



Contents lists available at ScienceDirect

# The Journal of Logic and Algebraic Programming

journal homepage: [www.elsevier.com/locate/jlap](http://www.elsevier.com/locate/jlap)



## 1 Cryptographically-enforced hierarchical access control with multiple 2 keys<sup>☆</sup>

3 Jason Crampton

4 *Information Security Group, Royal Holloway, University of London, UK*

### ARTICLE INFO

#### Article history:

Received 1 April 2009

Accepted 1 April 2009

Available online xxx

#### Keywords:

Information flow policy

Cryptographic access control

Key assignment scheme

Multi-key assignment scheme

### ABSTRACT

Hierarchical access control policies, in which users and objects are associated with nodes in a hierarchy, can be enforced using cryptographic mechanisms. Protected data is encrypted and authorized users are given the appropriate keys. Lazy re-encryption techniques and temporal hierarchical access control policies require that multiple keys may be associated with a node in the hierarchy. In this paper, we introduce the notion of a multi-key assignment scheme to address this requirement. We define bounded, unbounded, synchronous, and asynchronous schemes. We demonstrate that bounded, synchronous schemes provide an alternative to temporal key assignment schemes in the literature, and that unbounded asynchronous schemes provide the desired support for lazy re-encryption.

© 2009 Elsevier Inc. All rights reserved.

5

## 7 1. Introduction

8 One of the fundamental security services in computer systems is *access control*, a mechanism for constraining the inter-  
9 action between (authenticated) users and protected resources. Generally, access control is implemented by an authorization  
10 service, which includes an authorization decision function for deciding whether a user request to access a resource should  
11 be permitted or not. The output of an authorization decision function is usually determined by evaluating the access request  
12 with respect to an authorization policy.

13 It is sometimes appropriate, however, to enforce an authorization policy using cryptographic techniques. In such circum-  
14 stances, protected data (objects) are encrypted and authorized users are given the appropriate cryptographic keys. Such  
15 an approach is useful when objects are: read often, by many users; written once, or rarely, by the owner of the data; and  
16 transmitted over unprotected networks. Fu et al. [10] cite content distribution networks, such as Akami and BitTorrent, as  
17 examples where some kind of cryptographic access control is particularly suitable.

18 A no-read-up information flow policy [9] controls access to resources by comparing security labels assigned to users and  
19 objects; a user may have read access to an object (causing an information flow from object to user) if and only if the user's  
20 security label is at least as high as that of the object. Such a policy can be implemented by associating a symmetric encryption  
21 key with each security label and encrypting objects with the appropriate key. A user is given (or can derive) the encryption  
22 key for each label less than or equal to his security label. Since the seminal paper of Akl and Taylor [1], there has been a  
23 considerable amount of research into so-called *key assignment schemes*, which seek to minimize the number of keys required  
24 by each user (see the recent survey paper [7] for an overview of such schemes).

<sup>☆</sup> A shorter version of this paper appeared in the Proceedings of the 12th Nordic Workshop on Secure IT Systems (NordSec 2007). Programme committee co-chair Úlfar Erlingsson acted as guest editor for this revised version.

E-mail address: [jason.crampton@rhul.ac.uk](mailto:jason.crampton@rhul.ac.uk)

25 *Motivation and related work.* One of the main practical problems with cryptographic access control in general, and with key  
 26 assignment schemes in particular, is that huge amounts of data may need to be re-encrypted when a key is changed. There  
 27 are a number of reasons why keys might change, of which the most are important are user revocation, key compromise and  
 28 access control policy requirements.

29 One way of alleviating the burden of object re-encryption following a user revocation is to re-encrypt objects with the  
 30 new key only when the content changes. This type of approach, sometimes called *lazy re-encryption*, has been used in  
 31 cryptographic file systems [4,13]. A consequence of this is that a user may require two or more keys for objects in the same  
 32 protection domain: one or more keys for objects that have not yet been re-encrypted and one for those that have. The burden  
 33 of object re-encryption is particularly onerous in key assignment schemes: if a user with security label  $l$  is revoked, then the  
 34 key for every security label  $l' \leq l$  has to be changed and every object with security label  $l'$  has to be re-encrypted. Therefore,  
 35 lazy re-encryption is a particularly appropriate technique for key assignment schemes and one that has not previously been  
 36 studied.

37 Some key assignment schemes incorporate a temporal element (see [3] for a survey of such schemes): Bertino et al., for  
 38 example, have used cryptographic access control to restrict access to subscription-based content, with each time period using  
 39 a different key to encrypt content [6]. In this case, a user may require two or more keys for objects in the same protection  
 40 domain: one key for objects encrypted in the current period and one for each of the preceding periods. All such research has  
 41 assumed that all keys (including those for future time intervals) are issued when the scheme is instantiated. Although this  
 42 simplifies key distribution, there is, however, a very good reason for not doing so: if a user pays subscriptions periodically  
 43 and defaults on the payments, it will be necessary to update all the keys in the scheme in order to revoke the defaulter's  
 44 authorization.

45 *Contributions.* In this paper, we examine extensions to existing key assignment schemes that permit several keys to be  
 46 associated with each security label and illustrate how such schemes can be used to support lazy re-encryption and temporal  
 47 access control policies. The main contributions of the paper are: to extend lazy revocation schemes to hierarchical access  
 48 control policies (in contrast to the Unix-like “owner-group-world” policies of previous work), and to permit the timed release  
 49 of new encryption keys for temporal hierarchical access control policies. The schemes presented in this paper are derived  
 50 from techniques used for existing key assignment schemes, in which each label is associated with a single encryption key.  
 51 They also make use of some of the ideas of Fu et al. [10] and work on one-time passwords [11].

52 *Outline.* In the next section, we summarize the characteristics of information flow policies and existing key assignment  
 53 schemes. We then describe our new schemes, how they support multiple keys for each security label, and discuss the  
 54 potential applications of each type of scheme. We conclude the paper with a summary of our contributions and some ideas  
 55 for future work.

56 **2. Preliminaries**

57 A *partially ordered set* (or *poset*) is a pair  $(L, \leq)$ , where  $\leq$  is a reflexive, anti-symmetric, transitive binary relation on  $L$ .  
 58 We say  $x$  covers  $y$ , denoted  $y \triangleleft x$ , if  $y < x$  and there does not exist  $z \in L$  such that  $y < z < x$ .  $L$  is a *total order* if for all  $x, y \in L$ ,  
 59 either  $x \leq y$  or  $y \leq x$ .

