

Graph-Decomposition-Based Frameworks for Subset-Cover Broadcast Encryption and Efficient Instantiations

Nuttapong Attrapadung and Hideki Imai

Imai Laboratory, Institute of Industrial Science, University of Tokyo,
4-6-1 Komaba, Meguro-ku, Tokyo 153-8505, Japan
nuts@imailab.iis.u-tokyo.ac.jp, imai@iis.u-tokyo.ac.jp

Abstract. We present generic frameworks for constructing efficient broadcast encryption schemes in the subset-cover paradigm, introduced by Naor et.al., based on various key derivation techniques. Our frameworks characterize any instantiation completely to its underlying *graph decompositions*, which are purely combinatorial in nature. This abstracts away the security of each instantiated scheme to be guaranteed by the generic one of the frameworks; thus, gives flexibilities in designing schemes. Behind these are new techniques based on (trapdoor) RSA accumulators utilized to obtain practical performances.

We then give some efficient instantiations from the frameworks. Our first construction improves the currently best schemes, including the one proposed by Goodrich et.al., without any further assumptions (only pseudo-random generators are used) by some factors. The second instantiation, which is the most efficient, is instantiated based on RSA and directly improves the first scheme. Its ciphertext length is of order $O(r)$, the key size is $O(1)$, and its computational cost is $O(n^{1/k} \log^2 n)$ for any (arbitrary large) constant k ; where r and n are the number of revoked users and all users respectively. To the best of our knowledge, this is the first explicit collusion-secure scheme in the literature that achieves both ciphertext size and key size independent of n simultaneously while keeping all other costs efficient, in particular, sub-linear in n . The third scheme improves Gentry and Ramzan's scheme, which itself is more efficient than the above schemes in the aspect of asymptotic computational cost.

Keywords: Broadcast Encryption, Revocation Scheme, Subset-cover, Optimal Key Storage.

1 Introduction

Broadcast encryption (BE) involves 1 broadcaster and n receivers. Each receiver is given a unique private key. The broadcaster is given a private broadcaster key. The broadcaster wishes to broadcast messages to a designated set $P \subseteq N = \{1, \dots, n\}$ of receivers. Any receivers in P should be able to decrypt the broadcast message using only its private key while a coalition $F \subseteq N \setminus P$ (revoked users)

should not be able to do so. Such a scheme is motivated largely by pay-TV systems, the distribution of copyrighted materials such as CD/DVD. Broadcast encryption schemes were first formalized by Fiat and Naor [13]. Since then, many variants of the basic problem were proposed. The arguably most challenging variant is the one which considers the case where P can be an arbitrary subset in N while the collusion is considered the full one, $N \setminus P$, and also that the private key stored by each user is fixed from the initialization time (stateless receiver). The main goal is to construct efficient schemes that satisfy the above variant and require only small size of both the header of broadcast and the private key as a function of n or $r := n - |P|$. The *header* is the encapsulation of session key that is used to encrypt data.

An efficient solution which is considered a ground work to many consequences is the Complete (binary) Subtree scheme (CS) by Naor et al. [18]. Schemes which were considered the current state of the art (before two very recent works, see below) are: (i) Pseudo-random sequences generator (PRSG) based schemes such as the Subset Difference scheme (SD) [18], its refinement—the Layered SD scheme (LSD) [14], and their somewhat generalizations in [4]. (ii) RSA accumulator based schemes such as Asano’s scheme [2], and its optimal generalizations in [3,11]. See Table 1 for the efficiency comparison. No scheme above could achieve simultaneous small header size independent of n , small key size of order $O(\log n)$, while keeping computational cost and all other costs grow only sub-linear in n .

More recently, Goodrich et al. [12] and Wang et al. [20] independently propose more efficient schemes that break the above barrier. In particular, they achieve simultaneously header size of order $O(r)$ and key size of $O(\log n)$, and computational cost of $O(n^{1/k})$ for arbitrary constant k . (In fact, in [20] only the case when $k = 1, 2$ is considered).

In this paper, we propose generic frameworks for constructing broadcast encryption and give some efficient instantiations. One of our instantiations (Instantiation 2 in Table 1) achieves not only small header size as of order $O(r)$ but also small key size as $O(1)$ with no extra non-secret storage, while keeping computational cost $O(n^{1/k} \log^2 n)$ which grows only sub-linear in n . Thus this is the first scheme that achieves header and private key size independent of n while keeping computational cost sub-linear in n , with no extra non-secret storage. The contributions in more detail are described below.

1.1 Our Contributions

In the general subset-cover paradigm of [18], which includes almost all of the above schemes, it has been *implicitly* understood that one can separate the design of such a scheme into two seemingly orthogonal problems namely: designing combinatorial set system which enables subset covering (this step determines the header size), and defining computational key derivation (this step determines the private key size and computational cost). This is first explicitly characterized by Gentry-Ramzan [11] for the case of Akl-Taylor’s RSA based key derivation [1].

Table 1. Comparison among previous schemes and our instantiations. (k is an arbitrary parameter, a is an arbitrary constant).

	Header size		Priv. key size	Comp. cost (bit complexity)	
	Complexity	\leq		Prime-gen	Others
CS [18]	$O(r \log(\frac{n}{r}))$		$\log n + 1$	-	$O(\log \log n)$
PRSG or OWF -based ↓					
SD [18]	$O(r)$	$2r-1$	$O(\log^2 n)$	-	$O(\log n)$
LSD [14]	$O(r)$	$2kr-k$	$O(\log^{1+1/k} n)$	-	$O(\log n)$
GST04 [12]	$O(r)$	$4kr$	$2 \log n$	-	$O(n^{1/k})$
WNR04 [20]	$O(r)$	$4r$	$2 \log n$	-	$O(n^{1/2})$
Instantiation 1	$O(r)$	$2kr$	$\leq \log n + 1$	-	$O(n^{1/k})$
RSA Accumulator -based ↓					
Asano [2]	$O(r \log_a(\frac{n}{r})+r)$		1	$O(2^a \log_a^5 n)$	$O(2^a \log_a^2 n)$
GR04 [11]	$O(r \log_a(\frac{n}{r})+r)$		1	$O(a \log_a^5 n)$	$O(a \log_a^2 n)$
Instantiation 3	$O(r \log_a(\frac{n}{r})+r)$		1	$O(1)$	$O(a \log n)$
$(SD)^{acc}$	$O(r)$	$2r-1$	1	$O(n \log^4 n)$	$O(n)$
Instantiation 2	$O(r)$	$2kr$	1	$O((\log^5 n)/k^5)$	$O((n^{1/k} \log^2 n)/k)$

FRAMEWORK. In this paper, we characterize the two orthogonal components in general. We then explicitly present three generic sub-frameworks for computational key derivation component (*generic* as arbitrary set systems are applicable): PRSG based technique (re-formalizing from [4] so as to be consistent with presentations here), non-trapdoor- and trapdoor- RSA Accumulator based techniques. The non-trapdoor RSA based one is a new optimal generalization of Akl-Taylor’s technique and is further improved by the trapdoor RSA based one.

The main issue is that we characterize three sub-frameworks so that such instantiations in these frameworks and their resulting efficiencies will depend solely on properties related to *graph decompositions* of the set systems being instantiated; while in the same time the security will be guaranteed *automatically* from the general frameworks. The PRSG based framework will be based on *tree decomposition*, and the two RSA based frameworks will be based on *chain decomposition*; both are purely combinatorial. Therefore the whole paradigm abstracts away the computational security issues and reduces the problem to only pure combinatorics. Moreover it allows modularity in designing a scheme: it is a matter of finding a set system which yields a good header size in the first step, and then finding a graph decomposition of that set system that yields a good private key size and computational cost.

As for the generic efficiency characterization, both RSA based frameworks achieve key size of $O(1)$ for all instances. One generic property of the trapdoor based framework that makes it superior to the non-trapdoor based one is that when restricting to the same asymptotic resources and instantiating the same set system (or to be more precise, its hierarchical version and itself respectively), if the non-trapdoor based one allows n users in the scheme, then the trapdoor based one will allow n^k users for any (arbitrary large) constant k . Indeed, the costs due to prime generation are exactly the same (not only asymptotically).

