

Constructing Key Assignment Schemes from Chain Partitions

Jason Crampton, Rosli Daud and Keith M. Martin

Technical Report
RHUL-MA-2010-10
16 April 2010



Department of Mathematics
Royal Holloway, University of London
Egham, Surrey TW20 0EX, England

<http://www.rhul.ac.uk/mathematics/techreports>

Abstract

In considering a problem in access control for scalable multimedia formats, we have developed new methods for constructing a key assignment scheme. Our first contribution is to improve an existing cryptographic access control mechanism for scalable multimedia formats. We then show how our methods can be applied to a chain partition to develop alternative mechanisms for scalable multimedia formats and how these methods can themselves be extended to create a new type of key assignment scheme.

1 Introduction

Scalable multimedia formats, such as MPEG-4 [15] and JPEG2000 [6], consist of two components: a non-scalable base component and a scalable enhancement component. Decoding the base component will yield low quality results. The quality of the decoded data can be improved by decoding the enhancement component as well as the base component. The enhancement component may comprise multiple “orthogonal” layers, orthogonal in the sense that each layer controls a distinct aspect of the quality of the encoded content. The MPEG-4 FGS (fine granularity scalability) format [15], for example, has a bit-rate layer and a peak signal-to-noise ratio (PSNR) layer.

Zhu *et al.* proposed a layered access control scheme for MPEG-4 FGS called SMLFE (scalable multi-layer FGS encryption) [16]. The purpose of SMLFE is to provide different end-users with access to the same content at different levels of quality (by controlling access to the enhancement component).

SMLFE assumes that each enhancement frame is decomposed into different *segments*, each of which is associated with some bit-rate level and some PSNR level. In other words, the *enhancement component stream* (a sequence of enhancement frames) is split into a number of distinct segment streams. Each of these segment streams is encrypted with a different key, and the ability of an end-user (or, more accurately, the decoder available to the end-user) to reconstruct the enhancement component is determined by the keys that are accessible to the user.

However, SMLFE had a number of inadequacies and subsequent research sought to address these deficiencies [11, 12, 13]. This later research uses a labeling technique, which associates each segment with a k -tuple and then uses iterative hashing to derive key components for each segment. Most of these labeling schemes suffer from the distinct disadvantage that different users can combine their respective key components to derive keys for which

no single user is authorized. The one exception [12] uses a very complicated labeling process that makes it very difficult to reason about the properties of the scheme (including whether it is secure against colluding users or not). Our first contribution is to construct a labeling scheme that can be proved to be secure against colluding users and has other significant advantages over existing schemes. We discuss labeling schemes in Sec. 3.

We then consider alternative approaches to the problem of layered access control for scalable multimedia formats. Our second contribution is to define several schemes in Sec. 4 that make use of chain partitions. One of our constructions makes use of the labeling scheme we introduce in Sec. 3. The constructions in Sec. 4 have demonstrable advantages, in the context of layered access control, over labeling schemes and existing approaches to cryptographic access control.

It can be shown that the enforcement of layered access control for scalable multimedia formats can be regarded as an instance of a *key assignment scheme*. Such schemes are used to enforce a no-read-up information flow policy using cryptographic techniques. A recent survey of such schemes proposed a classification into four generic types of scheme [7]. These schemes offer different trade-offs in terms of the amount of storage required and the complexity of key derivation. Our final contribution is to show that the schemes in Sec. 4 can be generalized to create new types of generic key assignment schemes. These generic schemes offer different trade-offs from existing schemes, which may prove useful for certain applications.

We conclude the paper with some suggestions for future work. Before proceeding further, we introduce some relevant background material.

2 Background

In this section, we first recall some relevant concepts from mathematics and cryptography. The section concludes with a more formal statement of the problem of layered access control and a discussion of its relationship to work on key assignment schemes.

2.1 Definitions and Notation

A *partially ordered set* (or *poset*) is a pair (X, \leq) , where \leq is a reflexive, anti-symmetric, transitive binary relation on X . X is a *total order* (or *chain*) if for all $x, y \in X$, either $x \leq y$ or $y \leq x$. We say $A \subseteq X$ is an *antichain* if for all $x, y \in A$, $x \not\leq y$ and $x \not\geq y$. We may write $y < x$ if $y \leq x$ and $y \neq x$, and we may write $x \geq y$ if $y \leq x$.

The (directed, acyclic) graph (X, \leq) would include all “reflexive edges” and all “transitive edges”, so it is customary to represent a poset using a smaller set of edges. We say x covers y , denoted $y \triangleleft x$, if $y < x$ and there does not exist $z \in X$ such that $y < z < x$. Then the *Hasse diagram* of a poset (X, \leq) is defined to be the (directed, acyclic) graph (X, \triangleleft) [8]. A simple Hasse diagram is shown in Fig. 1(a). Note that all edges in the diagram are assumed to be directed upwards.

A *partition* of a set X is a collection of sets $\{Y_1, \dots, Y_k\}$ such that (i) $Y_i \subseteq X$ (ii) $Y_1 \cup \dots \cup Y_k = X$, and (iii) $Y_i \cap Y_j = \emptyset$ if and only if $i = j$. The *greatest common divisor* of x and y is written $\gcd(x, y)$; we say x and y are *co-prime* if $\gcd(x, y) = 1$.

We assume the existence of an *RSA key generator* [14], a randomized algorithm that takes a security parameter k as input and outputs a triple (N, e, d) such that:

- $N = pq$, where p and q are distinct odd primes;
- $e \in \mathbb{Z}_{\phi(N)}^*$, where $\phi(N) = (p - 1)(q - 1)$, $e > 1$, and $\gcd(e, \phi(N)) = 1$;
- $d \in \mathbb{Z}_{\phi(N)}^*$, where $ed \equiv 1 \pmod{\phi(N)}$.

Let $h : \{0, 1\}^* \rightarrow \{0, 1\}^\ell$ be a hash function and let $k \geq 0$ be an integer. Then we define the *iterative hash function* $h^k : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$ in the following way:

$$h^k(x) = \begin{cases} x & \text{if } k = 0, \\ h(h^{k-1}(x)) & \text{otherwise.} \end{cases}$$

2.2 Key Assignment Schemes

We now rephrase the problem at hand in more formal terms and consider in general terms how this problem is related to existing work on *key assignment schemes*. Let us assume that we are concerned with a scalable multimedia format with two distinct layers (such as bit-rate and PSNR), containing m and n levels respectively.