60 The *Hasse diagram* of a poset is the directed graph  $(L, \triangleleft)$  [8]. A simple Hasse diagram is shown in Fig. 1a (on page 2).  
 61 Note that  $d \leq a$ , but this “transitive edge” is not included in the graph. Note also that all edges are assumed to be directed  
 62 upwards.

63 An *order ideal*  $I \subseteq L$  has the property that if  $x \in I$  and  $y \leq x$  then  $y \in I$ . Similarly, an *order filter*  $F \subseteq L$  has the property that  
 64 if  $x \in F$  and  $y \geq x$  then  $y \in F$ . The order ideal generated by  $x \in L$ , denoted  $\downarrow x$ , is defined to be  $\{y \in L : y \leq x\}$ . It is easy to show  
 65 that the set complement of an order ideal in  $L$  is an order filter (and vice versa) [8].

66 Henceforth we adopt the following conventions: given a poset  $L$  and  $x, y \in L$ , we may write  $y < x$  if  $y \leq x$  and  $y \neq x$ , and  
 67 we may write  $x \geq y$  if  $y \leq x$ ;  $x$ , when used as input to some (cryptographic) function, will denote a string identifying  $x \in L$ ;  
 68  $E$  denotes a symmetric encryption algorithm, and  $E_k(p)$  denotes the encryption of plaintext  $p$  with key  $k$ ;  $s \parallel t$  denotes the




Fig. 1. An example of the Akl-Taylor scheme.

69 concatenation of strings  $s$  and  $t$ ;  $m \mid n$  means that integer  $m$  divides integer  $n$  without remainder;  $(m, n)$  denotes the greatest  
70 common divisor of integers  $m$  and  $n$ ; in particular,  $(m, n) = 1$  means that  $m$  and  $n$  are co-prime.

71 We assume the existence of an RSA key generator, a randomized algorithm that takes a security parameter  $k$  as input and  
72 outputs a triple  $(n, e, d)$  such that:

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

76 We also assume the existence of an associated hash function, which maps elements of  $\mathbb{Z}_n^*$  (in some standard representation)  
77 to bit-strings of some desired length.

## 78 2.1. Information flow policies

79 **Definition 1.** An *information flow policy* is a tuple  $(L, \leq, U, O, \lambda)$ , where:  $(L, \leq)$  is a (finite) partially ordered set of security  
80 labels;  $U$  is a set of users;  $O$  is a set of (protected) objects;  $\lambda : U \cup O \rightarrow L$  is a security function that associates users and objects  
81 with security labels.

82 Such policies were of particular interest in the mid-1970s through to the mid-1980s, and formed part of the Bell-LaPadula  
83 security model [5]. The basic semantics of the policy is that a user  $u$  can read an object  $o$  if  $\lambda(u) \geq \lambda(o)$ . Henceforth, we will  
84 represent an information flow policy  $(L, \leq, U, O, \lambda)$  as a pair  $(L, \leq)$  with the tacit understanding that  $U, O$  and  $\lambda$  are given.

85 Clearly, one way of implementing such a policy is to encrypt an object with security label  $y \in L$  with a key  $\kappa(y)$ , and to  
86 provide all users with security label  $x \geq y$  with the key  $\kappa(y)$ . Henceforth we assume that the policy will be implemented by  
87 encrypting objects and distributing keys to users, enabling them to decrypt the objects to which they should have access.

88 A consequence of using cryptographic methods to control access is that users may cache decrypted data objects locally.  
89 Hence, revoking a user's assignment to a security level does not necessarily have an immediate effect on that user's ability  
90 to read documents he has previously accessed. For this reason, lazy re-encryption – only re-encrypting data objects when  
91 they are modified – is a reasonable strategy to employ in this context.

## 92 2.2. Key assignment schemes

93 A *key assignment scheme* (KAS) is a means of implementing an information flow policy using cryptographic techniques.  
94 The simplest such scheme provides the user  $u$  the set of keys  $\{\kappa(y) : y \leq \lambda(u)\}$ . More sophisticated KASs seek to minimize  
95 the number of keys that need to be distributed to users. This entails either making certain additional information publicly  
96 available or providing each user with additional secret information (or both).

97 A recent paper formalized the characteristic features of a key assignment scheme [7]. In simple terms, there must be  
98 a way of choosing keys for each  $x \in L$ , a method for constructing public information, and a method for deriving keys. The  
99 paper also identified five generic constructions for key assignment schemes and compared these schemes according to the  
100 following criteria: amount of public storage required, amount of private storage required, complexity of key derivation, and  
101 complexity of key updates. We discuss two of these schemes in the next two sections. In both schemes a user  $u$  has a single  
102 key  $\kappa(\lambda(u))$ .

### 103 2.2.1. IKE KAS

104 An *iterative key encrypting* (IKE) KAS uses public information to enable a user to iteratively derive keys for which he is  
105 authorized. The trusted center:

- 106 • chooses  $\kappa(x)$ ,  $x \in L$ , at random from the key space;
- 107 • publishes  $\{E_{\kappa(x)}(\kappa(y)) : y \leq x, y \in L\}$ .

108 Informally, the public information “encrypts” edges in the Hasse diagram: for any edge  $(x, y)$ ,  $\kappa(y)$  will be encrypted with  
109  $\kappa(x)$  and stored in the public information. Key derivation is performed by traversing a path in the Hasse diagram decrypting  
110 keys along the way. More formally, if  $y < x$ , then there exists a path  $y = z_0 < z_1 < \dots < z_n = x$ ,  $n \geq 1$ , so a user with security  
111 label  $x$  can successively derive keys  $\kappa(z_{n-1}), \dots, \kappa(z_0) = \kappa(y)$ . Assuming a suitable encryption method is chosen, no user  
112 can feasibly compute keys for which he is not authorized.

113 The AFB scheme, due to Atallah, Frikken and Blanton is an example of an IKE KAS [2]. The key space is  $\{0, 1\}^l$ , for some  
114 integer  $l$ , and the public information is defined to be  $\{\kappa(y) \oplus h(\kappa(x) \parallel y) : y < x\}$ , where  $h : \{0, 1\}^* \rightarrow \{0, 1\}^l$  is a hash function.