EFFICIENT INSTANTIATIONS. For the combinatorial set system component, all of our schemes are based on new set systems we call Subset Incremental chain (SIC) and Layered-SIC (LSIC) which are designed so to achieve small header size as being $O(r)$ while intrinsically have graph decompositions with good properties. For the computational key derivation component, we instantiate the LSIC set system by presenting their graph decompositions, resulting in various concrete schemes upon each sub-framework as follows. We use the notation $(X)^y$ to denote an instantiation of the set system X using the y -based framework. Denote $\text{LSIC}[k]$ as LSIC with parameter k . Note that $\text{LSIC}[1] = \text{SIC}$.

Instantiation 1 : $(\text{LSIC}[k])^{\text{PRSG}}$. This scheme directly improves the scheme of [12,20] (and it is fair to compare with since the same assumption, PRSG, or equivalently one-way function, was used). In particular it can reduce some overheads, albeit only within constant terms in the worst case: the worst-case key sizes are half of those in [12,20]. Indeed the key size in our scheme is non-uniform among users; some users are even required to store only constant-size keys (cf. Theorem 4, 6, and Eq.(4)). Our scheme also reduces the computational cost from [12], but only in the average case (the worst-case costs are asymptotically the same).

Instantiation 2 : $(\text{LSIC}[k])^{\text{acc}}, (\text{LSIC}[k])^{\text{tacc}}$. Note that $(\text{t})\text{acc}$ is for (trapdoor) accumulator. The performance of this scheme is as mentioned previously. It is the first scheme that achieves header and private key size independent of n while keeping computational cost sub-linear in n , with no extra non-secret storage. The number of primes used per user is optimal as being $O(\log n)$ for $(\text{LSIC}[k])^{\text{acc}}$ and further reduced to $O((\log n)/k)$ for $(\text{LSIC}[k])^{\text{tacc}}$ (so that the on-the-fly prime generation cost is $O((\log^5 n)/k^5)$). Had one used the non-optimal Akl-Taylor's framework as put forth to the context of BE by [2,3,11], it would be $O(n^{1/k} \log n)$ which is super-logarithmic (and the prime generation cost would be $O(n^{1/k} \log^5 n)$).

Instantiation 3 : $(\text{LSIC}[\log_a n])^{\text{tacc}}$. This scheme improves Gentry and Ramzan's scheme [11], which itself is more efficient than the above schemes in the aspect of asymptotic computational cost. Our scheme reduces *poly-logarithmic* cost due to prime generation, which was the dominant cost, to only a *constant* one without affecting the other parameters. Among the *constant-key-size* schemes with header size $O(r \log_a(n/r) + r)$ and no extra non-secret storage, this is the first one in the literature that achieves $O(\log n)$ overall computational cost. (And in fact, ours uses only a constant number of primes). The previous improvement for this class of schemes was done by [11] to improve [2] but only in the constant term involving a . (See Table 1).

1.2 Other Related Works

Very recently, Boneh et.al. [7] propose a *public-key* broadcast encryption scheme which achieves size $O(1)$ for both header and private key. However, the size of the public key to be used by an encrypter, which is also the non-secret storage needed for the decrypter, is $O(n)$. Moreover, the computational cost is $O(n - r)$

(albeit with small coefficient). The second scheme in [7] reduces the non-secret storage size to $O(\sqrt{n})$ but with the price of the increased header size as $O(\sqrt{n})$, and not independent of n anymore. Boneh and Silverberg [6] show that n -linear maps can be used to construct an optimal public-key scheme with constant private key, public key, and header size. However, there are currently no known constructions for such a map for $n > 2$. Most recently, Jho et.al. [15] propose some efficient schemes with small header size when r is not too small. However, their schemes do not enjoy practical *asymptotic* performances as either the header size is $c_1r + c_2n = O(n)$ (for some constant c_1, c_2) or the key size is $\binom{n-1}{k} = O(n^k)$ (where $k \geq 2$) for their best two schemes.

2 Framework and Some Preliminaries

2.1 Framework

We refer to [18] for the definitions and the security notions for private-key broadcast encryption. Now we recap the subset-cover framework [18] separately into two components as follows.

Combinatorial Set System Component. We first redefine a set system which is useful for such a scheme in this framework called **complement-cover set system**. Such a set system is a family of subsets of a universe with the property that every subset of the universe can be efficiently partitioned to a union of some collection of subsets in the family.

Definition 1. (COMPLEMENT-COVER SET SYSTEM). *For a map $c : \mathbb{Z}_{>0}^2 \rightarrow \mathbb{Z}_{>0}$, a set system $\mathcal{S} = \{S_1, \dots, S_m\}$ over a base set $N = \{1, \dots, n\}$ is c -complement-cover if there is a polynomial-time algorithm such that upon input any subset $R \subset N$, outputs $\{S_{i_1}, \dots, S_{i_t}\}$ for some $1 \leq i_1, \dots, i_t \leq m$ such that $N \setminus R = \bigcup_{j=1}^t S_{i_j}$ and that $t \leq c(n, |R|)$. \square*

As usual n, r is the number of all users and revoked users respectively. Such a $c(n, r)$ -complement-cover set system yields a broadcast encryption scheme in the subset-cover framework with the header size $c(n, r)$. The scheme is as follows. The broadcaster defines a subset key for each subset in the family. Each user stores a set of keys in such a way that he can derive all the keys of subsets (in the family) that he is a member. (Thus, the easiest way to do is to store them all. However to reduce the storage of keys, it would be better to store only some and derive the others from those stored keys on the fly. Such derivation patterns are predefined by the broadcaster.) To revoke the set R of users, the broadcaster just let a header to be a session key encrypted with each key of subsets in the partition of $N \setminus R$. Thus the header size is $c(n, r)$. We often denote $c_X(n, r)$ for $c(n, r)$ of the set system \mathcal{S}_X , where X is the name of that set system.

Computational Key Derivation Component. We formalize the specification on key derivations in the context of access control scheme as the following. Denote by $k(S)$ the subset key for $S \in \mathcal{S}$ and $p(u)$ the private key of $u \in N$.

Informally, the security of such a scheme requires that with $\mathbf{p}(u)$, one can derive $\mathbf{k}(S)$ if and only if $u \in S$; moreover, the collusion $N \setminus S$ cannot derive it.

Definition 2 (ACCESS CONTROL SCHEME, AC). . *An Access Control Scheme AC for a set system \mathcal{S} over a base set N is a 2-tuple of polynomial-time algorithms (Keygen, Derive), where:*

Keygen(1^λ): *Takes as input a security parameter 1^λ . It returns all $\mathbf{k}(S_i)$'s, all $\mathbf{p}(u)$'s, and public parameter **pub**.*

Derive($\langle u, \mathbf{p}(u) \rangle, S_i, \mathbf{pub}$): *Takes as input $u \in N$, the key $\mathbf{p}(u)$, $S_i \in \mathcal{S}$, and **pub**. It returns $\mathbf{k}(S_i)$ if $u \in S_i$, or special symbol \perp otherwise. \square*

Naor et al. [18] proved that BE in the subset-cover paradigm whose the access control component is secure in the sense of Key-Indistinguishability (KIND) is secure in the standard notion, namely IND-CCA1. Dodis and Katz [10] use the technique involving multiple encryption to obtain a generic scheme which is IND-CCA2-secure. Key-Intractability (KINT) can be defined analogously. These definitions are captured in the full version of this paper due to limited space here. Also note that there is a simple conversion from KINT-secure scheme to KIND-secure one. Thus KIND or KINT is sufficient for the security of the scheme.