Define $R_{m,n} = \{(x, y) : 1 \leq x \leq m, 1 \leq y \leq n\}$ and define $(x_1, y_1) \leq (x_2, y_2)$ if and only if $x_1 \leq x_2$ and $y_1 \leq y_2$. Then $(R_{m,n}, \leq)$ is a partially ordered set. Each segment (and segment stream) represents a distinct protected object and is labeled with a pair (i, j) indicating the corresponding levels in the bit-rate and PSNR layers, respectively. Each pair $(i, j) \in R_{m,n}$ is associated with an encryption key $\kappa_{i,j}$. Segment streams are encrypted with the corresponding key. Each user is authorized to access layered multimedia

of some quality $q_{i,j}$, which implies that such a user must be able to compute $\kappa_{x,y}$ for all $x \leq i$ and all $y \leq j$ in order to decode the relevant segment streams. Figure 1(a) illustrates the poset $R_{3,4}$.




Figure 1: $R_{3,4}$: an example of a poset used in layered access control for scalable multimedia formats and its canonical decomposition into chains

Clearly the access control requirements described above closely resemble the “no-read-up” component of an *information flow policy* [5, 9]. There are many schemes in the literature for enforcing an information flow policy using cryptographic techniques (see the survey paper of Crampton *et al.* [7], for example). Given a security lattice (L, \leq) , a set of subjects U , a set of protected objects O , and a security function $\lambda : U \cup O \rightarrow L$, we define a set of cryptographic keys $\{\kappa(x) : x \in L\}$. Then, adopting a cryptographic approach to policy enforcement, we encrypt object o with (symmetric) key $\kappa(\lambda(o))$. In order to correctly implement the information flow policy, a user u with security label should be given, or be able to derive, $\kappa(y)$ for all $y \leq x$. There are several generic approaches: one is to give u the set of keys $\{\kappa(y) : y \leq \lambda(u)\}$. More commonly, we give u a single key $\kappa(\lambda(u))$ and publish additional information that enables the user to derive $\kappa(y)$ whenever $y < \lambda(u)$.

Any cryptographic enforcement scheme for an information flow policy should satisfy two criteria.

- The scheme is *correct* if for all y , $\kappa(x)$ can be derived from $\kappa(y)$ and the public information if $y \geq x$.¹

¹It should be emphasized here that “derived from” means “derived from in a feasible amount of time”. Very few cryptographic schemes provide unconditional security in an information-theoretic sense; rather, they guarantee with a high probability that a scheme is secure against an adversary with reasonable resources. The interested reader is referred to the literature for a more detailed discussion of these issues [14].

- The scheme is *collusion secure* if, for all $x \in X$ and all $Y \subseteq X$ such that for all $y \in Y$, $y \not\geq x$, it is not possible to derive $\kappa(x)$ from $\{\kappa(y) : y \in Y\}$ and the public information. Note that this definition includes the case of a singleton subset Y , which corresponds to a single user “colluding” to recover a key for which she is not authorized.²

Clearly, the problem of enforcing layered accessed control for scalable multimedia formats can be addressed by defining an appropriate key assignment scheme for the partially ordered set $(R_{m,n}, \leq)$. However, because of the particularly simple structure of $R_{m,n}$, in the next two sections we consider some key assignment schemes that are tailored to the problem of layered access control for scalable multimedia formats. In Sec. 3, we consider *labeling schemes*, in which each key is defined by a set of *key components*, each of which is obtained by iteratively hashing some secret value. In Sec. 4, we consider some alternative approaches using chain partitions of $R_{m,n}$.

3 A New Labeling Scheme for Layered Access Control

Apart from SMLFE [16], all existing schemes for layered access control (to our knowledge) associate a distinct k -tuple with each element of $R_{m,n}$ [11, 12, 13]. This k -tuple is used to construct k key components using iterative hashing. We write $\phi(x, y) \in \mathbb{Z}^k$ to denote the label assigned to $(x, y) \in R_{m,n}$ and we write $\phi_i(x, y)$ to denote the i th co-ordinate of $\phi(x, y)$. In this section, we first summarize the basic technique and then describe our new labeling scheme and compare it to existing work.

First we introduce some additional definitions. Let $a = (a_1, \dots, a_k)$ and $b = (b_1, \dots, b_k)$ be elements of \mathbb{Z}^k . Then we define $(a_1, \dots, a_k) \leq (b_1, \dots, b_k)$ (in \mathbb{Z}^k) if and only if $a_i \leq b_i$ for all i , and we define $a - b = (a_1 - b_1, \dots, a_k - b_k)$. We say a is *positive* if $a_i \geq 0$ for all i .

Labeling schemes have the property that $(x, y) \geq (x', y')$ in $R_{m,n}$ if and only if $\phi(x', y') - \phi(x, y)$ is positive. It is this property that ensures the

²Recent work has introduced the notions of *key recovery* and *key indistinguishability* [2]. A proof that a scheme is secure against key recovery is analogous to proving that a scheme is collusion secure. The main difference is that collusion security assumes that colluding users will try to compute a key using the particular methods of key derivation associated with the scheme, whereas a proof of security against key recovery establishes that the recovery of a key is as difficult as solving some known hard problem. While formal security proofs of this nature are certainly important in modern cryptographic research, space constraints mean they are out of scope for this paper.

correctness of each scheme, since $\phi(x', y') - \phi(x, y)$ is used to construct k secrets per node using iterative hashing.

The content provider, hereafter called the *scheme administrator*, chooses a hash function $h : \{0, 1\}^* \rightarrow \{0, 1\}^\ell$ and k secrets $\sigma_1, \dots, \sigma_k \in \mathbb{Z}^\ell$. Then the secret $\sigma_{x,y}$ assigned to $(x, y) \in R_{m,n}$ comprises k key components:

$$\sigma_{x,y} \stackrel{\text{def}}{=} (h^{\phi_1(x,y)}(\sigma_1), \dots, h^{\phi_k(x,y)}(\sigma_k)).$$

For brevity, we may abuse notation and write $h^{\phi(x,y)}(\sigma)$ to denote $\sigma_{x,y}$. We define the key assigned to (x, y) to be

$$\kappa_{x,y} \stackrel{\text{def}}{=} h(h^{\phi_1(x,y)}(\sigma_1) \parallel \dots \parallel h^{\phi_k(x,y)}(\sigma_k)),$$

where $s_1 \parallel s_2$ denotes the concatenation of s_1 and s_2 .³ Again, we may abuse notation and write $h(\sigma_{x,y})$ to denote $\kappa_{x,y}$.