### 115 2.2.2. The Akl–Taylor scheme

116 The second generic scheme is a *node-based key assignment scheme*. Key derivation in a node-based KAS is direct and the  
117 user is only required to store a single key. The disadvantage of a node-based KAS is that key updates may be more disruptive  
118 than in an IKE KAS, because keys are defined in terms of public information. The earliest KAS (of any type) is due to Akl and  
119 Taylor; it remains the best known, simplest, and most effective node-based KAS [1]. The trusted center

- 120 • runs the RSA key generator to obtain  $(n, e, d)$ , of which only  $n$  is used;
- 121 • chooses a secret  $s$  at random from  $\mathbb{Z}_n^*$ ;
- 122 • chooses a mapping  $\nu : L \rightarrow \mathbb{N}^*$  such that  $\nu(x) \mid \nu(y)$  if and only if  $y \leq x$ ;
- 123 • defines  $\kappa(x) = s^{\nu(x)} \pmod n$ ;
- 124 • publishes  $\{n\} \cup \{\nu(x) : x \in L\}$ .

A user with security label  $x$  can derive  $\kappa(y)$ ,  $y \leq x$ , using private information  $\kappa(x)$  and public information  $\nu(x)$  and  $\nu(y)$ , by computing

$$(\kappa(x))^{\frac{\nu(y)}{\nu(x)}} = (s^{\nu(x)})^{\frac{\nu(y)}{\nu(x)}} = s^{\nu(y)} = \kappa(y).$$

Akl and Taylor proved that no user (or set of users) can derive a key for which he is not authorized, provided  $\nu$  is chosen appropriately. We discuss the security of key assignment schemes in Section 2.2.4. The canonical way of constructing  $\nu$  is to associate a distinct prime  $\alpha(x)$  with each  $x \in L$  and to define

$$\nu(x) = \prod_{y \not\leq x} \alpha(y). \tag{1}$$

- 125 An example of this construction is shown in Fig. 1.  $L$  is shown in Fig. 1a; the small italic numbers in Fig. 1b denote the primes  
 126 associated with each element of  $L$ ; and the larger numbers denote the value of  $\nu(x)$  for each  $x \in L$ . Note that  $\nu(x) \mid \nu(y)$  if  
 127 and only if  $x \geq y$ .

MacKinnon et al. identified a different initial labeling  $\alpha : L \rightarrow \mathbb{N}$ , based on a decomposition of  $L$  into disjoint chains [14]. In other words, every element in  $L$  is a member of precisely one chain. We associate a distinct prime  $\alpha_i$  with the  $i$ th chain. Then for  $x \in L$ , the  $j$ th element in the  $i$ th chain (where the maximal element in each chain is the first element), we define

$$\alpha(x) = \alpha_i^j \quad \text{and} \quad \nu(x) = \text{lcm}\{\alpha(y) : y \not\leq x\}.$$

- 128 An example of this labeling, which we call the MTMA labeling (after its creators, MacKinnon, Taylor, Meijer and Akl), is shown  
 129 in Fig. 1c. The solid lines denote a chain ( $a > b > d$ ) associated with the prime 2; the remaining node ( $c$ ) forms a trivial chain  
 130 associated with the prime 3. Note that the  $\nu$  values satisfy condition (2).

### 131 2.2.3. Implementation issues

It has long been claimed that key derivation in the Akl–Taylor scheme may be expensive in practice (see [2,12], for example) because it appears that a user must compute

$$\kappa(y) = \kappa(x)^{\frac{\nu(y)}{\nu(x)}}$$

- 132 to obtain  $\kappa(y)$  from  $\kappa(x)$ , where  $\nu(y)$  and  $\nu(x)$  are the products of primes. This means that an expensive division operation  
 133 is required followed by an exponentiation. However, it was noted by Crampton et al. [7] that the  $\nu$  values were simply used  
 134 to encode the structure of the poset, and that this encoding could be done in such a way as to avoid expensive division  
 135 operations during key derivation.

Consider the definition of  $\nu$  in the canonical labeling:

$$\nu(x) = \prod_{y \not\leq x} \alpha(y).$$

- 136 The product is taken over all elements that are not less than or equal to  $x$ . The set of elements less than or equal to  $x$  is (by  
 137 definition) the order ideal generated by  $x$ . Hence, the product used to compute  $\nu(x)$  is taken over the complement of the  
 138 order ideal  $\downarrow x$ , which is an order filter.

We now choose some enumeration of the elements of  $L = \{x_1, \dots, x_{|L|}\}$  and for each  $x \in L$ , define the *characteristic tuple* (of  $x$ )  $\chi_x = (b_1, \dots, b_{|L|})$ , where  $b_i = 0$  if  $x_i \leq x$  and  $b_i = 1$  otherwise. Hence, we may represent the elements of the filter  $L \setminus \downarrow x$  by  $\chi_x$ . Moreover, we may publish this information and the mapping  $\alpha$  associating each  $x_i \in L$  with a prime  $\alpha(x_i)$ . Now,

$$\nu(x) = \prod_{i=1}^{|L|} (\alpha(x_i))^{b_i},$$

where  $b_i$  is the  $i$ th element of  $\chi_x$ . Hence,

$$\frac{\nu(y)}{\nu(x)} = \frac{\prod_{i=1}^{|L|} (\alpha(x_i))^{b'_i}}{\prod_{i=1}^{|L|} (\alpha(x_i))^{b_i}},$$

where  $b'_i$  is the  $i$ th element of  $\chi_y$ . Hence

$$\frac{\nu(y)}{\nu(x)} = \prod_{i=1}^{|L|} (\alpha(x_i))^{b'_i},$$

139 where  $b'_i = b'_i - b_i$ . In other words, no division is required. We simply take the co-ordinatewise difference of the characteristic  
140 tuples of  $x$  and  $y$  and then multiply the appropriate primes together.

141 Taking our running example in Fig. 1 and enumerating the elements of  $L$  in alphabetical order, the order ideal generated  
142 by  $b$  is  $\{b, d\}$  and the order ideal generated by  $d$  is  $\{d\}$ . The characteristic tuples of  $b$  and  $d$  are  $(1, 0, 1, 0)$  and  $(1, 1, 1, 0)$ ,  
143 respectively. To compute  $\kappa(d)$  from  $\kappa(b)$  we compute  $(1, 1, 1, 0) - (1, 0, 1, 0)$  to obtain  $(0, 1, 0, 0)$ , which means that we  
144 have to raise  $\kappa(b)$  to the power of the prime associated with  $b$  to obtain  $\kappa(d)$ . This can be confirmed by looking at  
145 Fig. 1b. In short, no division of exponents is required. That is, there exist efficient implementations of the Akl–Taylor  
146 scheme.

