated by \mathcal{A}_{SE} . Hence, if $\lambda(ID) \geq \lambda^C(o)$, \mathcal{A}_{SE} returns \perp and otherwise it uses its knowledge of the KAS to respond genuinely by running Algorithm 6.
5. \mathcal{A}_{VC} eventually outputs a choice of label which \mathcal{A}_{SE} can register him for in the same manner as in the oracle query above, and \mathcal{A}_{VC} is again given oracle access which is handled as before. It eventually outputs a challenge computation, and if this is not the computation chosen by \mathcal{A}_{SE} at the beginning of the game, then the game is aborted. Otherwise, the choice of computation label is valid since the game has not already been aborted during the oracle queries.
6. \mathcal{A}_{SE} should now register two entities: a delegator C and a verifier V . However, these will not be required in the following, so \mathcal{A}_{SE} may simulate registering them with labels $\lambda^C(o)$ and $\lambda^V(o)$ respectively and update the public parameters accordingly without actually requiring the correct keys for these entities (as the adversary will not see any output from these entities other than their presence in the lists in the public parameters). \mathcal{A}_{SE} can also run the FnInit and Certify algorithms as written.

7. \mathcal{A}_{SE} must now run **ProbGen** for the challenge computation o to generate an encoded input σ_o . To do so, it will make use of the LoR oracle provided by \mathcal{C} in the IND-CPA game. \mathcal{A}_{SE} will run lines 1, 3 and 4 as written to generate a problem encoding σ'_o . It sets $m_0 = \sigma'_o$ and chooses another message m_1 of the same length uniformly at random from the message space. These are queried to the LoR oracle which will return the encryption of message m_b for the challenger's random choice of b . \mathcal{A}_{SE} can then use this response to form σ_o and form RK_o using the verification KAS key which it holds.
8. All relevant information is passed to \mathcal{A}_{VC} who is also given oracle access. This is again handled in the same manner as before. Eventually, \mathcal{A}_{VC} will produce an encoded output θ_o .

Observe that if the challenger's random bit $b = 0$, then this is precisely **Game 1** and \mathcal{A}_{VC} wins with non-negligible probability δ . If, on the other hand, $b = 1$ then \mathcal{A}_{VC} is provided with the encryption of a random message which does not relate to the computation at all. In this case, the best that \mathcal{A}_{VC} can do is to either attack the verification key, which would break the Public Verifiability of the RPVC scheme (even having seen the verification key, as it is public, a server should not be able to persuade a verifier of an incorrect result), or to guess the correct θ_o . Using the RPVC scheme of Alderman et al. [2], attacking the verification key would amount to inverting the one-way function, whilst guessing θ_o would involve guessing the encrypted plaintext contained in σ'_o (which \mathcal{A}_{VC} never sees). Thus, if $b = 1$, then the advantage of \mathcal{A}_{VC} is $\epsilon = \max \left\{ Adv_{\mathcal{A}_{VC}}^{Invert}(OWF, g, 1^\kappa), \frac{1}{2^{|m|}} \right\}$ where m is the message encrypted in σ'_o and is of length $|m| = poly(\kappa)$, both of which are negligible.

Thus, \mathcal{A}_{SE} may run the verification process using the keys that he owns in the verification KAS. If verification succeeds, \mathcal{A}_{SE} may output a guess $b' = 0$, and guess $b' = 1$ otherwise. If $b = 0$, then this guess is correct with non-negligible probability δ . If $b = 1$, then this guess is correct with non-negligible probability $1 - \epsilon$. Thus, if \mathcal{A}_{SE} guessed the correct label and hence did not abort the game, \mathcal{A}_{SE} breaks the IND-CPA security of the symmetric encryption scheme with non-negligible probability. However, as we assumed that \mathcal{SE} was $\text{IND-CPA} \wedge \text{INT-PTXT}$ secure, this is a contradiction and hence an adversary with non-negligible advantage against **Game 1** cannot exist. Thus, we conclude, the overall advantage against the Authorized Computation game is the sum of the distinguishing advantage between **Game 0** and **Game 1**, and the advantage in the reduction to IND-CPA, both of which we have shown to be negligible. Therefore, the overall advantage against the Authorized Computation game is negligible. \square