Denote $(X)^y$ to be the access control scheme for set system \mathcal{S}_X that is constructed via AC framework y . Denote $\text{KeySize}_{(X)^y}(u)$ to be the number of keys of u (i.e., $|\mathbf{p}(u)|$, when $\mathbf{p}(u)$ is treated as a set) and $\text{CompCost}_{(X)^y}$ to be the worst-case computational cost for Derive. We also refer $(X)^y$ as a BE scheme via the complement-cover set system \mathcal{S}_X . For any y , $\text{HeaderSize}_{(X)^y}(n, r) = c_X(n, r)$.

2.2 Some Terminology

Viewing Set system as Poset. A set system is partially ordered by the inclusion relation (\subset). Interpreting a set system as a partially ordered set (poset) is useful when defining key derivations in AC. Intuitively, Derive algorithm implies that whenever $S_i \subset S_j$, anyone who can access $\mathbf{k}(S_i)$ is allowed to access $\mathbf{k}(S_j)$.

Terminology for Posets, Graphs. The terminology for posets and graphs used in this paper is quite standard one (cf.[9]) (with some exceptions, see below). Here we review some. A graph is a pair $G = (V, E)$ of sets satisfying $E \subseteq \binom{V}{2}$. V is the set of vertices (or nodes), usually denoted $V(G)$, E is the set of edges, usually denoted $E(G)$. Often, we abuse notation $v \in G$ to mean $v \in V(G)$. A tree is a connected acyclic graph. We often denote $x = \text{parent}_T(y)$ if x is the parent of y in tree T . A directed graph is a pair $G = (V, E)$ of sets satisfying $E \subseteq V \times V$, i.e., an edge is an ordered pair. A directed acyclic graph (DAG) is a directed graph with no directed cycle in it. A notation of chain $x \rightarrow y \rightarrow z$ means a directed graph which $E = \{x, y, z\}$, $V = \{(x, y), (y, z)\}$ and is generalized naturally.

An inclusion poset \mathcal{S} can be represented by a DAG G by setting $V = \mathcal{S}$, $E = \{(S, S') : S \subset S'; S, S' \in \mathcal{S}\}$. This is called the maximal representation, denoted $\text{DAG}_{\max}(\mathcal{S})$. The minimal representation, denoted $\text{DAG}_{\min}(\mathcal{S})$, is the one with $E = \{(S, S') : S \subset_c S'; S, S' \in \mathcal{S}\}$ where we say $S \subset_c S'$ iff there is no $S'' \in \mathcal{S}$ such that $S \subset S'' \subset S'$.




Fig. 1. Toy example 1 and its graph decompositions

In our context¹, a graph decomposition (often denoted \mathcal{G}) of a poset \mathcal{S} is a family of connected subgraphs whose sets of nodes partition the set of all nodes in the $\text{DAG}_{\max}(\mathcal{S})$. (Thus we sometimes say \mathcal{G} is a graph decomposition of $\text{DAG}_{\max}(\mathcal{S})$). When each subgraph is a tree whose edges are directed away from the root, we call it a tree decomposition (often denoted \mathcal{T}). When each graph is a directed chain whose edges are directed in the same direction, we call it a chain decomposition (often denoted \mathcal{C}). An induced graph decomposition is one in which each subgraph is an induced subgraph. Fig.1 shows graph decompositions of the set system for toy example 1, $\mathcal{S}_{\text{toy1}} = \{\{1\}, \{2\}, \{3\}, \{4\}, \{1, 2\}, \{2, 3\}, \{2, 4\}, \{3, 4\}, \{1, 2, 3\}\}$. From now we abuse some notations, often in figures, e.g., writing 12 or 1, 2 instead of $\{1, 2\}$ if it causes no confusion. Note that every chain decomposition is a tree decomposition.

We will fix BT to be the complete binary tree of n leaves labeled $1, \dots, n$ from left to right. The level of node in BT is the distance from root to it. For a fixed node, its left (resp., right) nodes are those nodes with the same level and appear on the left (resp., right). BT will be used only to help defining set systems and should not be confused with the graph representations of posets of set systems.

3 New Set Systems

3.1 Subset Incremental Chain (SIC) Set System

The SIC Set System. For $i, j \in N = \{1, \dots, n\}$ and $i < j$, denote

$$i \rightarrow j := \{\{i\}, \{i, i + 1\}, \dots, \{i, \dots, j\}\},$$

$$i \leftarrow j := \{\{j\}, \{j, j - 1\}, \dots, \{j, \dots, i\}\},$$

and $(i \rightarrow i) = (i \leftarrow i) := \{\{i\}\}$. Consider the binary tree BT. For a node v in BT, let l_v (resp., r_v) be the leftmost (resp., rightmost) leaf under v . We define the set system SIC (of n users) by letting

$$\mathcal{S}_{\text{SIC}} = \bigcup_{v \in \text{BT}_L} (l_v + 1 \leftarrow r_v) \cup \bigcup_{v \in \text{BT}_R} (l_v \rightarrow r_v - 1) \cup (1 \rightarrow n) \cup (2 \leftarrow n), \quad (1)$$

where BT_L (resp., BT_R) are the set of internal nodes which are left (resp., right) children. An informal visual view of \mathcal{S}_{SIC} is shown in Fig.2, where the union of all the collections written there is the only important information.

¹ Our notions for tree and chain decompositions are *not* standard ones (cf.[9]). Instead the notions introduced here might be named as *tree cover* and *path cover*, resp.




Fig. 2. Set system SIC defined by the union of all the collections written at each node

Theorem 1. \mathcal{S}_{SIC} is $(2r)$ -complement-cover set system.

Proof. We call a set of the form $\{i, i + 1, \dots, j\}$ for some $i \leq j$ a consecutive set. We first claim that any consecutive set, say $A = \{i, \dots, j\}$, can be partitioned to no more than 2 sets in \mathcal{S}_{SIC} ; then prove it as follows. Let a be the least common ancestor node of the leaves i and j in BT, denoted $\text{lca}(i, j) = a$. Let s be the least ancestor of a which is in BT_L if $a \in \text{BT}_R$ and which is in BT_R if $a \in \text{BT}_L$. Let x, y be the left and right children of a . First if $i = 1$ then $A \in (1 \rightarrow n) \subseteq \mathcal{S}_{\text{SIC}}$; else if $j = n$ then $A \in (2 \leftarrow n) \subseteq \mathcal{S}_{\text{SIC}}$ (since $2 \leq i$). Now assume $i \neq 1, j \neq n$. We list all possible cases of (i, j) as follows. Let $*$ be an unspecified value.

1. If $(i = l_a; j = *; a \in \text{BT}_L)$ then $A \in (l_s \rightarrow r_s - 1) \subseteq \mathcal{S}_{\text{SIC}}$ (since $i = l_s; j < r_s - 1$; and $s \in \text{BT}_R$),
2. If $(i = *; j = r_a; a \in \text{BT}_R)$ then $A \in (l_s + 1 \leftarrow r_s) \subseteq \mathcal{S}_{\text{SIC}}$ (since $j = r_s; l_s + 1 < i$; and $s \in \text{BT}_L$),
3. If $(i = l_a; j \neq r_a; a \in \text{BT}_R)$ then $A \in (l_a \rightarrow r_a - 1) \subseteq \mathcal{S}_{\text{SIC}}$ (since $j \leq r_a - 1$),
4. If $(i \neq l_a; j = r_a; a \in \text{BT}_L)$ then $A \in (l_a + 1 \leftarrow r_a) \subseteq \mathcal{S}_{\text{SIC}}$ (since $l_a + 1 \leq i$),
5. If $(i \neq l_a; j \neq r_a; a \in *)$ then $A = P \cup Q; P = \{i, \dots, r_x\}, Q = \{l_y, \dots, j\}$, and we have $P, Q \in \mathcal{S}_{\text{SIC}}$ (since
 - $\text{lca}(i, r_x) = x$, thus (i, r_x) will fall to the case 2 or 4 and $P \in \mathcal{S}_{\text{SIC}}$;
 - $\text{lca}(l_y, j) = y$, thus (l_y, j) will fall to the case 1 or 3 thus $Q \in \mathcal{S}_{\text{SIC}}$.