We now observe that a similar type of representation can be used for the MTMA labeling. Recall that  $L$  is partitioned  
into  $i$  chains, the  $i$ th chain being associated with a distinct prime  $\alpha_i$ ,  $\alpha(x) = \alpha_i^j$  and  $\nu(x) = \text{lcm}\{\alpha(y) : y \not\leq x\}$ . In other words,  
 $\nu(x) = \prod_{i=1}^w \alpha_i^{a_i}$ , where  $a_i \in \mathbb{N}$ . In this context, the characteristic tuple of  $x$  is defined to be  $(a_1, \dots, a_w) \in \mathbb{N}^w$ . In our example,  
we have two chains and we have

$$\chi_a = (0, 0), \quad \chi_b = (1, 1), \quad \chi_c = (2, 0), \quad \chi_d = (2, 1).$$

147 As before, division of exponents corresponds to subtraction of tuples. To derive  $\kappa(d)$  from  $\kappa(b)$ , for example, we compute  
148  $(2, 1) - (1, 1) = (1, 0)$ , which implies that we raise  $\kappa(b)$  to the power of the prime associated with the first chain.

#### 149 2.2.4. Security considerations

150 The fact that  $\nu(x) \mid \nu(y)$  if, and only if,  $x \geq y$  implies a user  $u$  with security label  $x$  cannot derive a key  $\kappa(y)$ ,  $y \not\leq x$ , unless  
151  $u$  can compute roots modulo  $n$ . This is known as the *RSA problem* and is believed to be as hard as factoring  $n$ : it is this *RSA*  
152 *assumption* on which the security of the RSA cryptosystem is based [15]. Akl and Taylor proved that no *set* of users can collude  
153 to derive a key for which none of them is authorized, provided  $\nu$  is chosen appropriately. In particular, we have the following  
154 result.

155 **Theorem 2** (Akl and Taylor).  $\kappa(y)$  can be feasibly computed from a set of keys  $\{\kappa(x_1), \dots, \kappa(x_m)\}$  if, and only if,  $\text{gcd}\{\nu(x_1), \dots,$   
156  $\nu(x_m)\} \mid \nu(y)$ .

The result is proved by showing that if this property did not hold then it would be feasible to solve the RSA problem.  
Hence, we choose  $\nu$  in such a way that for all  $x \in L$ ,

$$\text{gcd}\{\nu(y) : y \not\leq x\} \nmid \nu(x). \tag{2}$$

157 The canonical labeling defined in Eq. (1),  $\nu(x) = \prod_{y \not\leq x} \alpha(y)$ , satisfies property (2), as does the MTMA labeling.

158 Atallah et al. examined the security of IKE KASs [2], introducing the notions of *key recovery* and *key indistinguishability*,  
159 which are analogous to plaintext recovery and plaintext indistinguishability in conventional security analyses. Informally, a  
160 scheme is secure against key recovery if it is infeasible for a group of users to pool key information to obtain a key for which  
161 none of them is authorized. In this sense, the Akl–Taylor scheme is secure against key recovery.

162 Key indistinguishability is believed to be an important consideration in IKE KASs [10]. Informally, if an attacker  $a$  is given a  
163 key  $k$  that purports to be  $\kappa(x)$  for some  $x \in L$  and  $a$  is able to decide (under reasonable computational assumptions) whether  
164  $k$  is actually  $\kappa(x)$  with some probability of success significantly greater than  $1/2$ , then the scheme does not have the key  
165 indistinguishability property. Clearly the basic IKE KAS does not have this property, since an attacker that knows  $\kappa(y)$ , for  
166 some  $y \leq x$ , can determine whether  $k$  is indeed  $\kappa(x)$  by trying to derive  $\kappa(y)$  from  $k$  and the public information. If he succeeds  
167 in deriving  $\kappa(y)$ , then he knows that  $k = \kappa(x)$ .

168 The IKE KASs described in this paper do not have the key indistinguishability property, although Atallah et al. [2] and,  
169 independently, Fu et al. [10] have developed straightforward modifications to existing schemes to obtain this property.<sup>1</sup> We  
170 do not make similar modifications to our schemes for ease of exposition.

171 Ateniese et al. [3] extended the security analysis of Atallah et al. to temporal key assignment schemes, and also considered  
172 active and passive adversaries. They showed that an active adversary has no advantage over a passive one for either key  
173 recovery or key indistinguishability. In other words, the security properties of an IKE KAS appear to be dependent only on

<sup>1</sup> Informally, users with security label  $x$  are provided with a secret  $\sigma(x)$ , rather than an encryption key  $\kappa(x)$ . The secret  $\sigma(x)$  may be used to derive  $\kappa(x)$  (using some suitable publicly known one-way function) and  $\sigma(y)$  for all  $y \leq x$  (using existing techniques for key assignment schemes). The point being that an attacker cannot deduce  $\sigma(x)$  from  $\kappa(x)$  and hence is unable to use key derivation (as described above) to distinguish the genuine  $\kappa(x)$  from a random value.

174 the security properties of the algorithm used to encrypt keys, the results of which are made public. We hope to confirm these  
175 types of properties for our schemes in our future work.

### 176 3. Multi-key assignment schemes

177 We now introduce the concept of a *multi-key assignment scheme* (MKAS), which is a key assignment scheme that supports  
178 multiple keys for each security label. We assume there are a number of update events indexed by the natural numbers, and  
179 that the  $i$ th key is used to

- 180 • encrypt objects created between the  $i$ th and  $(i + 1)$ th updates, and
- 181 • re-encrypt existing objects whose contents change between the  $i$ th and  $(i + 1)$ th updates.

182 The  $i$ th interval is the period between the  $i$ th and  $(i + 1)$ th update.

183 If a user  $u$  is currently authorized to access an object  $o$ , then  $u$  is also authorized to access any object  $o'$  such that  
184  $\lambda(o) \geq \lambda(o')$ . This implies that  $u$  must be able to compute each key that has been used for  $\lambda(o')$ . We also require that a  
185 user  $v$  revoked at the  $j$ th key update cannot access any object created or amended since the  $j$ th update. We write  $\kappa_i(x)$  to  
186 denote the  $i$ th key for security label  $x$ .<sup>2</sup> In other words, an authorized user  $u$  in time interval  $i$  can derive any key in the set  
187  $\{\kappa_j(l) : l \leq \lambda(u), 0 \leq j \leq i\}$ .

188 There are two main considerations when designing a multi-key assignment scheme.

- 189 • *Is the number of updates determined in advance?* We say an MKAS is *bounded* if the number of updates is pre-determined,  
190 and *unbounded* otherwise.