Lemma 3. *Given a secure RPVC scheme, a KAS secure in the sense of Strong-Key Indistinguishability and an authenticated symmetric encryption scheme secure in the sense of $\text{IND-CPA} \wedge \text{INT-PTXT}$, let \mathcal{PVC}_{AC} be the PVC-AC scheme defined in Algorithms 1-10. Then \mathcal{PVC}_{AC} is secure in the sense of Authorized Verification (Game 7).*

Proof. Suppose \mathcal{A}_{VC} is an adversary with non-negligible advantage δ in the Authorized Verification game and \mathcal{A}_{SE} is an adversary playing the IND-CPA game with a challenger \mathcal{C} against \mathcal{SE} . We first define the following modified game and show that an adversary has negligible distinguishing advantage between the two. We can therefore employ an adversary against this modified game to break the IND-CPA security of the symmetric encryption scheme.

- **Game 0.** This is the Authorized Verification game as defined in Section 5.3.
- **Game 1.** This is identical to **Game 0**, except that we replace the key $\kappa_{\lambda^V(o)}$ for the challenge computation o in the verification poset with a random key κ^* drawn uniformly from the keyspace.

Game 0 to Game 1 This transition relies on the Strong Key-Indistinguishability of the KAS. The proof is very similar to that in the proof of Lemma 1, and so we leave this as an easy exercise. If \mathcal{A} can distinguish **Game 0** from **Game 1** with non-negligible advantage δ , then an adversary using \mathcal{A} as a subroutine can break the Strong Key-Indistinguishability of the KAS with the same non-negligible advantage δ . However, as the KAS is assumed S-KI secure, such an adversary may not exist and **Game 0** is indistinguishable from **Game 1** except with at most a negligible advantage $\epsilon \leq 1 - \delta$.

Reduction to IND-CPA Suppose \mathcal{A}_{VC} is an adversary with non-negligible advantage against the Authorized Verification game. We now show that we can construct an adversary \mathcal{A}_{SE} that breaks the IND-CPA security of the symmetric encryption scheme \mathcal{SE} using \mathcal{A}_{VC} as a subroutine. From the adversary's point of view, **Game 1** is indistinguishable from **Game 0** with negligible distinguishing advantage. Therefore, \mathcal{A}_{SE} may simulate **Game 1** instead to remove any information leakage from the KAS. Let \mathcal{C} be the IND-CPA challenger for \mathcal{A}_{SE} who in turn acts as the challenger in **Game 1** for \mathcal{A}_{VC} .

1. \mathcal{C} first chooses a bit β uniformly at random and generates a key $\kappa^* \leftarrow \text{SE.KeyGen}(1^\kappa)$. It sends the security parameter 1^κ to \mathcal{A}_{SE} .
2. \mathcal{A}_{SE} must now initialize **Game 1** for \mathcal{A}_{VC} . It will implicitly set the KAS key for the label $\lambda^V(o)$ to be κ^* . However, as \mathcal{A}_{VC} has not yet chosen his challenge labels, \mathcal{A}_{SE} must guess one of these labels in order to generate the public parameters and provide oracle access. \mathcal{A}_{SE} may guess $\lambda^V(o)$ with probability at least $\frac{2}{N}$ (as \mathcal{A}_{VC} will select two labels, and \mathcal{A}_{SE} must match one of them) for a poset of size N nodes. Assuming that the guess is correct, we proceed as follows.
3. \mathcal{A}_{SE} runs $\text{PVC}_{AC}.\text{Setup}$ as given in Algorithm 1 except that the key for the guessed label in the verification poset, $\kappa_{\lambda^V(o)}$ is implicitly set to be κ^* and the KAS is made consistent with this choice. Note that in the Authorized Verification game (and by extension **Game 1**) the adversary may not query any label that is an ancestor of the challenge label $\lambda^V(o_b)$ in the verification poset. Thus \mathcal{A}_{SE} can instantiate a KAS over the remaining nodes and simulate the public information for the ancestor set – as keys for this set cannot be derived, the public information need not be functionally correct. \mathcal{A}_{SE} can also use the security parameter to generate remaining keys in this set. From the adversarial point of view, this is indistinguishable from the real games.
4. \mathcal{A}_{SE} sends the resulting public parameters to \mathcal{A}_{VC} and provides oracle access as follows:
 - For queries to Fnlit , Certify , and Revoke \mathcal{A}_{SE} can simply call the relevant algorithm.
 - If a Register query is made for a label $\lambda(ID) \geq \lambda^V(o)$ then \mathcal{A}_{SE} aborts the game as \mathcal{A}_{VC} will not be able to choose $\lambda^V(o)$ as one of its challenge computations, and hence \mathcal{A}_{SE} 's guess was incorrect. Otherwise, \mathcal{A}_{SE} holds the relevant KAS keys and may respond by running Oracle Query 8.
 - A ProbGen query for a computation labelled $\lambda^V(o)$ can be handled by running Algorithm 5 with the exception of lines 3 and 6 which are simulated as follows. \mathcal{A}_{SE} makes an oracle query to the IND-CPA LoR oracle for the message pair $m_0 = m_1 = RK'_o$. For all other queries, \mathcal{A}_{SE} holds the correct keys and may honestly run Algorithm 5.
 - Compute queries can be handled by running Algorithm 6 since it relies only on the computation poset which is owned by \mathcal{A}_{SE} .