Correctness. By construction $(x, y) \geq (x', y')$ if and only if $\phi_i(x', y') - \phi_i(x, y)$ is positive. Now the i th key component of $\kappa_{x,y}$ is $h^{\phi_i(x,y)}(\sigma_i)$ and the i th key component of $\kappa_{x',y'}$ is $h^{\phi_i(x',y')}(\sigma_i)$. Hence, if $(x, y) \geq (x', y')$, then we simply hash the i th component of $\kappa_{x,y}$ a total of $\phi_i(x', y') - \phi_i(x, y)$ times to obtain the i th key component of $\sigma_{x',y'}$. Conversely, if $(x, y) \not\geq (x', y')$, then for some i , $\phi_i(x', y') - \phi_i(x, y) < 0$, which implies that we can only obtain the i th component of $\kappa_{x,y}$ by inverting h , which is computationally infeasible provided h is chosen appropriately.

The IWFK-1 scheme [13, §3.1.1], for example, simply defines $\phi(x, y)$ for $(x, y) \in R_{m,n}$ to be $(m - x, n - y)$. So, for example, $\phi(2, 4) = (1, 0)$ and $\phi(1, 1) = (2, 3)$ in $R_{3,4}$. Then $\sigma_{2,4} = (h(\sigma_1), \sigma_2)$ and $\sigma_{1,1} = (h^2(\sigma_1), h^3(\sigma_2))$. Hence, $\sigma_{2,4}$ can be used to derive $\sigma_{1,1}$ by hashing the first component of $\sigma_{2,4}$ once and hashing the second component twice.

Collusion Security. It is known that all but one of the schemes in the literature are not collusion secure. Indeed, it is trivial to find examples that break each of the schemes: in the IWFK-1 scheme for $R_{3,4}$, for example, $\sigma_{2,4} = (h(\sigma_1), \sigma_2)$ and $\sigma_{3,3} = (\sigma_1, h(\sigma_2))$; clearly these keys can be combined to recover $(\sigma_1, \sigma_2) = \sigma_{3,4}$. The IFAK scheme is claimed to be collusion secure [12], although no proof of this claim is given.

³The schemes in the literature simply define the “key” associated with (x, y) to be the concatenation of the key components. We take the hash of the concatenation of those components to make the distinction between key and key components clearer. It also means that we have fixed-length, short symmetric keys, determined by the size of the h ’s output.

3.1 The CDM Scheme

We now explain how our scheme works, which we call the CDM scheme for ease of reference.

Definition 1 Let $(x, y) \in R_{m,n}$. Then we define the CDM label of (x, y) to be

$$\phi_{\text{CDM}}(x, y) \stackrel{\text{def}}{=} (\underbrace{n - y, \dots, n - y}_x, \underbrace{n, \dots, n}_{m-x}).$$

Henceforth, we will simply write $\phi(x, y)$ to denote the CDM labeling of $(x, y) \in R_{m,n}$. Note the CDM labeling has m components. We now state several elementary results concerning the properties of the CDM labeling.

Proposition 2 Let $(x, y), (x', y') \in R_{m,n}$. Then $\phi(x', y') - \phi(x, y)$ is positive if and only if $(x, y) \geq (x', y')$.

Proof First assume that $(x, y) \geq (x', y')$ and consider $\phi(x', y') - \phi(x, y)$. By definition, $\phi(x', y') = (n - y', \dots, n - y', n, \dots, n)$ and $\phi(x, y) = (n - y, \dots, n - y, n, \dots, n)$. We consider three cases:

1. For all $i \leq x'$, $\phi_i(x, y) = n - y \leq n - y' = \phi_i(x', y')$ since $y' \leq y$.
2. For all $x' < i \leq x$, $\phi_i(x, y) = n - y \leq n = \phi_i(x', y')$.
3. Finally, for all $i > x$, $\phi_i(x, y) = n = \phi_i(x', y')$.

In other words, $\phi_i(x, y) - \phi_i(x', y') \geq 0$ for all i .

Now suppose that $(x, y) \not\geq (x', y')$. Then either $x \not\geq x'$ or $y \not\geq y'$ (or both). In the case that $x \not\geq x'$, we have that $x < x'$ and hence $\phi_i(x, y) = n > n - y' = \phi_i(x', y')$, $x < i \leq x'$. Hence, $\phi(x', y') - \phi(x, y)$ is not positive. In the case that $y \not\geq y'$, we have that $y < y'$ and $\phi_1(x, y) = n - y > n - y' = \phi_1(x', y')$; hence $\phi(x', y') - \phi(x, y)$ is not positive. In summary, $(x, y) \not\geq (x', y')$ implies that $\phi(x, y) - \phi(x', y')$ is not positive. The result follows. ■

Proposition 3 Let $(x, y), (x', y') \in R_{m,n}$ such that $(x, y) \geq (x', y')$. Then $\sigma_{x', y'}$ can be derived from $\sigma_{x, y}$ using precisely $xy - x'y'$ hash computations.

Proof We know that $\phi(x', y') - \phi(x, y)$ is positive and, by construction, the number of hash computations required to derive $\sigma_{x', y'}$ from $\sigma_{x, y}$ is given by the sum

$$\sum_{i=1}^m \phi_i(x', y') - \phi_i(x, y),$$

where

$$\phi_i(x', y') - \phi_i(x, y) = \begin{cases} (n - y') - (n - y) & 1 \leq i \leq x', \\ n - (n - y) & x' < i \leq x, \\ n - n & \text{otherwise.} \end{cases}$$

Hence, the number of hash computations required is

$$x'(y - y') + (x - x')y = xy - x'y'.$$

■

Corollary 4 *The number of hash computations required is bounded by $mn - 1$.*

We now give some intuition behind the labeling and an example. The element $(x, y) \in R_{m,n}$ defines a sub-rectangle $R_{x,y}$. Removing $R_{x,y}$ truncates the first i chains and leaves the remaining chains intact. Our labeling simply records the lengths of the chains that are left following the removal of $R_{x,y}$. Hence, for example, $\phi(2, 4) = (0, 0, 4)$ and $\phi(1, 1) = (3, 4, 4)$. Note that $3 + 4 = 7$ operations are required to derive $\kappa_{1,1}$ from $\kappa_{2,4}$ (as we would expect from Proposition 3).