191 Bounded MKASs are well suited for temporal access control policies, because the number of time periods will usually  
192 be fixed and known in advance. They are less appropriate for schemes intended to support lazy re-encryption, since  
193 the number of updates will not be predictable. Unbounded MKASs are far more suitable for lazy re-encryption.

- 194 • *Can  $\kappa(x)$  be updated independently of  $\kappa(y)$ ?* We say an MKAS is *synchronous* if all keys are updated at the same time,  
195 and *asynchronous* otherwise.

196 Synchronous MKASs are particularly suitable for temporal access control policies, because each time interval requires  
197 a fresh set of keys, one for each security label. Asynchronous MKASs may well be more suitable for implementing lazy  
198 re-encryption, because it may only be necessary to update a small number of keys following the revocation of a user.

199 The key observation to make prior to introducing the schemes we have devised is to note that the requirements can  
200 be regarded as having two dimensions: a “policy dimension”, determined by  $L$ ; and a “temporal dimension”, which can be  
201 modeled as a chain of (update) events. In the next section we consider three different ways of realizing bounded schemes,  
202 each having different properties and advantages. In Section 3.2 we introduce two unbounded synchronous schemes. In  
203 Section 3.3 we describe two unbounded asynchronous schemes. We conclude with a summary and comparison of the  
204 respective advantages of our schemes.

#### 205 3.1. Bounded MKASs

206 We assume that there are a maximum of  $m$  updates, where  $m$  is chosen in advance by the scheme administrator, and that  
207 the keys for  $L' \subseteq L$  are updated. The basic technique is to construct an IKE KAS for  $(L, \leq)$  for each time interval. In addition, we  
208 provide a mechanism for computing  $\kappa_j(x)$  from  $\kappa_i(x)$  for any  $j < i$ . Note that the schemes described can be either synchronous  
209 or asynchronous; a synchronous scheme is obtained by setting  $L' = L$  at each update.

##### 210 3.1.1. Node-based iterative MKAS

211 In this scheme we choose the keys  $\kappa_0(x), \dots, \kappa_m(x)$  in such a way that  $\kappa_j(x)$  can be derived directly from  $\kappa_i(x)$ , for all  
212  $j < i$ . The construction is based on a very simple instance of the Akl-Taylor scheme in which  $L$  is a chain of  $m$  elements.

213 The trusted center first obtains a large compound integer  $n$  using an RSA key generator. For each  $x \in L$ , the trusted center

- 214 • chooses a secret  $s(x) \in \mathbb{Z}_n^*$ , such that for all  $y \in L$ ,  $(s(x), s(y)) \neq 1$  if and only if  $x = y$ ;
- 215 • defines  $\kappa_j(x) = (s(x))^{2^{m-j}} \pmod n$ ;
- 216 • defines an IKE for  $L$  and publishes  $\{E_{\kappa_0(x)}(\kappa_0(y)) : y \in L, y \triangleleft x\}$ .

217 We call the set of keys  $\kappa_0(x), \dots, \kappa_m(x)$  an *Akl-Taylor key chain*.

218 Let the set of keys changed at the  $i$ th update be  $L'_i$ . Following the  $i$ th update, for each  $x \in L'_i$ , the trusted center obtains the  
219 next key for  $x$  and updates the public information. In particular, the trusted center re-computes and replaces  $\{E_{\kappa_i(x)}(\kappa_i(y)) : y \in L', y \triangleleft x\}$   
220 and  $\{E_{\kappa_i(x)}(\kappa_i(y)) : y \in L', x \in L, y \triangleleft x\}$ . As a consequence, the amount of public information required by  
221 this scheme remains constant.

In order to recover  $\kappa_j(y)$ , given  $\kappa_i(x)$ , with  $y \leq x$  and  $j \leq i$ , the user first computes  $\kappa_i(y)$  using the public information. The user then computes

$$(\kappa_i(y))^{2^{i-j}} = (\kappa_i(y))^{2^{m-i}} = (s(y)^{2^{m-i}})^{2^{m-j}} = s(y)^{2^{m-j}} = \kappa_j(y).$$

<sup>2</sup> The initialization of the scheme is assumed to be the 0th “update event”.

222 In other words, the user first uses the current IKE KAS to compute the current key for the desired security label, and then  
223 computes the key for the desired time interval.

224 Note that the node-based KAS that is used for the keys  $\kappa_0(x), \dots, \kappa_m(x)$  has the property defined in (2). Hence, it is not  
225 possible for a user to compute a future key from an existing one.

226 3.1.2. Hash-chain iterative MKAS

In this scheme, keys are chosen in such a way that  $\kappa_j(x)$  can be derived iteratively from  $\kappa_i(x)$ . For each  $x \in L$ , the trusted center chooses a key  $\kappa_m(x) \in \{0, 1\}^l$  and defines  $\kappa_{i-1}(x) = h(\kappa_i(x))$ , where  $h : \{0, 1\}^* \rightarrow \{0, 1\}^l$  is a hash function. We call

$$\kappa_m(x), \kappa_{m-1}(x) = h(\kappa_m(x)), \dots, \kappa_i(x) = h^{m-i}(\kappa_m(x)), \dots, \kappa_0(x) = h^m(\kappa_m(x))$$

227 a *hash chain* for  $x$ . This type of approach has been used in one-time password schemes [11], and means that a user cannot  
228 compute a future key from an existing one.

229 In order to recover  $\kappa_j(y)$ , given  $\kappa_i(x)$ , with  $y \leq x$  and  $j < i$ , the user first computes  $\kappa_i(y)$  using the public information for  
230 the  $i$ th IKE KAS, as before. The user then iteratively derives  $\kappa_j(y)$ , by computing  $\kappa_{i-1}(y) = h(\kappa_i(y)), \dots, \kappa_j(y) = h(\kappa_{j+1}(y))$ .

231 3.1.3. Example

232 Fig. 2 illustrates how an asynchronous scheme would evolve over time, using the poset depicted in Fig. 1a. Labels for  
233 which new keys are generated are represented by unfilled nodes; existing keys are represented by filled nodes; edges that  
234 need to be updated are represented by broken edges. At any moment in time, the public information is used to compute  
235 current keys; previous keys are computed using an Akl-Taylor or hash chain.

236 Prior to any updates, a single key exists for each security label. A user with security label  $d$  is then removed from the  
237 scheme, meaning that  $\kappa(d)$  has to be refreshed. Since  $d$  is the minimal element in the poset, no other keys need to be changed.  
238 In the asynchronous scheme we must recompute the public information  $E_{\kappa(b)}(\kappa(d))$  and  $E_{\kappa(c)}(\kappa(d))$ , the "edge encryptions"  
239 for  $(d, b)$  and  $(d, c)$ . This is depicted in Fig. 2b.