5. \mathcal{A}_{VC} eventually outputs a choice of label which \mathcal{A}_{SE} can register him for as per the Register oracle query above, and \mathcal{A}_{VC} is again given oracle access which is handled as above. Eventually it chooses two challenge computations o_0 and o_1 and, if neither is the same as that chosen earlier by \mathcal{A}_{SE} , the game is aborted. Otherwise, the choices are valid since the game has not already been aborted during the oracle queries. Instead of choosing the bit b at random, it can be set such that o_b corresponds to \mathcal{A}_{SE} 's guess of challenge computation.
6. \mathcal{A}_{SE} now simulates registering two entities: a delegator C and a verifier V . However, as these will not be required in the following, \mathcal{A}_{SE} may simulate registering them with labels $\lambda^C(o_b)$ and $\lambda^V(o_b)$ respectively and update the public parameters accordingly without requiring valid keys for these entities (as the adversary will not see any output from these entities other than being listed in the public parameters). \mathcal{A}_{SE} can also run the Fnlit and Certify algorithms as written.
7. \mathcal{A}_{SE} now runs ProbGen for the challenge computation o_b to encode σ_{o_b} . It runs lines 1, 2,4 and 5 as written to generate a problem encoding σ_o and retrieval key RK'_o . It sets $m_0 = RK'_o$ and randomly chooses another message m_1 of the same length from the message space. These are given as input to the LoR oracle which returns the encryption of message m_β for the challenger's random choice of β . Using this response, \mathcal{A}_{SE} can create RK_o .
8. \mathcal{A}_{SE} must finally run Compute on the resulting encoded input, which it can do honestly as the algorithm does not rely on the verification poset. \mathcal{A}_{VC} then receives all relevant information and oracle access, which is again handled as before. Eventually, \mathcal{A}_{VC} outputs a guess b' of b .

Observe that if \mathcal{A}_{SE} guessed correctly and $\beta = 0$, then this is precisely **Game 1**, and \mathcal{A}_{VC} wins with non-negligible probability δ . Thus \mathcal{A}_{SE} wins with non-negligible probability $\frac{2\delta}{N}$. If $\beta = 1$, on the other hand, the encoded input provided to \mathcal{A}_{VC} is the encryption of a random message completely unrelated to the retrieval key. In this case, by the Blind Verification property of the underlying RPVC construction, \mathcal{A}_{VC} may only have a negligible advantage ϵ (over random guessing) at learning the result of the computation. Thus, if \mathcal{A}_{VC} 's output $b' = b$, then \mathcal{A}_{SE} should output a guess of $\beta' = 0$, and otherwise should guess $\beta' = 1$. If $\beta = 0$ then \mathcal{A}_{SE} wins with probability $\frac{2\delta}{N}$, and if $\beta = 1$ then \mathcal{A}_{SE} wins with probability $1 - \epsilon$. Thus, \mathcal{A}_{SE} wins in both cases with non-negligible probability.