These proved the claim. Now we are back to the proof, it is obvious that $N \setminus R$ can be partitioned to no more than r consecutive sets if 1 or $n \in R$; or to no more than $r + 1$ such sets otherwise. In the former case, the partition size to sets in \mathcal{S}_{SIC} is $\leq 2r$; while in the latter case (where $\{1, \dots, s\}$ and $\{t, \dots, n\}$ for some s, t are included in the partition), it is $\leq 1(1) + 2(r - 1) + 1(1) = 2r$. \square

Intuitively, SIC has graph decompositions with good properties since each collection in the union of Eq.(1) forms a chain of subset. This will become clearer in the next section. The set system LSIC below generalizes SIC.

3.2 Layered SIC (LSIC) Set Systems

The LSIC[k] Set System. We view BT consisting of subtrees (also binary and complete) of $n^{1/k}$ leaves so that there are exactly k layers of such subtrees, where




Fig. 3. Set system $\text{LSIC}[k]$, $k = 2$, as the union of all collections written at each node

$k \log n$. We will call such subtree an “atomic” subtree (to distinguish from other kinds of subtrees in BT). Informally, each atomic subtree contributes sets to $\mathcal{S}_{\text{LSIC}}$ as in the SIC set system for that subtree, albeit each leaf in the subtree represents all the leaves under it in BT. More formally, for node z in BT, let $A_z := \{l_z, l_z + 1, \dots, r_z\}$ (i.e., all the leaves under z). Let us consider the leaves u, v in an atomic subtree where v is some node on the right of u . We denote $u^{(+1)}, u^{(+2)}$ (and so on) be the next one, two (and so on) *right* leaves to u in that atomic subtree. Denote $u^{(-1)}, u^{(-2)}$ analogously. Denote

$$\begin{aligned}
 u \rightarrow v &:= \{A_u, A_u \cup A_{u^{(+1)}}, \dots, A_u \cup \dots \cup A_v\}, \\
 u \leftarrow v &:= \{A_v, A_v \cup A_{v^{(-1)}}, \dots, A_v \cup \dots \cup A_u\}.
 \end{aligned}$$

Let l'_w, r'_w be the leftmost and rightmost leaves under w in the atomic subtree and not w itself; for example, $l'_{\text{root}} = a, r'_{\text{root}} = d$ and $l'_a = 1, r'_a = 4$ in Fig.3. Let \mathbf{A} be the set of all nodes which are the roots of atomic subtrees but excluding the root of BT. We define $\text{LSIC}[k]$ analogously to Eq.(1) by letting

$$\begin{aligned}
 \mathcal{S}_{\text{LSIC}[k]} = \bigcup_{v \in \text{BT}_L \cup \mathbf{A}} (l'_v)^{(+1)} \leftarrow r'_v \cup \bigcup_{v \in \text{BT}_R \cup \mathbf{A}} (l'_v \rightarrow r'_v)^{(-1)} \\
 \cup (l'_{\text{root}} \rightarrow r'_{\text{root}}) \cup (l'_{\text{root}})^{(+1)} \leftarrow r'_{\text{root}}. \quad (2)
 \end{aligned}$$

Intuitively, each $v \in \mathbf{A}$ has two collections $(l'_v)^{(+1)} \leftarrow r'_v, (l'_v \rightarrow r'_v)^{(-1)}$ attached since it is the root of an atomic subtree, which SIC applies (cf. Eq.(1) and Fig.2).

Theorem 2. $\mathcal{S}_{\text{LSIC}[k]}$ is $(2kr)$ -complement-cover set system for a constant k ; and $\mathcal{S}_{\text{LSIC}[\log_a n]}$ is $O(r \log_a(n/r) + r)$ -complement-cover set system for a constant a .

Note that when $k = \log_a n$, from the former claim we already have that $\mathcal{S}_{\text{LSIC}[\log_a n]}$ is $(2r \log_a n)$ -complement-cover, but the claim above gives a sharper bound.

Proof. First we will prove that $\mathcal{S}_{\text{LSIC}[k]}$ is $(2kr)$ -complement-cover. Let ST_R denote the Steiner tree of a set of leaves $R \subseteq N$, i.e., the subtree of BT that consists of all paths from the root to each leaf in R . We call a node v special if $v \in \mathbf{A}$. We “color” a node if it is special but is not in ST_R and all of its special ancestors are in ST_R . Denote \mathbf{C} the set of all color nodes. Hence

$N \setminus R = \bigcup_{v \in C} A_v = \bigcup_{j=1}^k \bigcup_{v \in L_j \cap C} A_v$ where we denote L_j to be the set of all special nodes in the j -th special layer away from root (i.e., at distance $j(\log n)/k$ from the root). It suffices to prove that for each special layer j , the set $Y_j := \bigcup_{v \in L_j \cap C} A_v$ can be partitioned to at most $2r$ sets in the family $\mathcal{S}_{\text{LSIC}}$. Denote x_i to be the number of uncolored special nodes in the i -th atomic subtrees from left to right in this j -th layer. From Theorem 1, it is easy to deduce that Y_j can be partitioned to at most $2(x_1 + x_2 + \dots + x_p)$ sets in $\mathcal{S}_{\text{LSIC}}$, where p is the last atomic subtree in this layer (in fact, $p = n^{(j-i)/k}$). But we have $x_1 + \dots + x_p \leq r$ since the Steiner tree of r leaves passes through all these uncolored special nodes. This proves the claim.

Next we will prove that $\mathcal{S}_{\text{LSIC}[\log_a n]}$ is $O(r \log_a(n/r) + r)$ -complement-cover. We first give the definition of Stratified Subset-Difference set system with each atomic subtree of a leaves (SSD_a): $\mathcal{S}_{\text{SSD}_a} = \{A_u \setminus A_v : u \text{ is an ancestor of } v \text{ in the same atomic subtree}\}$. It is known [11] that $\mathcal{S}_{\text{SSD}_a}$ is $(O(r \log_a(n/r) + r))$ -complement-cover. Using a similar approach as when proving Theorem 1, it is not hard to see that each $A_u \setminus A_v$ can be partitioned to at most 2 sets in $\mathcal{S}_{\text{LSIC}[\log_a n]}$. (The proof is omitted here due to space). Combining these we have that $\text{LSIC}[\log_a n]$ has $c_{\text{LSIC}[\log_a n]}(n, r) = 2c_{\text{SSD}_a}(n, r) = O(r \log_a(n/r) + r)$. \square

4 Key Derivation Based on PRSG

4.1 Reformalize the PRSG Based Framework of [4]

Framework Idea (review). In this framework, we use pseudo-random sequence generators to derive keys from one subset to another. The correctness of access control schemes allows this to be done only if the first set is included in the latter (e.g., $\{1\} \subset \{1, 2\}$). Thus such derivations can be defined in correspondence with directed edges in a graph decomposition of $\text{DAG}_{\max}(\mathcal{S})$, in which all the inclusion relations in \mathcal{S} are included. One exception is that there should be no node with $\text{indegree} > 1$ in any graph in the decomposition since it would imply a collision of PRSG, which should be computable by neither broadcasters nor adversaries. Therefore, all the valid decompositions are *tree decompositions*, of which the class includes all graph decompositions of the poset that allow $\text{indegree} \leq 1$ for all nodes. Each user then stores keys for subsets which he is in and are closest to the root of that tree. For the toy example 1 in Fig.1, our paradigm with the tree decomposition in the figure namely $\mathcal{T}_{\text{toy1}}$ allows the user 2 to store only the keys at 2, 24.

Note that in order to be provably secure in the KIND sense, it is mandatory to make an adaptation so that keys are not derived from another key *directly*. Instead, one should use intermediate keys denoted $t(S)$ for $S \in \mathcal{S}$; how to use this is explained in the construction. This was neglected in many recent schemes that use similar one-way derivation approaches.

The Construction $(X)^{\text{PRSG}}$. This is based solely on a tree decomposition, say \mathcal{T} , of the poset \mathcal{S}_X . The scheme applies to an arbitrary complement-cover set system X .