The geometric intuition behind the scheme also provides some understanding of why our scheme is collusion secure.

Proposition 5 *Let $(x_1, y_1), \dots, (x_j, y_j) \in R_{m,n}$ such that $(x_i, y_i) \not\geq (x, y)$ for all i . Then there exists t , $1 \leq t \leq m$, such that $\phi_t(x, y) < \phi_t(x_i, y_i)$ for all i .*

Proof First note that we may assume $x_i \neq x_j$ if $i \neq j$. (If they are not, then $x_i = x_j = z$ and $(z, y_i) < (z, y_j)$ or $(z, y_i) > (z, y_j)$ or $(z, y_i) = (z, y_j)$. In each case we can omit one or other of (x_i, y_i) or (x_j, y_j) .) Hence, we may assume, without loss of generality, that $x_1 < \dots < x_j$. Hence, $y_1 > \dots > y_j$ (otherwise, $(x_i, y_i) < (x_j, y_j)$ and we can omit (x_i, y_i) since κ_{x_j, y_j} can derive κ_{x_i, y_i}). In the worst case $j = m$. Then $x = x_t$ for some t and $y > y_t$ (otherwise $(x_t, y_t) \geq (x, y)$).⁴ Then $\phi_t(x, y) = n - y < n - y_t$. Moreover, $\phi_t(x_i, y_i) = n > n - y_t$ for $i < t$ and $\phi_t(x_i, y_i) \geq n - y_t > \phi_t(x, y)$ for $i \geq t$. ■

Hence, no set of m colluding users can recover the t th component of $\sigma_{x,y}$. In other words, we have the following corollary.

Corollary 6 *The CDM scheme is collusion secure.*

⁴The proof can easily be modified if we do not make the assumption that $j = m$, but the exposition is easier with this assumption. Moreover, the assumption has the effect of revealing as much information as possible to colluding users.

3.2 Related Work

Table 1 provides a summary of the four schemes in the literature for layered access control (IWFK-1 [13, §3.1.1], IWFK-2 [13, §3.2.3], IFAK [12], and HIFK [11]), presented in chronological order and identified by the initial letters of the authors' surnames. Each component of $\sigma_{x,y} = h^{\phi(x,y)}(\sigma)$ is a distinct secret key component, as each component has to be hashed independently of the others. Hence, we believe it is appropriate to minimize the number of key components and the number of derivation steps that are required. The table reports precise storage requirements (given by the number of key components k) and worst case derivation (in terms of the number of hash computations required).

All of these schemes are correct, but only the IFAK scheme is claimed to be collusion secure, in the sense that a set of collaborating users cannot combine the secret components of their respective keys (and possibly use iterative hashing) to derive a key for which no one of them was authorized.

Scheme	k	Key derivation	Collusion secure
IWFK-1	2	$m + n - 2$	N
IWFK-2	3	$m + 2n - 3$	N
IFAK	$m + n - 1$	$\frac{1}{2}(m + n - 2)(m + n - 1)$	Y
HIFK	3	$2m + 2n - 4$	N
CDM	m	$mn - 1$	Y

Table 1: A summary of labeling schemes for layered access control

The characteristics of our scheme are shown in the last row of the table. Our scheme is collusion secure under the same assumptions that the IFAK scheme is (claimed to be) secure. However, we use a smaller value of k and we require fewer derivation steps. Moreover, we have a systematic and easily implementable way of generating our labels (unlike the IFAK scheme); because of this we can also compute the number of derivation steps required for any $(x, y), (x', y') \in R_{m,n}$ and prove that our scheme is collusion secure. IFAK, in contrast, has an extremely complicated labeling scheme, which makes it difficult to reason about (i) the number of derivation steps required in the general case (ii) the collusion security of the scheme.

4 New Schemes for Layered Access Control

In this section, we propose a number of key assignment schemes for implementing layered access control for scalable multimedia formats. These schemes assume that the poset $(R_{m,n}, \leq)$ has been partitioned into chains. Dilworth's Theorem [10] asserts that every partially ordered set (X, \leq) can be partitioned into w chains, where w is the *width* of X .⁵

Evidently, there are many different ways to partition the poset $R_{m,n}$ into chains, but we choose a particular partition that enables us to define two very simple schemes. We assume without loss of generality that $m \leq n$, and we define the *canonical* partition of $R_{m,n}$ into chains to be $\{C_1, \dots, C_m\}$, where $C_i = \{(i, j) : 1 \leq j \leq n\}$. Figure 1(b) illustrates the canonical partition of $R_{3,4}$ into chains.

4.1 Schemes with No Public Information

Generally, key assignment schemes rely on public information for key derivation [7]. An interesting feature of the schemes in the previous section is that no public information is used. In this section we consider two different schemes that require no public information: one in which the user can derive a key iteratively, the other in which the user derives keys directly. The difference between the two schemes lies in the cryptographic primitives that are used.

4.1.1 A Scheme Based on Hash Functions.

The scheme administrator first selects a family of m hash functions $h_i : \{0, 1\}^* \rightarrow \{0, 1\}^\ell$, $1 \leq i \leq m$. The scheme administrator also selects m secret values, $\sigma_1, \dots, \sigma_m \in \{0, 1\}^\ell$, where σ_i is associated with chain C_i . The scheme administrator then computes a secret key for each element in $R_{m,n}$, where $\kappa_{i,j}$ is defined to be $h_i^{n-j}(\sigma_i)$. Then a user authorized for content quality $q_{i,j}$ is given the keys $\{\kappa_{x,j} : 1 \leq x \leq i\}$. For reasons that will be apparent from the above description, we call this a *multiple-key iterated hash scheme*.