The overall advantage against the Authorized Verification game is the sum of the distinguishing advantage between **Game 0** and **Game 1**, and the advantage in the reduction to IND-CPA, both of which we have shown to be negligible. Therefore, the overall advantage against the Authorized Verification game is negligible. \square

Lemma 4. *Given a secure RPVC scheme, a Key-Indistinguishability secure KAS and an IND-CPA secure symmetric encryption scheme, let \mathcal{PVC}_{AC} be the PVC-AC scheme defined in Algorithms 1-10. Then \mathcal{PVC}_{AC} is secure in the sense of Weak Input Privacy (Game 8).*

Proof. Assume that \mathcal{A}_{VC} is an adversary with non-negligible advantage δ in the Authorized Computation game. We show that we can use this to construct an adversary, \mathcal{A}_{KI} , with non-negligible advantage in the Key-Indistinguishability game. Let \mathcal{C} be the challenger for \mathcal{A}_{KI} and let \mathcal{A}_{KI} act as the challenger for \mathcal{A}_{VC} .

1. \mathcal{C} begins by running

$$\kappa_{\mathcal{P}_C} \leftarrow \text{MakeKeys}(1^\kappa, \mathcal{P}_C),$$

$\omega_{\mathcal{P}_C} \leftarrow \text{MakeSecrets}(1^\kappa, \mathcal{P}_C)$, and

$\text{Pub}_{\mathcal{P}_C} \leftarrow \text{MakePublicData}(1^\kappa, \mathcal{P}_C)$.

It initialises the list Q to be empty and sends $\text{Pub}_{\mathcal{P}_C}$ to \mathcal{A}_{KI} .

2. \mathcal{A}_{KI} first sets $L = \epsilon$ and $o = \perp$. Then it runs lines 5 to 8 of Algorithm ?? and sets $\text{MK}_{AC} = (\perp, \perp, \kappa_{\mathcal{P}_V}, \omega_{\mathcal{P}_V})$. It also runs Algorithm 1 as given, and sends PP_{VC} and PP_{AC} to \mathcal{A}_{VC} .
3. Now, \mathcal{A}_{VC} is provided with oracle access which \mathcal{A}_{KI} can simulate as follows:
 - (a) $\text{FnInit}(\text{PP}_{VC}, \text{MK}_{VC}, \cdot, \cdot)$, $\text{Certify}(\text{PP}_{VC}, \text{MK}_{VC}, \cdot, \cdot, \cdot)$ and $\text{Revoke}(\text{MK}_{VC}, \cdot, \cdot, \cdot)$ can all be run as specified in the construction as these rely on the underlying RPVC scheme only.
 - (b) $\text{Register}(\text{PP}_{VC}, \text{MK}_{VC}, \text{MK}_{AC}, \cdot, \cdot)$: \mathcal{A}_{VC} will query for an identity ID and associated security label $\lambda(ID)$. \mathcal{A}_{KI} will follow Algorithm 3 but will make use of **Corrupt** oracle queries to \mathcal{C} whenever it requires a value of $\kappa_{\lambda^C(ID)}$ or $\omega_{\lambda^C(ID)}$. That is, \mathcal{A}_{KI} will issue an oracle query $\mathcal{O}^{\text{Corrupt}}(\lambda^C(ID), \perp)$, to which \mathcal{C} will respond by adding $\lambda^C(ID)$ to the query list Q and returning $(\kappa_{\lambda^C(ID)}, \omega_{\lambda^C(ID)})$ as well as updating $L = L \cup \lambda^C(ID)$. Note that keys and secrets for verification policies are owned by \mathcal{A}_{KI} already.

\mathcal{A}_{VC} will eventually output a choice of computational label $\lambda^C(\mathcal{A}_{VC})$.