Keygen : (Subset keys) At a root S of a tree in \mathcal{T} , let $\mathbf{t}(S) \leftarrow \{0, 1\}^\lambda$. For each node S (either root or non-root of a tree in \mathcal{T}) whose all children are S_{i_1}, \dots, S_{i_d} where d is the outdegree of S , we define the following recurrence relation:

$$\mathbf{t}(S_{i_1}) \parallel \dots \parallel \mathbf{t}(S_{i_d}) \parallel \mathbf{k}(S) \leftarrow \text{PRSG}_{d+1}(\mathbf{t}(S)), \quad (3)$$

where $|\mathbf{t}(S_{i_1})| = \dots = |\mathbf{t}(S_{i_d})| = |\mathbf{k}(S)| = \lambda$ bits; $\text{PRSG}_j : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{j\lambda}$.

(User keys) For $u \in N$, we define $\mathbf{p}(u) = \{\mathbf{t}(S) \mid u \in S; u \notin \text{parent}_G(S), G \in \mathcal{T}\}$.
Derive : Find the tree where S is in and then use Eq.(3) to derive $\mathbf{k}(S)$.

Characterizing Efficiency. Let $\text{RN}_{\mathcal{T}}(u) = |\{S \mid u \in S; u \notin \text{parent}_G(S), G \in \mathcal{T}\}|$ and call it the *reachability number* of u in \mathcal{T} (since it is the minimal number of sufficient nodes such that when traversing from these nodes in the edge direction we meet all $S \in \mathcal{S}$ such that $u \in S$). Let $\text{DD}_{\mathcal{T}}$ = the depth of the deepest trees. We have

$$\text{KeySize}_{\langle X \rangle_{\text{prsg}}}(u) = \text{RN}_{\mathcal{T}}(u), \quad \text{CompCost}_{\langle X \rangle_{\text{prsg}}} = \text{DD}_{\mathcal{T}}. \quad (4)$$

Theorem 3. ([4]) $\langle X \rangle^{\text{prsg}}$ is secure in the sense of *KIND* assuming secure PRSG.

4.2 PRSG Based Instantiation for SIC, LSIC

Instantiating SIC. It suffices to define a tree decomposition of \mathcal{S}_{SIC} and the concrete scheme will follow automatically from the general construction of the framework. We choose the following natural one and prove that it is the optimal decomposition for SIC. For $i \leq j \in N$, define a graph $\mathbf{G}(i \rightarrow j)$ as $\{i\} \rightarrow \{i, i+1\} \rightarrow \dots \rightarrow \{i, \dots, j\}$; $\mathbf{G}(i \leftarrow j)$ as $\{j\} \rightarrow \{j, j-1\} \rightarrow \dots \rightarrow \{j, \dots, i\}$. Let $\mathcal{T}_{\text{SIC}} = \{\mathbf{G}(l_v+1 \leftarrow r_v) \mid v \in \text{BT}_{\text{L}}\} \cup \{\mathbf{G}(l_v \rightarrow r_v-1) \mid v \in \text{BT}_{\text{R}}\} \cup \{\mathbf{G}(1 \rightarrow n), \mathbf{G}(2 \leftarrow n)\}$ (5)

Let $\langle x \rangle$ denotes the binary representation of x . We have the following theorem.

Theorem 4. The tree decomposition \mathcal{T}_{SIC} yields minimal $\max_{u \in N} \text{RN}_{\mathcal{T}}(u)$, indeed we have

$$\text{RN}_{\mathcal{T}_{\text{SIC}}}(u) = \begin{cases} \log n + 2 - f(\langle u - 1 \rangle) & ; 2 \leq u \leq n \\ 1 & ; u = 1, \end{cases}$$

where $f(y) :=$ the number of the same consecutive least significant bits of y . In particular, $\max_{u \in N} \text{RN}_{\mathcal{T}_{\text{SIC}}}(u) = \log n + 1$. We also have $\text{DD}_{\mathcal{T}_{\text{SIC}}} = n$.

Proof. We define $F_v = l_v + 1 \leftarrow r_v$ if $v \in \text{BT}_{\text{L}}$ and $l_v \rightarrow r_v - 1$ if $v \in \text{BT}_{\text{R}}$. \mathcal{T}_{SIC} is really a tree decomposition since $\{F_v : v \in \text{BT}_{\text{L}} \cup \text{BT}_{\text{R}}\} \cup \{(1 \rightarrow n), (2 \leftarrow n)\}$ can be proved to be a pairwise non-intersecting family (somewhat straightforwardly). Next we prove the formula for $\text{RN}_{\mathcal{T}_{\text{SIC}}}(u)$. For $u \in N \setminus \{1\}$, only possible trees in \mathcal{T}_{SIC} that u appears are those graphs $\mathbf{G}(F_v)$ for internal nodes v on the path from the leaf u to the root in BT , and $\mathbf{G}(1 \rightarrow n), \mathbf{G}(2 \leftarrow n)$. Each graph $\mathbf{G}(\cdot)$ that u appears contribute one key for u . Thus $\text{RN}_{\mathcal{T}_{\text{SIC}}}(u)$ is at most $(\log n - 1) + 2$. Let

$u, w_1, \dots, w_{\log n}, \text{root}$ be the nodes on that path. Due to symmetry, we assume w.l.o.g. that $w_1, \dots, w_{z-1} \in \text{BT}_L$ and $w_z \in \text{BT}_R$. Now it is easy to see that

- for $1 \leq j \leq z-1$: $\mathsf{G}(F_{w_j}) = \mathsf{G}(l_{w_j} + 1 \leftarrow r_{w_j})$ does not contain $u (= l_{w_j})$;
- for $j = z$: $\mathsf{G}(F_{w_j}) = \mathsf{G}(l_{w_z} \rightarrow r_{w_z} - 1)$ contains $u (= l_{w_z})$;
- for $z < j \leq \log n$: $\mathsf{G}(F_{w_j})$ contains u (since $l_{w_j} < u < r_{w_j}$),

and that $z = f(\langle u - 1 \rangle)$. Thus $\text{RN}_{\mathcal{T}_{\text{SIC}}}(u) = (\log n - 1) + 2 - (f(\langle u - 1 \rangle) - 1)$ as desired. Now we prove that \mathcal{T}_{SIC} is optimal (obtaining minimal $(\max_{u \in N} \text{RN}_{\mathcal{T}}(u))$ among all \mathcal{T} of SIC). Observe that for all \mathcal{T} of SIC, $\sum_{u \in N} \text{RN}_{\mathcal{T}}(u) = \sum_{S \in \mathcal{S}_{\text{SIC}}} |\{u : u \in S, u \notin \text{parent}_G(S), G \in \mathcal{T}\}| \geq |\mathcal{S}_{\text{SIC}}| = n \log n + 1$. Hence $\max_{u \in N} \text{RN}_{\mathcal{T}}(u) \geq \lceil \frac{n \log n + 1}{n} \rceil = \log n + 1$. Our decomposition matches this bound. \square

The number of keys at each user is not uniform as recorded in the corollary below. While sharing some similarities with our scheme, the basic schemes in [12,20] assign one-way chains in both left and right directions at each node in BT while we use only one direction and exploit some symmetries. This can be an intuition as to why we can reduce key size at least 2 times (and up to $\log n$ in the best case, user 1). Those schemes can be considered as instantiations in our framework, but with storage-redundancies in the sense that the set systems extracted from their schemes are sets with repetition. Moreover, the scheme of [12] can also be shown to be derivation-redundant since its derivation graph as exposed in our framework contains loop edges. (See our full paper).