240 Then a user with security label  $b$  is removed from the scheme, meaning that  $\kappa(b)$  has to be refreshed. Since  $d < b$ , we  
241 must also update  $\kappa(d)$ . In addition, we must re-compute all edges from elements immediately greater than either  $b$  or  $d$ .  
242 This is shown in Fig. 2c. Notice that if, instead of removing the user with security label  $b$  from the scheme, we changed her  
243 security label to  $c$ , then the update set is  $\{b\}$  rather than  $\{b, d\}$  (since  $d < c$ ); in this case, we would only update  $\kappa(b)$  and  
244 re-compute  $E_{\kappa(a)}(\kappa(b))$  and  $E_{\kappa(b)}(\kappa(d))$ .

245 3.1.4. MTMA MKAS

246 Finally, we introduce a bounded scheme based on the MTMA labeling. We note that we can regard the set of keys  
247  $\{\kappa_i(l) : l \in L, 0 \leq i \leq m\}$  as being associated with the nodes in the poset  $L \times T$ , where  $T$  is a chain of  $m$  elements. We write  
248  $(x, i)$  for the  $i$ th element in the chain for node  $x$ . It is obvious, therefore, that we can decompose  $L \times T$  into  $|L|$  chains, each  
249 containing  $m + 1$  elements.

250 Hence, we can define an MKAS in which all keys are defined using the Akl-Taylor scheme with the MTMA labeling, where  
251 each node is associated with a different prime. More specifically, the trusted center first obtains a large compound integer  $n$   
252 and chooses a prime  $\alpha(x)$  for each  $x \in L$ . Then the trusted center

- 253 • selects a secret  $s \in \mathbb{Z}_n^*$ ;
- 254 • defines  $\alpha(x, i) = \alpha(x)^{m+1-i}$ ,  $0 \leq i \leq m$ ;
- computes and publishes

$$v(x, i) = \prod_{\substack{y \not\leq x, \\ j > i}} \alpha(y, j);$$




Fig. 2. Key updates in an asynchronous scheme.

- defines  $\kappa_i(x) = s^{v(x,i)}$ .
- If a user  $u$  has  $\kappa_i(x)$ , then she can directly compute

$$\kappa_j(y) = \kappa_i(x)^{\frac{v(y,j)}{v(x,i)}}$$

for any  $y \leq x$  and any  $j \leq i$ . In other words, this scheme has the advantage that key derivation is always performed in a single step, at the expense of requiring  $\mathcal{O}(m|L|)$  storage for public information. Moreover,  $u$  cannot compute  $\kappa_j(y)$  for any  $y \not\leq x$  or any  $j > i$  since this scheme is just an instance of the Akl–Taylor scheme (using the MTMA labeling) for the poset  $L \times T$ .

### 3.2. Unbounded synchronous MKASs

We now assume that the number of updates is not known in advance, and that all keys are updated at the same time. Note that any unbounded synchronous MKAS can be used as a bounded synchronous MKAS.

A feature of the Akl–Taylor KAS is that it is possible to update all the keys simply by choosing a new system secret. No public information needs to change, unlike any IKE KAS. Using this observation, we propose two synchronous MKAS. In the first scheme, user  $u$  is required to keep all earlier versions of  $\kappa(\lambda(u))$ . In the second scheme,  $u$  only requires a single key.

#### 3.2.1. Node-based MKAS

Initially, the trusted center

- runs the RSA key generator to obtain  $n$ ;
- chooses a system secret  $s_0 \in \mathbb{Z}_n^*$  at random;
- defines a mapping  $v : L \rightarrow \mathbb{N}^*$  such that  $v(x) \mid v(y)$  if and only if  $y \leq x$ ;
- publishes  $\{n\} \cup \{v(x) : x \in L\}$ ;
- defines  $\kappa_0(x) = s_0^{v(x)} \pmod n$ .

Following the  $i$ th update event, the trusted center

- chooses a new system secret  $s_i \in \mathbb{Z}_n^*$ , where  $(s_i, s_j) = 1$  for all  $j < i$ ;
- defines  $\kappa_i(x) = s_i^{v(x)} \pmod n$ .

A user with security label  $x$  can derive  $\kappa_j(y)$ ,  $y \leq x$ , using  $\kappa_j(x)$  and the public information  $v(x)$  and  $v(y)$ .

Two significant advantages of this scheme are that the number of update events does not need to be defined before the scheme is instantiated, and the public information does not need to be updated. Of course, the user now accumulates keys as the scheme develops, so private storage requirements increase and the user is required to manage a number of different keys.

#### 3.2.2. Node-based RSA MKAS

A refinement of the scheme described above is to use a deterministic method for generating new system secrets. Of course, we must ensure that future keys cannot easily be computed from existing ones. The scheme is initialized in the same way as in Section 3.2.1, except that the trusted center publishes  $e$  as well as  $n$ . Following the  $i$ th update event, the trusted center

- computes a new system secret  $s_i = s_{i-1}^d \pmod n$ ;
- defines  $\kappa_i(x) = s_i^{v(x)} \pmod n$ .

Note that

$$(\kappa_i(y))^e = (s_i^{v(y)})^e = (s_i^e)^{v(y)} = ((s_{i-1}^d)^e)^{v(y)} = (s_{i-1}^{de})^{v(y)} = s_{i-1}^{v(y)} = \kappa_{i-1}(y).$$

Hence the user now only needs to retain the most recent key: he can derive  $\kappa_j(y)$  from  $\kappa_i(x)$  ( $y \leq x, j \leq i$ ) by first computing  $\kappa_i(y)$  from the public information, and then iteratively computing  $\kappa_{j'}(y)$ ,  $i > j' \geq j$ . The advantages of this scheme are that the user only requires a single key and the public information is fixed. The disadvantage is that key derivation is iterative if the key is for an earlier time interval.

Note that a user cannot feasibly compute a future key from an existing one, as this would require the ability to compute  $d$ , which in turn requires the ability to factorize  $n$ . In other words, the security of future keys relies on the same assumptions as the RSA cryptosystem.

### 3.3. Unbounded asynchronous MKASs

It would be convenient if we could adapt the schemes described in Section 3.2 to construct an unbounded asynchronous MKAS. However, such an adaptation is not possible, in general, because keys are a function of public information and a secret value.