4. To register \mathcal{A}_{VC} as a server with this security label, \mathcal{A}_{KI} will make use of its **Corrupt** oracle queries in the Key Indistinguishability game – that is, it will query $\mathcal{O}^{\text{Corrupt}}(\lambda^C(\mathcal{A}_{VC}), \perp)$ in Game 1. \mathcal{C} will add $\lambda^C(\mathcal{A}_{VC})$ to the list Q and return the relevant key and secret information $(\kappa_{\lambda^C(\mathcal{A}_{VC})}, \omega_{\lambda^C(\mathcal{A}_{VC})})$. Additionally, \mathcal{A}_{KI} updates the list to $L = L \cup \lambda^C(\mathcal{A}_{VC})$, runs $\text{RPVC.Register}(\text{PP}_{VC}, \text{MK}_{VC}, \mathcal{A}_{VC})$ and outputs $SK'_{\mathcal{A}_{VC}}$. Finally, \mathcal{A}_{KI} forwards $SK_{\mathcal{A}_{VC}} = (SK'_{\mathcal{A}_{VC}}, \kappa_{\lambda^C(\mathcal{A}_{VC})}, \omega_{\lambda^C(\mathcal{A}_{VC})})$ to \mathcal{A}_{VC} .
5. \mathcal{C} now chooses a computation o uniformly at random from the space of all possible computations of F that the adversary is not authorised to compute i.e. we require $\lambda^C(o) \not\leq \lambda^C(\mathcal{A}_{VC})$. Note that the choice of o will dictate the input data x . If not such o may be found then the game ends with a loss for the adversary as they did not select a valid attack identity.
6. \mathcal{A}_{KI} sends \mathcal{C} its choice of challenge security label in the Key Indistinguishability game $v^* = \lambda^C(o)$. Notice that this choice is allowed since $\lambda^C(o)$ is not a descendent of any security labels given in a **Corrupt** oracle query (those labels added to Q). This is necessarily so since the only **Corrupt** queries thus far have been for the adversary's label $\lambda^C(\mathcal{A})$ in Step ??, and those labels given in a **Register** query by \mathcal{A}_{VC} . Now, by the condition in Step ?? we know $\lambda^C(o)$ is certainly not less than or equal to $\lambda^C(\mathcal{A})$. Also, by the restriction on \mathcal{A}_{VC} 's oracle queries, it is not allowed to query for any $\lambda^C(ID) \geq \lambda^C(o)$ as this would constitute a trivial win.

\mathcal{C} chooses a random bit b and returns $\kappa_{\lambda^C(o)}$ if $b = 0$ or chooses a random key from the keyspace otherwise, and sends this key, κ^* , to \mathcal{A}_{KI} .
7. \mathcal{A}_{KI} must now simulate a client and a verifier. It does this by first registering two such entities, C and V respectively, that are authorised to compute and verify o respectively. It sets $\lambda^C(C) = \lambda^C(o)$ and hence $\kappa_{\lambda^C(C)} = \kappa^*$. Note that this is the only part of SK_C that

is required, and in particular $\omega_{\lambda^C(C)}$ is not required (as this is only needed for deriving keys, and S will only need to use κ^*).

To register C , \mathcal{A}_{KI} sets $SK_C = (\kappa^*, \perp, \kappa_{\lambda^V(C)}, \omega_{\lambda^V(C)})$. Furthermore, to register V , \mathcal{A}_{KI} sets $SK_V = (\kappa_{\lambda^V(V)}, \omega_{\lambda^V(V)})$.

8. \mathcal{A}_{KI} runs

$$\begin{aligned} (PK_F, L_F) &\leftarrow \text{RPVC.FnInit}(\text{PP}_{VC}, \text{MK}_{VC}, F), \\ (EK_{F,S}, L_F) &\leftarrow \text{RPVC.Certify}(\text{PP}_{VC}, \text{MK}_{VC}, F, L_F, S), \\ (\sigma'_x, VK_{F,x}, RK'_{F,x}) &\leftarrow \text{RPVC.ProbGen}(x, PK_F) \text{ for } x \text{ specified in } o, \\ \sigma_x &= (\lambda^C(o), \text{SE.Encrypt}(\sigma'_x, \kappa^*)) \\ RK_{F,x} &= (\lambda^V(o), \text{SE.Encrypt}(RK'_{F,x}, \kappa_{\lambda^V(o)})) \end{aligned}$$

\mathcal{A}_{KI} sends $(\sigma_y, VK_{F,x}, \lambda(o), SK_A, PK_F, L_F)$ to \mathcal{A}_{VC} .

9. \mathcal{A}_{VC} again gets oracle access which \mathcal{A}_{KI} responds to as follows:

- (a) As before $\text{FnInit}(\text{PP}_{VC}, \text{MK}_{VC}, \cdot, \cdot)$, $\text{Certify}(\text{PP}_{VC}, \text{MK}_{VC}, \cdot, \cdot, \cdot)$ and $\text{Revoke}(\text{MK}_{VC}, \cdot, \cdot, \cdot)$ can all be run as given in the construction.
- (b) $\text{Register}(\text{PP}_{VC}, \text{MK}_{VC}, \text{MK}_{AC}, \cdot, \cdot)$: \mathcal{A}_{VC} queries for an identity ID and associated security label $\lambda(ID)$. \mathcal{A}_{KI} follows Algorithm 3 making use of Corrupt oracle queries to \mathcal{C} whenever it requires a value of $\kappa_{\lambda^C(ID)}$ or $\omega_{\lambda^C(ID)}$. That is, \mathcal{A}_{KI} will issue an oracle query $\mathcal{O}^{\text{Corrupt}}(\lambda^C(ID), \lambda^C(o))$. \mathcal{C} will respond by adding $\lambda^C(ID)$ to the query list Q and returning $(\kappa_{\lambda^C(ID)}, \omega_{\lambda^C(ID)})$ if and only if $\lambda^C(ID) \not\cong \lambda^C(o)$ and updating the list L accordingly. \mathcal{A}_{KI} can respond with verification policy keys and secrets since it already owns them.