Corollary 5. *In the scheme $(\text{SIC})^{\text{prsg}}$, there are exactly 2^x users who store exactly $x + 2$ keys for $0 \leq x \leq (\log n) - 1$ and exactly 1 user who stores 1 key.*

Instantiating LSIC. Before describing our default tree decomposition of $\mathcal{S}_{\text{LSIC}}$, denoted $\mathcal{T}_{\text{LSIC}[k]}$, we first describe a more straightforward one, denoted $\mathcal{T}'_{\text{LSIC}[k]}$, which is constructed, informally, as the union of all \mathcal{T}_{SIC} applied to each atomic subtree in BT. More formally, we can define $\mathsf{G}(u \rightarrow v)$ for u, v which are leaves in the same atomic subtree, analogously as before, by letting $\mathsf{G}(u \rightarrow v) = A_u \rightarrow A_u \cup A_{u(+1)} \rightarrow \dots \rightarrow (A_u \cup \dots \cup A_v)$, and analogously for $\mathsf{G}(u \leftarrow v)$. Without going into details, we can define $\mathcal{T}'_{\text{LSIC}[k]}$ from Eq.(2) in an analogous way when we defined \mathcal{T}_{SIC} in Eq.(5) from Eq.(1).

Now $\mathcal{T}_{\text{LSIC}[k]}$ is constructed by an observation that $\mathsf{G}(l'_v \rightarrow r'_v{}^{(-1)})$ and $\mathsf{G}(v \rightarrow *)$ can be combined into one chain (and in particular, one tree) since the maximum element in the former, $A_{l'_v} \cup \dots \cup A_{r'_v{}^{(-1)}}$, is included in A_v , the minimum element of the latter. For $v \in \text{BT}_R \cup \{\text{root}\}$, let w_1, \dots, w_m be the sequence of nodes in $\text{BT}_L \cap A$ such that $w_1 = l'_v$; for $1 \leq i \leq m-1$, $w_{i+1} = l'_{w_i}$; and $l_v = l'_{w_m}$, then define $\bar{\mathsf{G}}(l'_v \rightarrow x) := \mathsf{G}(l'_{w_m} \rightarrow r'_{w_m}{}^{(-1)}) \rightarrow \dots \rightarrow \mathsf{G}(l'_{w_1} \rightarrow r'_{w_1}{}^{(-1)}) \rightarrow \mathsf{G}(l'_v \rightarrow x)$ where x is some right node of l'_v . (Here, ‘ \rightarrow ’ means to connect the chains). The definition for $\bar{\mathsf{G}}(x \leftarrow r'_v)$ for $v \in \text{BT}_L \cup \{\text{root}\}$ can be done analogously. Now we define

$$\begin{aligned} \mathcal{T}_{\text{LSIC}[k]} = \{ & \bar{\mathsf{G}}(l'_v{}^{(+1)} \leftarrow r'_v) | v \in \text{BT}_L \} \cup \{ \bar{\mathsf{G}}(l'_v \rightarrow r'_v{}^{(-1)}) | v \in \text{BT}_R \} \\ & \cup \{ \bar{\mathsf{G}}(l'_{\text{root}} \rightarrow r'_{\text{root}}), \bar{\mathsf{G}}(l'_{\text{root}}{}^{(+1)} \rightarrow r'_{\text{root}}) \}. \end{aligned} \quad (6)$$




Fig. 4. The tree decomposition $\mathcal{T}_{\text{LSIC}[k]}$ of the set system $\text{LSIC}[k]$ (see Fig.3). A more simple decomposition $\mathcal{T}'_{\text{LSIC}[k]}$ is the one without the thick red edges.

The abstraction of this decomposition may disguise the simplicity of the scheme; in Fig.4 we thus give an explicit example when $n = 16$ and $k = 2$ (cf. Fig.3).

The following theorem and corollary can be proved by an elementary counting argument based on Theorem 4. We omit the proof to the full version of this paper.

Theorem 6. *The tree decomposition $\mathcal{T}_{\text{LSIC}[k]}$ yields*

$$\text{RN}_{\mathcal{T}_{\text{LSIC}[k]}}(u) = \log n + 1 + k - g_k(\langle u - 1 \rangle)$$

where $g_k(\langle x \rangle) := f(0||\langle x_1 \rangle) + f(b_1||\langle x_2 \rangle) \cdots + f(b_{k-1}||\langle x_k \rangle)$ where we parse $\langle x \rangle$, with padding of 0s on the left so to have length $\log n$ bits, as $\langle x_1 \rangle || \cdots || \langle x_k \rangle$ so that each $\langle x_i \rangle$ has length $(\log n)/k$ bits; b_j is the least significant bit of $\langle x_j \rangle$. In particular, $\max_{u \in N} \text{RN}_{\mathcal{T}_{\text{LSIC}[k]}}(u) = \log n + 1$. We also have $\text{DD}_{\mathcal{T}_{\text{LSIC}[k]}} = kn^{1/k}$.

As an example, user 4 will store 2 keys: $k(1234), k(4)$ (see Fig.4). This can be calculated as $|\text{p}(4)| = 4 + 1 + 2 - (f(0||00) + f(0||11)) = 2$ (Note $\langle 4 - 1 \rangle = 0011$).

Corollary 7. *In $(\text{LSIC}[k])^{\text{prsg}}$, exactly $\sum_{j=0}^{x-1} \binom{k}{j} C(x-1, j, (\log n)/k) 2^{x-1-j}$ users store exactly x keys for $2 \leq x \leq (\log n) + 1$ and exactly 1 user stores 1 key where $C(a, b, c)$ is the number of integer compositions (ordered partitions) of a into b positive integers, each $\leq c$.²*

5 Key Derivation Based on Non-trapdoor RSA

5.1 The New Non-trapdoor RSA Based Framework

Framework Idea. We first briefly review the access control scheme of Akl-Taylor [1]. There, each $S \in \mathcal{S}$ is assigned a publicly known prime. The key of S is defined as $k(S) = s^{\prod_{T: S \not\rightarrow T} p_T}$ modulo an RSA modulus, where s is a secret; and $S \not\rightarrow T$ means (S, T) is not an edge in $\text{DAG}_{\text{max}}(\mathcal{S})$. Each user u just stores $k(\{u\})$. The terms in the exponents are arranged so that even any collusion cannot compute keys that are not supposed to be computable by them. However, the number of primes used in the above schemes are too large as $|\mathcal{S}|$. Such

² For example $C(5, 3, 2) = 3$ since $5 = 1 + 2 + 2 = 2 + 1 + 2 = 2 + 2 + 1$. The exact formula of $C(a, b, c)$ is quite complicated and is shown in [19].

primes will be stored as non-secret storage or derived on-the-fly.³ We propose a new paradigm which makes uses of *prime powers* so that the number of primes used becomes optimal. We will see shortly that assigning prime powers depends essentially on a *chain decomposition* of $\text{DAG}_{\max}(\mathcal{S})$. Indeed, the number of primes used will be exactly the number of chains; and each node in the same chain will correspond to the same prime but with a distinct power. For the toy example 1 in Fig.1, our new paradigm with the chain decomposition $\mathcal{C}_{\text{toy1}}$ will result in only 5 primes used while the Akl-taylor’s needs 9 primes. We will describe how to assign those powers over primes by an incidence matrix. We formalize the notion of incidence matrices that admit a secure scheme as *maximin matrix*:

Maximin Matrix. An $n \times m$ matrix $\{a_{ij}\}$ where $a_{ij} \in \mathbb{Z}_{\geq 0}$ is called a maximin matrix for set system X if for all $S \in \mathcal{S}_X$, there exists $j: 1 \leq j \leq m$ such that $\max_{i \in S} a_{ij} < \min_{i \in N \setminus S} a_{ij}$. We give a formal treatment of RSA functions as accumulators and our construction first, then explain later.

RSA Accumulators. We fix a function $f : \mathbf{U}_f \times \mathbf{E}_f \rightarrow \mathbf{U}_f$ to be an RSA function: $f(x, e) := x^e \bmod \eta$ where $\eta = pq, p = 2p' + 1, q = 2q' + 1$ and p, q, p', q' are distinct odd primes. We restrict that \mathbf{U}_f is the set of quadratic residues and \mathbf{E}_f is the set of primes not equal to p', q' . We say f is generated from an RSA function generator $\text{GRSA}(1^\lambda)$. The function f is an instance of *RSA accumulators*, first proposed in [5], which has a *quasi-commutative* property: for all $x \in \mathbf{U}_f$, and $e_1, e_2 \in \mathbf{E}_f$, $f(f(x, e_1), e_2) = f(f(x, e_2), e_1)$. If $E = \{e_1, \dots, e_h\}$ where each $e_i \in \mathbf{E}_f$, then we denote $f(x, E) := f(f(\dots f(x, e_1), \dots), e_h)$. Note that a set E is threaten as a multi-set, where the repetition of members is important. We thus denote a repetition of a member e which occurs t_e times as $t_e \triangleleft e$. For example, $f(x, \{s \triangleleft e_1, t \triangleleft e_2\}) = x^{(e_1^s \cdot e_2^t)}$.