As we have seen in Section 3.2, it is easy to change all keys in an Akl–Taylor scheme by changing the secret value. However, changing the prime associated with one element of  $L$  may cause changes to a number of different keys. To see this, consider the poset in Fig. 1 and the effect of changing the prime associated with  $b$  to 11. This has the effect of changing both  $\kappa(c)$  and  $\kappa(d)$  because  $v(c)$  and  $v(d)$  are both defined in terms of the prime associated with  $b$ .




Fig. 3. Dynamic information flow policies for an unbounded asynchronous MKAS.

301 3.3.1. Iterative RSA MKAS

302 Hence, we need to adopt an alternative approach to develop an unbounded asynchronous MKAS. Informally, we use a  
303 method similar to that in Section 3.2 to create key chains for each  $x \in L$ , and link these chains using an IKE KAS. More formally,  
304 the trusted center

- 305 • runs the RSA key generator to produce  $(n, e, d)$ ;
- 306 • for each label  $x \in L$ , chooses secret  $\kappa_0(x)$  at random from  $\mathbb{Z}_n^*$ ;
- 307 • publishes  $(n, e)$ ;
- 308 • constructs an IKE KAS for  $L$ .

309 Following the  $i$ th key revocation for security label  $x$ , the trusted center

- 310 • computes,  $\kappa_i(x) = \kappa_{i-1}^d \pmod n$ ;
- 311 • replaces  $\kappa_{i-1}(x)$  in the IKE KAS with  $\kappa_i(x)$ ;
- 312 • updates the public information for the IKE KAS.

Note that

$$(\kappa_i(y))^e = ((\kappa_{i-1}(y))^d)^e = \kappa_{i-1}(y). \quad (3)$$

313 Suppose now that user  $u$  with  $\kappa_i(x)$  wishes to derive  $\kappa_j(y)$ ,  $x \geq y, i > j$ . Then  $u$  can compute  $\kappa_j(y)$  using  $\kappa_i(x)$  and the public information  
314 in the IKE KAS, and can now iteratively compute  $\kappa_j(y)$  using the computation defined in (3).

315 3.3.2. Iterative MKAS

316 The second scheme employs a single IKE KAS. The difference is that  $L$  now evolves over time. Using the example depicted in  
317 Fig. 2, each of the three Hasse diagrams would represent different posets, rather than the same poset with different numbers  
318 of keys associated with each of the nodes. Following an update to  $\kappa(d)$ , we would construct an IKE KAS for the poset depicted  
319 in Fig. 3a. Similarly, following an update to  $\kappa(b)$ , we would construct an IKE KAS for the poset depicted in Fig. 3b.

320 More formally, for each node  $x$  in the update set, we need to:

- 321 • insert a new node  $x'$ ;
- 322 • add a new edge  $(x, x')$ ; and
- 323 • for each edge  $(x, y)$ , we need to insert a new edge  $(x', y)$  and delete  $(x, y)$ .

324 A key for  $x'$  needs to be selected and the public information for the new edges computed and published.

325 The advantage of this scheme is its simplicity. The disadvantage is that the number of nodes and edges, and hence the  
326 amount of public information, grows over time.

327 3.4. Summary

328 We summarize the characteristics of our schemes in Fig. 4. The basic design features of each scheme are described in  
329 Fig. 4a. Note, in particular, the MTMA scheme (Section 3.1.4), which uses a single method for both the policy and temporal  
330 dimensions. We only indicate whether a scheme is unbounded or not and whether a scheme is asynchronous or not: any  
331 unbounded scheme can be used as a bounded scheme, and any asynchronous scheme can be used as a synchronous scheme.

332 Fig. 4b summarizes a number of operational characteristics, including those identified by Crampton et al. [7] as being  
333 of interest when assessing key assignment schemes. The derivation column includes the key derivation method within the  
334 policy and temporal dimensions: we write “Dir” to denote direct key derivation and “Ind” to denote indirect key derivation.  
335 Some schemes require public information to be updated following a change of keys; this is indicated in the column headed  
336 “Public Update”. Storage requirements indicate the number of items of data that need to be stored in the worst case following  
337  $u$  updates;  $E$  denotes the edge set in the Hasse diagram of  $L$ . The storage required is proportional to the value given in the  
338 table (the constant of proportionality being determined by the cryptographic primitives chosen to implement the scheme).

Scheme	Dimension		Unbounded	Asynchronous
	Policy	Temporal		
§3.1.1	IKE	Akl-Taylor	No	Yes
§3.1.2	IKE	Hash chain	No	Yes
§3.1.4	Akl-Taylor (MTMA labeling)		No	Yes
§3.2.1	Akl-Taylor	n/a	Yes	No
§3.2.2	Akl-Taylor	RSA	Yes	No
§3.3.1	IKE	RSA	Yes	Yes
§3.3.2	IKE	IKE	Yes	Yes

Scheme	Derivation		Public Update	Storage	
	Policy	Temporal		Public	Private
§3.1.1	Ind	Dir	Yes	E	1
§3.1.2	Ind	Ind	Yes	E	1
§3.1.4	Dir		No	m L	1
§3.2.1	Dir	n/a	No	L	u
§3.2.2	Dir	Ind	No	L	1
§3.3.1	Ind	Ind	Yes	E	1
§3.3.2	Ind	Ind	Yes	u E	1

Fig. 4. Characteristics of our schemes. (a) Design characteristics. (b) Operational characteristics.

339 3.5. Discussion

340 It is difficult to make a definitive statement about which of these schemes is the best, because “best” may depend on the  
 341 context in which a scheme is to be deployed. In an ideal world, a scheme would: have direct key derivation and low storage  
 342 costs; be unbounded and asynchronous; not require updates to public information; and each user would have a single key.  
 343 Of course, none of these schemes have all these features.

344 It is not clear whether asynchronous schemes offer any significant advantages: again, this will probably depend on context.  
 345 It should be noted that many updates for label  $l$  will require updates for all  $l' \leq l$ . In some scenarios, therefore, it may be  
 346 simpler to use a synchronous scheme. Nevertheless, if we can assume (or know) that most of the update activity will be  
 347 for relatively junior security labels, then there may well be a distinct advantage to using this incremental approach. This  
 348 may well be a reasonable assumption: there will probably be fewer users with relatively senior labels, and those users are  
 349 probably less likely to be re-assigned to new security labels or removed from the scheme.

350 Of the bounded schemes, the one based on the MTMA labeling has some appealing features, particularly compared to the  
 351 other two bounded schemes: key derivation is direct and no updates are required for public information. As we showed in  
 352 Section 2.2.3, there exist implementations of the Akl–Taylor scheme with the MTMA labeling that permit fast key derivation.