\mathcal{A}_{VC} eventually outputs a guess, x' , for x .

10. By the IND-CPA property of the symmetric encryption scheme we hide the attributes attached to the intermediate encodings in ProbGen . If $b = 0$ then κ^* is a real key corresponding to $\lambda^C(o)$ and therefore the encryption of σ'_x will be performed using a valid key and will look correct to \mathcal{A}_{VC} . Therefore, to \mathcal{A}_{VC} the system looks consistent and correct, and hence he will have non-negligible advantage δ to guess x correctly, as we assumed. If $x' = x$, \mathcal{A}_{KI} will output a guess $b' = 0$.

Otherwise, κ^* is a random key and the encryption of the intermediate encoded input σ'_x will not be correct, so \mathcal{A}_{VC} may not possess the assumed advantage and \mathcal{A}_{KI} should output a guess $b' = 1$.

We conclude that \mathcal{A}_{KI} at least succeeds whenever \mathcal{A}_{VC} succeeds i.e. with non-negligible probability δ . However, our construction assumes that the KAS is secure in the sense of Key Indistinguishability, and hence \mathcal{A}_{VC} cannot have a non-negligible success probability. \square

We conclude that combining the results of the Lemmas 1 to 4 proves Theorem 1.

7 Conclusion

We have motivated the need for the cryptographic enforcement of access control policies in the setting of outsourced computation, particularly in the multi-user setting. As developments in VC continue towards such settings, it is vital to enable restrictions to be placed on: the computations that delegators can outsource (both from the perspective of separation of duties, and considering a server providing differing levels of service for different users); the computations a server may perform (such that certain computations, over sensitive data say, may only be performed by a server satisfying a policy); and the verifiers that may learn the output of the result (e.g. ensuring that read access to the newly generated data is handled consistently with the sensitivity of the inputs). We have shown example graph-based access control policies for these scenarios, as well as providing a formal definitional framework, security models, and a provably secure construction built from Key Assignment Schemes.

Future work will consider alternate enforcement mechanisms, such as authentication protocols that enforce graph-based authorisation policies to achieve ticket-based access control to a computational service, or the use of dual-policy ABE to outsource computations (using the KP-ABE policy) and enforcing access control (using the CP-ABE policy). Specific VC scenarios may lead to interesting access control models. One particularly applicable setting for the enforcement of access control policies is *verifiable searchable encryption*. Consider a remote database host that returns verifiably correct results to user queries (computations). In practice it is unlikely that all users should have unrestricted access to the entire database. It is imperative that only authorized users may perform specific queries (those relating solely to their duties and to data for which they have clearance) and that results remain protected to prevent data leakage.

Acknowledgements

We thank Christopher Dearlove, Rachel Player and Gordon Procter for helpful discussions. The first author acknowledges support from BAE Systems Advanced Technology Centre under a CASE Award.

This research was partially sponsored by US Army Research laboratory and the UK Ministry of Defence under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the U.S. Government, the UK Ministry of Defense, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

- [1] S. G. Akl and P. D. Taylor. Cryptographic solution to a problem of access control in a hierarchy. *ACM Trans. Comput. Syst.*, 1(3):239–248, 1983.
- [2] J. Alderman, C. Cid, J. Crampton, and C. Janson. Revocation in publicly verifiable outsourced computation. Cryptology ePrint Archive, Report 2014/640, 2014. <http://eprint.iacr.org/>.