The Construction $(X)^{\text{acc}}$.

Keygen : Run a GRSA to obtain a description of $f : \mathbf{U}_f \times \mathbf{E}_f \rightarrow \mathbf{U}_f$. Pick a random secret $s \in \mathbf{U}_f$. For $1 \leq j \leq m$, pick an element $p_j \in \mathbf{E}_f$. Let **pub** consist of all p_j ’s and $\{a_{ij}\}$; indeed we let user derive prime p_j only when necessary by predetermining the intervals of those primes (see below). Let

$$\begin{aligned} \mathbf{p}(u) &= f(s, \{a_{uj} \triangleleft p_j : 1 \leq j \leq m\}), \\ \mathbf{k}(S) &= f(s, \{(\max_{i \in S} a_{ij}) \triangleleft p_j : 1 \leq j \leq m\}). \end{aligned} \tag{7}$$

for user $u \in N$ and set $S \in \mathcal{S}_X$.

Derive : Compute $\mathbf{k}(S) = f(\mathbf{p}(u), \{(\max_{i \in S} a_{ij}) \triangleleft p_j : 1 \leq j \leq m\})$.

Theorem 8. $(X)^{\text{acc}}$ is *KINT-secure* assuming the strong RSA assumption.

First it is easy to see that the correctness holds: Derive is computable. Next we will give an intuition as to why for each $S \in \mathcal{S}$, the collusion of all users

³ In the latter, a sequence of integers $\{x_j\}$ is pre-specified by the broadcaster and p_i is defined to be the first prime in $[x_i, x_{i+1})$; the program to recognize $\{x_j\}$ has negligible size (cf. [2]). More primes imply more computational cost on-the-fly.

from $N \setminus S$ cannot compute the key of S . Informally, the best they can do is to obtain the value with the same base s and the exponent term being GCD of all the exponent terms of the keys for users in $N \setminus S$, which is $\prod_{j=1}^m p_j^{\min_{i \in N \setminus S} a_{ij}}$ (by the well-known trick involving using the extended Euclid’s algorithm). To be able to compute the key of S , it must divide $\prod_{j=1}^m p_j^{\max_{i \in S} a_{ij}}$. But this will not happen due to the property of the maximin matrix.

Constructing a Maximin Matrix. Consider a chain decomposition $\mathcal{C} = \{G_1, \dots, G_m\}$ of \mathcal{S}_X . For each chain $G_j : S_1 \rightarrow \dots \rightarrow S_l$, construct j -th column by letting

$$a_{ij} := \begin{cases} 0 & \text{if } i \in S_1 \\ w & \text{if } i \in S_{w+1} \setminus S_w \\ l & \text{otherwise} \end{cases} \quad (8)$$

Proposition 9. *The above construction is a maximin matrix. Moreover, \mathcal{C} with the minimum number of chains will imply the maximin matrix with the minimum m , the number of all primes used.*

Proof. We will prove that the construction by Eq.(8) is a maximin matrix for X . Consider arbitrary $S \in \mathcal{S}$, observe that there is a chain $G_j : S_1 \rightarrow \dots \rightarrow S_l$ and some $w, 0 \leq w \leq l - 1$, such that $S = S_{w+1}$ (since \mathcal{C} is a chain decomposition). For all $i \in S$ we have $0 \leq a_{ij} \leq w$ by the construction. For all $i' \in N \setminus S$ we have $w > a_{i'j}$ also by the construction. This implies $\max_{i \in S} a_{ij} \leq w < \min_{i' \in N \setminus S} a_{i'j}$ which is what we wanted to prove. To prove the second claim, it is sufficient to prove the converse of the first claim: from any maximin matrix for X one can construct a chain decomposition in which the number of chains is less than or equal to the number of columns of the matrix. The proof idea is essentially the same as the first, thus we omit the detail to the full version of this paper. \square

Characterizing Efficiency. We will generate primes on the fly using the technique in [2] (cf. footnote 3). Without going into detail, this technique requires computational cost $O(\log^4 P)$ to generate one prime, and produces each prime of size $O(P \log P)$, where P is the number of all primes needed in such a scheme. In our scheme, $P = m$. Note that only when $P = O(1)$, it is worthless to use this technique; we just store the least P primes (which requires only negligible storage) so the cost for prime generation in this case is $O(1)$.

Using the notation defined earlier, we have that $\text{RN}_{\mathcal{C}}(u)$ represents the number of chains in \mathcal{C} that u appears; and $\text{DD}_{\mathcal{C}}$ represents the length of the longest chain in \mathcal{C} . The number of all chains in \mathcal{C} is $|\mathcal{C}|$ (and= m). We obtain:

$$\text{KeySize}_{(X)\text{acc}}(u) = 1, \quad \text{CompCost}_{(X)\text{acc}} = O(\text{MC}_{\mathcal{C}}^{\text{acc}} + \text{PC}_{\mathcal{C}}^{\text{acc}}),$$

where $\text{MC}_{\mathcal{C}}^{\text{acc}}(u), \text{PC}_{\mathcal{C}}^{\text{acc}}(u)$ are the cost due to Modular exponentiation and on-the-fly Prime generation for user u respectively and $\text{MC}_{\mathcal{C}}^{\text{acc}} := \max_{u \in N} \text{MC}_{\mathcal{C}}^{\text{acc}}(u)$, $\text{PC}_{\mathcal{C}}^{\text{acc}} := \max_{u \in N} \text{PC}_{\mathcal{C}}^{\text{acc}}(u)$. Such costs depend solely on \mathcal{C} and can be characterized as:

$$\text{MC}_{\mathcal{C}}^{\text{acc}}(u) = O(\text{DD}_{\mathcal{C}} \cdot (\log |\mathcal{C}|) \cdot \text{RN}_{\mathcal{C}}(u)), \quad \text{PC}_{\mathcal{C}}^{\text{acc}}(u) = O((\log^4 |\mathcal{C}|) \cdot \text{RN}_{\mathcal{C}}(u)).$$

The analysis are as follows. The cost of modular exponentiation for computing Derive is logarithm in the exponent term which is $\prod_{j=1}^m p_j^{(\max_{i \in S} a_{ij} - a_{uj})}$. To determine its complexity, observe that $\max_{i \in S} a_{ij} = a_{uj}$ for all but only $\text{RN}_C(u)$ terms of j due to Eq.(8) and the fact that u appears only $\text{RN}_C(u)$ chains. Also, observe that $\max_{i \in S} a_{ij} - a_{uj} \leq \text{DD}_C$ due to Eq.(8). Each p_j is $O(m \log m)$, hence has bit length $O(\log m)$. Combining these, we get $\text{MC}_C^{\text{acc}}(u)$ as above. The cost for prime-generation above follows from the fact that the number of primes to be generated when deriving keys are $\text{RN}_C(u)$.

Remark 1. The MC of our scheme is asymptotically optimal among all non-trapdoor RSA-accumulator based paradigms (if there are any others) since it matches the lower bound in [11], which states that the optimal MC is of the same order as the number of subsets (in the set system) that one user is in, albeit here we calculate in bit complexity which includes the size of primes.

Remark 2. The Akl-Taylor’s scheme [1] is a special case of our framework where the trivial chain decomposition (the collection of all one-node chains) is used.