Correctness. We first show that a user can derive all keys for which she is authorized. Suppose that a user is authorized for quality $q_{i,j}$. (Equivalently, the user is associated with label $(i, j) \in R_{m,n}$.) Henceforth, we will simply write $u_{i,j}$ for such a user. Then $u_{i,j}$ must be able to derive all keys in the

⁵The width of X is the cardinality of the largest antichain in X . Clearly, any partition into chains must contain at least w chains. It is harder to prove that no more than w are required.

rectangle $R_{i,j}$. Now, by construction, $u_{i,j}$ has $\kappa_{x,j}$ for all $x \leq i$. Moreover, $\kappa_{x,y} = h_x^{j-y}(\kappa_{x,j})$, $1 \leq y < j$. Hence, a user $u_{i,j}$ can derive any key in $R_{i,j}$ in no more than $j - 1$ steps.

Collusion Security. Any “good” hash function will have the property that it is computationally hard to compute x given $y = h(x)$ (that is, *pre-image resistance*). Since keys are obtained by successively hashing elements in a chain, it is computationally hard to recover $\kappa_{i,j+1}$ from $\kappa_{i,j}$, since this would require the computation of the pre-image of $\kappa_{i,j}$. Hence, a user certainly cannot use a key from one key chain to derive a key higher up the same key chain (and hence for which she is not authorized), providing the scheme administrator chooses a suitable hash function. However, a user may have several keys: assuming that the key chains are independent – in the sense that knowledge of an element in C_i provides no information about any element in C_j , for all $j \neq i$ – then it is not possible for the user to derive any keys for which she is not authorized. We have chosen a different hash function for each chain in order to provide this key chain independence.

If two or more users collude – equivalently, if an adversary is able to obtain the keys of several users – then the set of keys available do not correspond to the nodes of a sub-rectangle (as they do for a single user). Suppose that an adversary (whether it is a group of colluding users or a single malicious entity) collectively has the keys $\kappa_{1,j_1}, \dots, \kappa_{m,j_m}$. Then κ_{i,j_i} cannot be used to recover κ_{i,j_i+k} for any $k > 0$ if h_i has pre-image resistance. Hence, assuming the independence of key chains, as before, we see that such an adversary has no additional advantage over a single user.

4.1.2 A Scheme Based on RSA.

In this scheme, we make use of a special case of the Akl-Taylor scheme [1], which can be applied to any poset. Specifically, we apply the scheme to each of the chains in the partition.

The scheme administrator first obtains m large compound integers N_1, \dots, N_m using an RSA key generator and makes these values public. For each chain C_i , the scheme administrator:

- chooses a secret $\sigma_i \in \mathbb{Z}_{N_i}^*$, such that for all σ_i and σ_j are co-prime if $i \neq j$;
- defines $\kappa_{i,j} = (\sigma_i)^{2^{n-j}} \bmod N_i$.

We call the sequence of keys

$$\kappa_{i,n} = (\sigma_i)^1, \kappa_{i,n-1} = (\sigma_i)^2 \bmod N_i, \dots, \kappa_{i,1} = (\sigma_i)^{2^{n-1}} \bmod N_i$$

an *RSA key chain*. As before, user $u_{i,j}$ is given the keys $\{\kappa_{x,j} : 1 \leq x \leq i\}$. Henceforth, for reasons of clarity and brevity, we will write x rather than $x \bmod N_i$, when N_i is clear from context.

Correctness and Collusion Security. Key derivation is quite different using RSA key chains. To obtain $\kappa_{x,y}$, where $x < i$ and $y < j$, the user selects $\kappa_{x,j}$ and then computes

$$(\kappa_{x,j})^{2^{j-y}} = (\kappa_{x,j})^{\frac{2^{n-y}}{2^{n-j}}} = ((\sigma_x)^{2^{n-j}})^{\frac{2^{n-y}}{2^{n-j}}} = (\sigma_x)^{2^{n-y}} = \kappa_{x,y}$$

To illustrate, consider Fig. 1(b) and suppose that the keys for C_2 are

$$\kappa_{2,4} = \sigma_2, \quad \kappa_{2,3} = \sigma_2^2, \quad \kappa_{2,2} = \sigma_2^4, \quad \kappa_{2,1} = \sigma_2^8.$$

Suppose we wish to derive $\kappa_{2,1}$ and we have $\kappa_{2,3}$. Then we compute

$$(\kappa_{2,3})^{2^{3-1}} = \kappa_{2,3}^4 = (\sigma_2^2)^4 = \sigma_2^8 = \kappa_{2,1}.$$

However, user with key $\kappa_{i,j}$ cannot derive $\kappa_{i,y}$ if $y > j$, since this would require the user to solve the *RSA problem*.⁶ Similarly, no collection of keys that includes $\kappa_{i,j}$ (but no key higher up the i th chain) can be used to derive $\kappa_{i,y}$.

4.2 Schemes with Single Keys

Most key assignment schemes in the literature require the end-user to store a single key. The multiple-key schemes described above clearly do not satisfy this criterion.

In this section, we describe schemes that only require the end user to store a single key. The trade-off is that such schemes require a certain amount of public information.

4.2.1 A Scheme Based on Hash Functions.

The scheme we now describe could be considered to be a hybrid of an iterative key encrypting key assignment scheme [7] and a hash chain. Atallah *et al.*, for example, define a concrete construction of an iterative key encrypting scheme [2].

In our scheme, the content provider selects m hash functions h_1, \dots, h_m and m secrets $\sigma_1, \dots, \sigma_m$, and defines key $\kappa_{i,j} = h_i^{n-j}(s_i)$, as before. Now,

⁶That is, given N , $y \in \mathbb{Z}_N^*$ and an integer $e > 0$ that is co-prime to $\phi(N)$, compute $y^{1/e} \bmod N$.

however, the content provider publishes enough information to enable the computation of $\kappa_{x,j}$ from $\kappa_{i,j}$ for all $x < i$, by publishing

$$\{\text{Enc}_{\kappa_{i,j}}(\kappa_{i-1,j}) : 1 < i \leq m, 1 \leq j \leq n\}.$$

Hence, we require $(m-1)n$ items of public information.

Correctness and Collusion Security. Again, it is very easy to demonstrate that a user $u_{i,j}$ can derive the key for any node in $R_{i,j}$. First, $u_{i,j}$ is given $\kappa_{i,j}$ and this key, in conjunction with the public information, can be used to derive $\kappa_{x,j}$ for all $x < i$. Moreover, $\kappa_{x,y}$ can be obtained from $\kappa_{x,j}$ by $j-y$ applications of h . Hence, $u_{i,j}$ can obtain $\kappa_{x,y}$ in no more than $i-1+j-1=i+j-2$ steps.