- [3] M. J. Atallah, M. Blanton, N. Fazio, and K. B. Frikken. Dynamic and efficient key management for access hierarchies. *ACM Trans. Inf. Syst. Secur.*, 12(3), 2009.
- [4] D. Bell and L. LaPadula. Secure computer systems: Mathematical foundations. Technical Report MTR-2547, MITRE Corporation, 1973.
- [5] M. Bellare and C. Namprempre. Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. *J. Cryptology*, 21(4):469–491, 2008.
- [6] J. Bethencourt, A. Sahai, and B. Waters. Ciphertext-policy attribute-based encryption. In *IEEE Symposium on Security and Privacy*, pages 321–334. IEEE Computer Society, 2007.
- [7] M. A. Bishop. *Computer Security. Art and Science*. Addison-Wesley Professional, 2002.
- [8] D. Boneh and M. K. Franklin. Identity-based encryption from the weil pairing. In J. Kilian, editor, *CRYPTO*, volume 2139 of *Lecture Notes in Computer Science*, pages 213–229. Springer, 2001.
- [9] A. Castiglione, A. D. Santis, and B. Masucci. Key indistinguishability vs. strong key indistinguishability for hierarchical key assignment schemes. Cryptology ePrint Archive, Report 2014/752, 2014. <http://eprint.iacr.org/>.
- [10] S. G. Choi, J. Katz, R. Kumaresan, and C. Cid. Multi-client non-interactive verifiable computation. In *TCC*, pages 499–518, 2013.
- [11] M. Clear and C. McGoldrick. Policy-based non-interactive outsourcing of computation using multikey FHE and CP-ABE. In P. Samarati, editor, *SECRYPT*, pages 444–452. SciTePress, 2013.
- [12] J. Crampton. Cryptographic enforcement of role-based access control. In P. Degano, S. Etalle, and J. D. Guttman, editors, *Formal Aspects in Security and Trust*, volume 6561 of *Lecture Notes in Computer Science*, pages 191–205. Springer, 2010.
- [13] J. Crampton, K. M. Martin, and P. R. Wild. On key assignment for hierarchical access control. In *CSFW*, pages 98–111. IEEE Computer Society, 2006.
- [14] A. L. Ferrara, G. Fuchsbauer, and B. Warinschi. Cryptographically enforced RBAC. In *2013 IEEE 26th Computer Security Foundations Symposium, New Orleans, LA, USA, June 26-28, 2013*, pages 115–129. IEEE, 2013.
- [15] E. S. V. Freire, K. G. Paterson, and B. Poettering. Simple, efficient and strongly KI-secure hierarchical key assignment schemes. In E. Dawson, editor, *Topics in Cryptology - CT-RSA 2013 - The Cryptographers' Track at the RSA Conference 2013, San Francisco, CA, USA, February 25-March 1, 2013. Proceedings*, volume 7779 of *Lecture Notes in Computer Science*, pages 101–114. Springer, 2013.
- [16] R. Gennaro, C. Gentry, and B. Parno. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In T. Rabin, editor, *CRYPTO*, volume 6223 of *Lecture Notes in Computer Science*, pages 465–482. Springer, 2010.
- [17] Google. Google Compute Engine – Cloud Computing & IaaS – Google Cloud Platform. <http://cloud.google.com/compute/>, 2014. [Online; accessed 23-October-2014].

- [18] V. Goyal, O. Pandey, A. Sahai, and B. Waters. Attribute-based encryption for fine-grained access control of encrypted data. In A. Juels, R. N. Wright, and S. D. C. di Vimercati, editors, *ACM Conference on Computer and Communications Security*, pages 89–98. ACM, 2006.
- [19] R. Ostrovsky, A. Sahai, and B. Waters. Attribute-based encryption with non-monotonic access structures. In P. Ning, S. D. C. di Vimercati, and P. F. Syverson, editors, *ACM Conference on Computer and Communications Security*, pages 195–203. ACM, 2007.
- [20] B. Parno, M. Raykova, and V. Vaikuntanathan. How to delegate and verify in public: Verifiable computation from attribute-based encryption. In R. Cramer, editor, *TCC*, volume 7194 of *Lecture Notes in Computer Science*, pages 422–439. Springer, 2012.
- [21] M. A. Rappa. The utility business model and the future of computing services. *IBM Systems Journal*, 43(1):32–42, 2004.
- [22] R. S. Sandhu and P. Samarati. Access control: principle and practice. *Communications Magazine, IEEE*, 32(9):40–48, 1994.
- [23] L. Xu and S. Tang. Verifiable computation with access control in cloud computing. *The Journal of Supercomputing*, 69(2):528–546, 2014.