353 As we have noted, temporal access control policies do not require an asynchronous scheme and generally will not  
 354 require the scheme to be unbounded. Hence, suitable choices for enforcing such policies would be the schemes described in  
 355 Sections 3.1.4 and 3.2.2. The trade-off here is between partial iterative key derivation and amount of public data. Generally,  
 356 the storage required for public data will be an insignificant fraction of the storage available (see [16] for a discussion of such  
 357 considerations), so direct key derivation is likely to be preferred. In contrast, if we wish to support lazy re-encryption, then  
 358 we should choose an asynchronous scheme. We will also probably require an unbounded scheme, so a suitable choice in  
 359 these circumstances would be one of the schemes described in Section 3.3.

360 4. Conclusion

361 Updating keys is known to be an issue in cryptographic access control when the same encryption key is used to encrypt  
 362 many different data objects. This problem occurs in cryptographic storage file systems, where techniques such as key  
 363 regression have been used so that re-encryption of data objects is not necessarily performed when the encryption key  
 364 is refreshed. The problem is even more pressing in key assignment schemes for hierarchical access control, in which a  
 365 number of different keys may need to change at the same point in time, meaning that even higher volumes of data may  
 366 need to be re-encrypted at some point. We have introduced the idea of multi-key assignment schemes and shown that such  
 367 schemes are useful in supporting lazy re-encryption and temporal access control policies. These are the first such schemes  
 368 in the literature, and represent a considerable advance on existing approaches to key updates in KASs.

369 Nevertheless, a number of opportunities for further work exist. One obvious shortcoming of our approach to enforcing  
 370 temporal access control policies is that a user with a key for  $x \in L$  can derive all previous keys for  $x$ . In some cases, it may  
 371 be that a user can only access material from the point at which she starts paying for content. In such cases, it must not be  
 372 possible for the user to derive keys for time points before she started paying. It is not clear whether a simple modification to  
 373 our schemes is sufficient to enforce such requirements. A priority, then, will be to investigate this question.

374 We noted in Section 2.2.4 that formal security analyses of KASs are beginning to appear in the literature. An informal  
375 analysis suggests that the schemes presented in this paper are secure against key recovery. However, none of the schemes  
376 has the property of key indistinguishability. Existing work suggests that a relatively simple extension to our schemes will be  
377 sufficient to obtain the key indistinguishability property. An important part of our future work will be to undertake a formal  
378 analysis of our schemes and to modify them, where necessary, using similar methods to those found in [2,3,10].

379 On a more practical note, it would be of considerable practical interest to adapt the architecture used in cryptographic file  
380 systems and implement a key assignment scheme for a hierarchical file system. Similarly, it would interesting to use XML  
381 encryption to implement hierarchical access control for XML data.

## 382 Acknowledgements

383 The author would like to thank the anonymous referees for their constructive and instructive comments.

## 384 References

- 385 [1] S.G. Akl, P.D. Taylor, Cryptographic solution to a problem of access control in a hierarchy, *ACM Trans. Comput. Syst.* 1 (3) (1983) 239–248.  
386 [2] M.J. Atallah, M. Blanton, N. Fazio, K.B. Frikken, Dynamic and efficient key management for access hierarchies, *ACM Trans. Inform. Syst. Security* 12 (3)  
387 (2009) 1–43.  
388 [3] G. Ateniese, A. De Santis, A.L. Ferrara, B. Masucci, Provably-secure time-bound hierarchical key assignment schemes, in: *Proceedings of the 13th ACM*  
389 *Conference on Computer and Communications Security*, 2006, pp. 288–297.  
390 [4] M. Backes, C. Cachin, A. Oprea, Secure key-updating for lazy revocation, in: *Proceedings of 11th European Symposium on Research in Computer Security*,  
391 2006, pp. 327–346.  
392 [5] D.E. Bell, L. LaPadula, *Secure Computer Systems: Mathematical foundations*, Technical Report MTR-2547, vol. I, Mitre Corporation, Bedford,  
393 Massachusetts, 1973.  
394 [6] E. Bertino, B. Carminati, E. Ferrari, A temporal key management scheme for secure broadcasting of XML documents, in: *Proceedings of the 8th ACM*  
395 *Conference on Computer and Communications Security*, 2002, pp. 31–40.  
396 [7] J. Crampton, K. Martin, P. Wild, On key assignment for hierarchical access control, in: *Proceedings of 19th Computer Security Foundations Workshop*,  
397 2006, pp. 98–111.  
398 [8] B.A. Davey, H.A. Priestley, *Introduction to Lattices and Order*, Cambridge University Press, Cambridge, United Kingdom, 1990.  
399 [9] D.E. Denning, A lattice model of secure information flow, *Communications of the ACM* 19 (5) (1976) 236–243.  
400 [10] K. Fu, S. Kamara, Y. Kohno, Key regression: enabling efficient key distribution for secure distributed storage, in: *Proceedings of the Network and*  
401 *Distributed System Security Symposium, NDSS 2006*, 2006.  
402 [11] N.M. Haller, The S/KEY one-time password system, in: *Proceedings of the 1994 Symposium on Network and Distributed System Security*, 1994, pp.  
403 151–157.  
404 [12] L. Harn, H.Y. Lin, A cryptographic key generation scheme for multilevel data security, *Comput. Security* 9 (6) (1990) 539–546.  
405 [13] M. Kallahalla, E. Riedel, R. Swaminathan, Q. Wang, K. Fu, Plutus: scalable secure file sharing on untrusted storage, in: *Proceedings of the FAST'03*  
406 *Conference on File and Storage Technologies*, 2003, pp. 29–42.  
407 [14] S.J. MacKinnon, P.D. Taylor, H. Meijer, S.G. Akl, An optimal algorithm for assigning cryptographic keys to control access in a hierarchy, *IEEE Trans.*  
408 *Comput. C-34* (9) (1985) 797–802.  
409 [15] R. Rivest, B. Kaliski, The R. Rivest, B. Kaliski, The RSA problem, in: H.C.A. van Tilborg (Ed.), *Encyclopedia of Cryptography and Security*, Springer, 2005,  
410 pp. 532–535.  
411 [16] H. Rowe, J. Crampton, Avoiding key redistribution in key assignment schemes, in: *Proceedings of the Fourth International Workshop on Mathematical*  
412 *Methods, Models, and Architectures for Computer Network Security*, 2007, pp. 127–140.