Collusion security follows from the fact that pre-image resistance of the hash function prevents the computation of $\kappa_{i,j+k}$ from $\kappa_{i,j}$ for any $k > 0$. The assumption that it is computationally hard to decrypt without knowledge of the secret key ensures that $\kappa_{i+k,j}$ cannot be derived from $\kappa_{i,j}$.

Trade-Offs. With additional public information, we can reduce the number of derivation steps further. We could, for example, publish information that enables the derivation of $\kappa_{x,j}$ from $\kappa_{i,j}$ in a single step, rather than $i-x$ steps. This requires an “edge” between every pair of nodes in each chain of length m ; that is, $\frac{1}{2}(m-1)mn$, increasing the public storage requirements by a factor of $m/2$. However, key derivation is reduced to no more than $(n-1)+1=n$ steps in the worst case.

At the extreme, we could publish information that always enables derivation in a single step, yielding a *direct key encrypting key assignment scheme* [7]. More complex trade-offs are also possible: Atallah *et al.*, for example, show that it is possible to define a scheme for a chain of m elements in which key derivation takes no more than two steps and $m \log_2 m$ items of public information are required [3], which increases public information by a multiplicative factor of $\log_2 m$.

4.2.2 A Scheme Based on RSA.

Finally, we note that we can use the CDM labeling (Definition 1) to construct a scheme with direct derivation for all keys. It is important to note that this scheme does not rely on the idea of encrypting edges, and is therefore quite different from the schemes described above.

Recall that we associate each $(x,y) \in R_{m,n}$ with a CDM label $\phi(x,y) \in \mathbb{Z}^m$. Moreover, $\phi(x',y') - \phi(x,y)$ is positive if and only if $(x,y) \geq (x',y')$.

In this new scheme the scheme administrator

- chooses primes $p_1, \dots, p_m \in \mathbb{Z}_N^*$ and makes them public;
- chooses a master secret $\sigma \in \mathbb{Z}_N^*$;
- defines

$$\pi(x, y) = \prod_{i=1}^m p_i^{\phi_i(x, y)};$$

- defines $\kappa_{x, y} = \sigma^{\pi(x, y)} \bmod N$.

Now let $(x, y) \geq (x', y')$. Then $\phi(x', y') - \phi(x, y)$ is positive and

$$\frac{\pi(x', y')}{\pi(x, y)} = \prod_{i=1}^m p_i^{\phi_i(x', y') - \phi_i(x, y)}$$

Hence,

$$(\kappa_{x, y})^{\frac{\pi(x', y')}{\pi(x, y)}} = (\sigma^{\pi(x, y)})^{\frac{\pi(x', y')}{\pi(x, y)}} = \sigma^{\pi(x', y')} = \kappa_{x', y'}$$

In other words, if $\phi(x', y') - \phi(x, y)$ is positive, we can compute $\kappa_{x', y'}$ from $\kappa_{x, y}$. Specifically, given $\kappa_{x, y}$:

1. compute $\phi(x, y)$ and $\phi(x', y')$, which is trivial if m and n are known;
2. compute $\phi(x', y') - \phi(x, y)$ and hence $\pi(x', y')/\pi(x, y)$;
3. finally, compute $\kappa_{x', y'}$.

We cannot compute $\kappa_{x'', y''}$ from $\kappa_{x, y}$ if $(x, y) \not\geq (x'', y'')$ since this would imply that $\phi(x'', y'') - \phi(x, y)$ is not positive and we would have to compute integral roots modulo N to compute $\kappa_{x'', y''}$. (In other words, solve the RSA problem.) Moreover, Proposition 5 implies that any adversary with keys $\kappa_{x_1, y_1}, \dots, \kappa_{x_j, y_j}$, such that $(x_i, y_i) \not\geq (x, y)$, would have to solve the RSA problem to compute $\kappa_{x, y}$.

4.3 Related Work

In Table 2, we summarize the properties of several schemes in the literature and compare them to the schemes we have introduced in this section. For our schemes, we report worst case private storage and derivation steps. The table includes, for ease of reference, the best labeling scheme from Sec. 3. We also include two generic key assignment schemes – IKE (iterative key encrypting) and DKE (direct key encrypting) [7] – which can be applied to

any poset and require the user to manage a single secret key. These schemes are recognized as having the most desirable balance between private storage, public storage and key derivation complexity. We write MKIH to denote the multiple-key iterative hash scheme and MKRSA to denote the multiple-key RSA scheme and replace ‘M’ with ‘S’ for the single-key analogues.

We write T_{Hsh} to denote the time take to compute a hash function, T_{Dec} to denote the time taken to decrypt an item of public information, and T_{Exp} to denote the time take to perform a modular exponentiation. We would expect that $T_{\text{Exp}} > T_{\text{Dec}} > T_{\text{Hsh}}$ for cryptographic primitives with similar levels of security. The table reports worst case storage costs and derivation times.

Scheme	Private storage	Public storage	Key derivation
CDM	m	0	$(mn - 1)T_{\text{Hsh}}$
IKE	1	$(m - 1)n + m(n - 1)$	$(m + n - 2)T_{\text{Dec}}$
DKE	1	$\frac{1}{4}mn((m + 1)(n + 1) - 4)$	T_{Dec}
MKIH	m	0	$(m - 1)T_{\text{Hsh}}$
MKRSA	m	0	T_{Exp}
SKIH	1	$(m - 1)n$	$(m - 1)T_{\text{Hsh}} + (n - 1)T_{\text{Dec}}$
SKRSA	1	m	T_{Exp}

Table 2: A summary of related work and a comparison with our schemes

It is clear that even the best labeling scheme (CDM) does not compare well with either the generic schemes in the literature or the schemes we have introduced in this section. The main reason for this is that the key components in the labeling schemes do not provide as much information about keys as the other schemes do. In MKIH, for example, a single key is required to derive all the keys on any particular chain in the canonical decomposition, in contrast to the labeling schemes.

It is also clear that the RSA-based schemes compare favorably with the respective iterative hashing schemes: the RSA-based schemes support direct derivation, require no more keys and require no more public information. However, it is worth noting that the primes that form the public storage in SKRSA will be relatively large compared to each item of public information in IKE or SKIH. (We might expect to be using 1024-bit primes in RSA-based schemes and only 128–256 bits for each item of encrypted data in an IKE scheme. So if n is small, as is likely to be the case in practice, there is likely to be little difference in actual public storage requirements.) We should also

note that modular exponentiation is a far more time-consuming operation than decryption, so this should be taken into consideration when making the comparison between SKRSA and IKE or SKIH.

5 New Key Assignment Schemes

Our original motivation was to construct better schemes for layered access control. However, it became apparent that the schemes described in the preceding section could be generalized to create key assignment schemes that could be applied to any poset. Moreover, the resulting schemes do not fit into the taxonomy of generic key assignment schemes proposed by Crampton *et al.* [7]. In this section, we describe briefly how two of our schemes for layered access control can be extended to create generic key assignment schemes.

Given a poset X , we first select a partition of X into chains $\{C_1, \dots, C_w\}$, where w is the width of X .⁷ We denote the length of C_i by ℓ_i , $1 \leq i \leq w$. We regard the maximum element of C_i as the first element in C_i and the minimum element as the last (or ℓ_i th) element.

Since $\{C_1, \dots, C_w\}$ is a partition of X , each $x \in X$ belongs to precisely one chain. So if x is the j th element of C_i , then we can represent x uniquely as the w -tuple $(0, \dots, j, 0, \dots, 0)$.

Let $C = x_1 \succ x_2 \succ \dots \succ x_m$ be any chain in X . Then we say any chain of the form $x_j \succ \dots \succ x_m$, $1 < j \leq m$ is a *suffix* of C . Now, for any $x \in X$, the set $\downarrow x \stackrel{\text{def}}{=} \{y \in X : y \leq x\}$ has non-empty intersection with one or more chains C_1, \dots, C_w . We now prove that the intersection of $\downarrow x$ and a chain C_i is a suffix of C_i . This result enables us to define the keys that should be given to a user with label x .

Proposition 7 *For all $x \in X$ and any chain $C \subseteq X$, either $\downarrow x \cap C$ is a suffix of C or $\downarrow x \cap C = \emptyset$.*

Proof Let $C = x_1 \succ \dots \succ x_m$ and suppose that $\downarrow x \cap C \neq \emptyset$. Then it must be the case that $x_j \in \downarrow x \cap C$ for some j and hence $x_j \in \downarrow x$. Moreover, since $\downarrow x = \{y \in X : y \leq x\}$ and \leq is transitive, we have that $x \geq x_j > x_{j+k}$, $1 \leq k \leq m - j$. Hence, $x_{j+k} \in \downarrow x$, $1 \leq k \leq m - j$. Hence $\downarrow x \cap C = \{x_j, x_{j+1}, \dots, x_m\}$ (and $\downarrow X \cap C$ is a suffix of C , as required). ■

⁷Unlike $R_{m,n}$, there is no canonical partition for an arbitrary poset X . At this stage, we do not consider what features a “good” partition might have. We return to this question towards the end of the section.

The above proposition indicates how we should allocate keys to users. Since $\{C_1, \dots, C_w\}$ is a partition of X into chains, $\{\downarrow x \cap C_1, \dots, \downarrow x \cap C_w\}$ is a disjoint collection of chain suffixes. Moreover, the keys for each element in X have been chosen so that the key for the j th element of a chain can be used to compute all lower elements in that chain. Hence, we can see that a user with label x must be given the keys for the maximal elements in the non-empty suffixes $\downarrow x \cap C_1, \dots, \downarrow x \cap C_w$. Given $x \in X$, let $\hat{x}_1, \dots, \hat{x}_w$ denote these maximal elements, with the convention that $\hat{x}_i = \perp$ if $\downarrow x \cap C_i = \emptyset$. Clearly the number of \hat{x}_i such that $\hat{x}_i \neq \perp$ is no greater than w . The above result and observations provide the foundations of both the schemes that follow.

5.1 Multiple-Key Iterated Hash Scheme

We first consider the use of iterated hashing. The scheme administrator

- selects a chain partition of X into w chains C_1, \dots, C_w ;
- selects w secret values $\sigma_1, \dots, \sigma_w$ and w hash functions h_1, \dots, h_w ;
- defines the key for the maximum element of chain C_i to be σ_i ;
- for each pair $x, y \in C_i$ such that $x < y$, defines $\kappa(x) = h_i(\kappa(y))$;
- for each $x \in X$, defines the private information for x to be $\{\kappa(\hat{x}_i) : \hat{x}_i \neq \perp\}$.

We denote the key for the j th element of C_i by $\kappa_{i,j}$. Clearly (as in Sec. 4), a user in possession of $\kappa_{i,j}$ can compute $\kappa_{i,y}$, for any $y > j$, by $y - j$ iterative hash computations.

Figure 2 illustrates a poset X of width 4 and one possible partition of X into 4 chains. In Fig. 2(b) we label each node with a 4-tuple indicating the relative position of the node in one of the chains.⁸ So, for example, x_8 has the label $(0, 2, 0, 0)$ because it is the second element in the second chain. Hence, if the second chain $(x_{11} \succ x_8 \succ x_4 \succ x_1)$ is associated with the secret value σ_2 , then $\kappa(x_{11}) = \sigma_2$ and $\kappa(x_8) = h(\sigma_2)$.

Clearly, the number of steps required for key derivation is bounded by the length of the longest chain in the partition. With this in mind, it might be sensible to choose a chain partition in which the chains are as similar in length as possible. In terms of correctness and collusion security, the MKIH scheme for arbitrary posets is no different from $R_{m,n}$.

⁸These labels have no connection with the CDM labeling defined in Sec. 3.




Figure 2: Partitioning an arbitrary poset into chains

5.2 Multi-Key RSA Scheme

In the second scheme, we use RSA key chains. The scheme administrator generates and publishes N_1, \dots, N_w , as before. As in the preceding section, the scheme administrator selects a chain partition of X and defines the key for the maximum element of the i th chain to be σ_i . Now, for each pair $x, y \in C_i$ such that $x < y$, the scheme administrator defines $\kappa(x) = (\kappa(y))^2 \bmod N_i$. Finally, the private information associated with $x \in X$ is defined to be $\{\kappa(\hat{x}_i) : \hat{x}_i \neq \perp\}$ (as in the preceding scheme).

5.3 Related Work

In Table 3, we summarize the differences between our schemes and existing generic key assignment schemes. We also illustrate these trade-offs for the poset and chain partition given in Fig. 2.

A *direct key encrypting* (DKE) key assignment scheme guarantees single-step key derivation by including a substantial amount of public information. In a DKE scheme, the encryption of $\kappa(y)$ with $\kappa(x)$ forms part of the public information whenever $y < x$. In contrast, an *iterative key encrypting* (IKE) key assignment scheme reduces the amount of public information but requires iterative key derivation; in the worst case, d steps may be required (where d is the length of the longest path in the Hasse diagram of X). In an IKE scheme, $\text{Enc}_{\kappa(x)}(\kappa(y))$ forms part of the public information whenever $y < x$. We write c for the cardinality of the cover relation \triangleleft and r for the cardinality of the order relation \leq . Every DKE or IKE key assignment scheme has the property that each user has a single key. In contrast, the *trivial* key assignment scheme,

in which the user is given every key that she requires clearly needs no public information to enable key derivation.

Scheme	Private storage				Public storage	Key derivation			
	x	x_{12}	x_{10}	x_9		x	x_{12}	x_{10}	x_9
Trivial	$\downarrow x$	12	7	5	0	1			
DKE	1				51	1			
IKE	1				14	d	5	4	3
MKIH	$\leq w$	4	3	3	0	$\leq d$	3	2	2
MKRSA	$\leq w$	4	3	3	0	1			

Table 3: A comparison of our schemes with existing generic key assignment schemes

Our schemes provide a trade-off between these three alternative schemes: users may have multiple keys ⁹⁾ but little or no public information is required and key derivation will generally be quicker than for an equivalent IKE scheme.¹⁰

6 Conclusion

We have shown how to construct a new type of generic key assignment scheme using chain partitions, the inspiration for the original constructions being provided by the problem of enforcing layered access control in scalable multimedia formats. Our schemes, both for layered access control and as generic key assignment schemes, compare favorably with those in the literature.

We have many ideas for future work. Of primary interest is whether we can prove that our schemes are secure against key recovery [2], a more exacting criterion than that of collusion security used in this paper. We also hope to gain some insight, either from a mathematical analysis or through experimental work, into what might be the best choice(s) of chain partition for an arbitrary poset. A third area for potential research is to generalize

⁹Schemes with multiple keys were usually disregarded in the early literature [7], although several recent schemes have made use of multiple keys [2, 4].

¹⁰Note that key derivation in our multi-key iterative hash scheme cannot be worse than key derivation in IKE and, in many cases, will be considerably better. As we observed earlier, it would be sensible to choose a chain partition in which all chains have approximately the same length. If this is possible, key derivation is approximately $|X|/w$. The poset in Fig. 2, for example, can be partitioned into 4 chains of length 3. Then any key can be derived in no more than 2 hops, whereas an IKE scheme would require 5 hops to derive $\kappa(x_1)$ from $\kappa(x_{12})$.

our constructions to more than two scalable components (most likely using a recursive construction with one of our schemes from Sections 3 and 4 as a base case).

References

- [1] S.G. Akl and P.D. Taylor. Cryptographic solution to a problem of access control in a hierarchy. *ACM Transactions on Computer Systems*, 1(3):239–248, 1983.
- [2] M.J. Atallah, M. Blanton, N. Fazio, and K.B. Frikken. Dynamic and efficient key management for access hierarchies. *ACM Transactions on Information and System Security*, 12(3):1–43, 2009.
- [3] M.J. Atallah, M. Blanton, and K.B. Frikken. Key management for non-tree access hierarchies. In *Proceedings of 11th ACM Symposium on Access Control Models and Technologies*, pages 11–18, 2006.
- [4] M.J. Atallah, M. Blanton, and K.B. Frikken. Efficient techniques for realizing geo-spatial access control. In *Proceedings of the 2007 ACM Symposium on Information, Computer and Communications Security*, pages 82–92, 2007.
- [5] D.E. Bell and L. LaPadula. Secure computer systems: Unified exposition and Multics interpretation. Technical Report MTR-2997, Mitre Corporation, Bedford, Massachusetts, 1976.
- [6] C. Christopoulos, A. Skodras, and T. Ebrahimi. The JPEG2000 still image coding system: An overview. *IEEE Transactions on Consumer Electronics*, 46(4):1103–1127, 2000.
- [7] J. Crampton, K. Martin, and P. Wild. On key assignment for hierarchical access control. In *Proceedings of 19th Computer Security Foundations Workshop*, pages 98–111, 2006.
- [8] B.A. Davey and H.A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, Cambridge, United Kingdom, 1990.
- [9] D.E. Denning. A lattice model of secure information flow. *Communications of the ACM*, 19(5):236–243, 1976.
- [10] R.P. Dilworth. A decomposition theorem for partially ordered sets. *Annals of Mathematics*, 51:161–166, 1950.

- [11] N. Hashimoto, S. Imaizumi, M. Fujiyoshi, and H. Kiya. Hierarchical encryption using short encryption keys for scalable access control of JPEG 2000 coded images. In *Proceedings of the 2008 IEEE International Conference on Image Processing*, pages 3116–3119, 2008.
- [12] S. Imaizumi, M. Fujiyoshi, Y. Abe, and H. Kiya. Collusion attack-resilient hierarchical encryption of JPEG 2000 codestreams with scalable access control. In *Proceedings of the 2007 IEEE International Conference on Image Processing*, volume 2, pages 137–140, 2007.
- [13] S. Imaizumi, O. Watanabe, M. Fujiyoshi, and H. Kiya. Generalized hierarchical encryption of JPEG 2000 codestreams for access control. In *Proceedings of the 2005 IEEE International Conference on Image Processing*, volume 2, pages 1094–1097, 2005.
- [14] J. Katz and Y. Lindell. *Introduction to Modern Cryptography*. Chapman & Hall/CRC, 2007.
- [15] W. Li. Overview of fine granularity scalability in MPEG-4 video standard. *IEEE Transactions on Circuits and Systems for Video Technology*, 11(3):301–317, 2001.
- [16] B.B. Zhu, S.M. Feng, and S. Li. An efficient key scheme for layered access control of MPEG-4 FGS video. In *Proceedings of the 2004 IEEE International Conference on Multimedia and Expo*, volume 1, pages 443–446, 2004.