5.2 Non-trapdoor RSA Based Instantiation for SIC, LSIC

Instantiating SIC, LSIC. We will state the result for LSIC so that the result for SIC can be obtained by setting $k = 1$. It suffices to define a chain decomposition of $\mathcal{S}_{\text{LSIC}[k]}$ and the concrete scheme will follow automatically. We choose a chain decomposition $\mathcal{C}_{\text{LSIC}[k]} = \mathcal{T}_{\text{LSIC}[k]}$ defined in Eq.(6). (Note that it is obvious that $\mathcal{T}_{\text{LSIC}[k]}$ was also a chain decomposition). A concrete example for $(\text{SIC})^{\text{acc}}$ is shown in Fig.5 for $n = 8$. As an example, the subset key $k(567) = s(p_1^6 p_2^1 p_3^1 p_4^3 p_5^2 p_6^1 p_7^1 p_8^3)$.

The following result follows directly from Theorem 4, 6 and the generic efficiency characterization of the framework with the fact that $|\mathcal{C}_{\text{LSIC}[k]}| = n$.

Corollary 10. $\text{MC}_{\mathcal{C}_{\text{LSIC}[k]}}^{\text{acc}} = O(kn^{1/k} \log^2 n)$ and $\text{PC}_{\mathcal{C}_{\text{LSIC}[k]}}^{\text{acc}} = O(\log^5 n)$.

Scheme $(\text{LSIC}[k])^{\text{acc}}$ has computational cost $O(\max\{kn^{1/k} \log^2 n, \log^5 n\})$. For trillion users ($n = 10^{12}$), choose k as low as 4 we have $4n^{1/4} \log^2 n < \log^5 n$ so that the computational cost is dominant by the latter, which is roughly as in Asano’s scheme (but ours enjoy exceptionally lower header size).




Fig. 5. Instantiating SIC ($n = 8$) by the non-trapdoor RSA accumulator based framework

Remark 3. If we instantiate with with Akl-Taylor’s, its chain decomposition has $\max_{u \in N} h_u = O(n^{1/k} \log n)$, and $m = O(2^k \cdot n^{1/k} (\log n)/k)$. Thus $PC = O(n^{1/k} \log^5(n))$, which is much worse than ours, $O(\log^5 n)$. Moreover, this cost always dominates over the optimal MC for LSIC, $O(n^{1/k} \log^2 n)$.

6 Key Derivation Based on Trapdoor RSA Accumulator

6.1 The New Trapdoor RSA Based Framework

Framework Idea. The framework in this section is applicable to a class of posets that we call *tree-stratifiable posets*. Informally, such a poset of this type is defined as one which can be considered as formed by a tree hierarchy of atomic posets (not necessarily homogeneous), as shown in Fig.6. There, the graph decomposition $\mathcal{G} = \{G_x, G_y, G_z, \dots\}$ is said to form a hierarchy represented by tree \mathcal{H} where $V(\mathcal{H}) = \{x, y, z, \dots\}$. Intuitively, such a graph decomposition is said to form a hierarchy if all the inclusion relations from every node in a lower subgraph (one with a lower index in the hierarchy), say G_y in the figure, to the next upper one in the hierarchy, G_x , are via a unique minimal node in that upper subgraph. Denote this minimal node as M_{G_y} . We will put a “dummy node” in each subgraph so that it will be the “representative” of that poset to reach that unique minimal node in the upper poset. (In the figure, the dummy node is D_{G_y} for subgraph G_y to reach M_{G_y}).

The idea for key derivations are as follows. First we define the key for each node in the highest sub-poset in the hierarchy by using the RSA-based framework in the last section. Recursively in a top-down fashion, we will define the set of keys corresponding to each lower sub-poset in the hierarchy. At some point, the set of keys for the nodes in G_x are defined. Then we define the “dummy key” for the dummy node in a next lower level sub-poset by applying a random permutation perm (w.l.o.g we will use the reverse direction) to the key of the minimal element in that upper sub-poset that it connects, that is, $k(D_{G_y}) = \text{perm}^{-1}(k(M_{G_y}))$. To define keys for the other nodes in this lower sub-poset (at G_y), we will again use the RSA-based framework for that sub-poset. However, this time the key for the dummy node has been already determined, while all the keys must agree with the relations of $(G'_y)^{\text{acc}}$, where G'_y is the modified




Fig. 6. The underlying idea for the trapdoor RSA based framework

subgraph that includes the dummy node, i.e., the relation of keys as defined in Eq.(7) instantiated to a poset that has G'_y as its representation. To solve this, it suffices to use the *trapdoor* of RSA. In this way, we can define keys recursively until reaching the lowest sub-posets. Users, on the other hand, do not have to use trapdoor since they only compute keys in the bottom-up fashion. Note that $(\text{perm}, \text{perm}^{-1})$ is a public permutation, such as any block cipher with a fixed known key. We will model perm as an ideal random permutation in the security proof (the random permutation model).

The idea of reducing the whole poset by instantiating RSA-based framework in each sub-poset results in the use of only small number of primes for the overall scheme since the same set of primes can be used across different instantiations for different sub-posets.

To formalize this, we first define some more notations. For a directed graph G , denote $V_{\min}(G)$ the set of all minimal elements of poset \mathcal{S} such that $\text{DAG}_{\min}(\mathcal{S}) = G$. $V_{\max}(G)$ is defined analogously. The definition below captures what we have explained in the framework idea. Essentially, the bijection π below maps $G_x \mapsto x$.

Definition 3. (TREE-STRATIFIABLE POSET) *An inclusion poset \mathcal{S} is called tree-stratifiable poset iff there exist an induced graph decomposition \mathcal{G} of \mathcal{S} and a tree \mathcal{H} with a bijection $\pi : \mathcal{G} \rightarrow V(\mathcal{H})$ such that for each $G \in \mathcal{G}$ if we define G' by letting $V(G') = V(G) \cup \{D_G\}$ and $E(G') = E(G) \cup \{(S, D_G) : S \in V_{\max}(G)\}$ where D_G is a dummy node; define $M_G := \bigcup_{S \in V_{\max}(G)} S$; and define a graph \mathcal{W} by letting $V(\mathcal{W}) = \bigcup_{G \in \mathcal{G}} V(G')$ and $E(\mathcal{W}) = \bigcup_{G \in \mathcal{G}} (E(G) \cup \{(D_G, M_G)\})$, then we have that (1) for all $G \in \mathcal{G}$, $M_G \in V_{\min}(\pi^{-1}(\text{parent}_{\mathcal{H}}(\pi(G))))$ and (2) $E(\text{DAG}_{\min}(\mathcal{S})) \subseteq E(\text{DAG}_{\max}(\mathcal{W}))$. \square*

Trapdoor RSA Accumulators. A trapdoor RSA function generator G_{tRSA} is the one that works exactly the same as G_{RSA} but in addition also outputs the *trapdoor* td which is $\phi(\eta)$ where ϕ is the Euler's phi function. With td , given the description of f , any $y \in \mathcal{U}_f$, and a (multi-)set of accumulated values E , one can efficiently compute $x \in \mathcal{U}_f$ such that $f(x, E) = y$. Denote such x by $f_{\text{td}}(y, E^{-1})$.

Towards formalizing the construction, we “normalize” each sub-poset $G \in \mathcal{G}$ so that its base set will be $B_G = \{1, \dots, |V_{\min}(G')|\}$ as follows. Construct $\gamma : V(G') \rightarrow 2^{B_G}$ by first picking an injective map $\tilde{\gamma} : V_{\min}(G) \rightarrow B_G$ then define for $S \in V(G')$, $\gamma : S \mapsto \{\tilde{\gamma}(U) : U \in V_{\min}(G), U \subseteq S\}$. Let $\mathcal{S}_G = \gamma(V(G'))$ (the set of all images by γ from $V(G')$) be the set system with the base set B_G .

The Construction $(X)^{\text{tacc}}$. For simplicity we will consider homogeneously stratifiable poset, i.e., each \mathcal{S}_G is isomorphic to each other (in the sense that its corresponding DAG is isomorphic), say the set system \mathcal{Y} . Let $\{a_{ij}\}_{1 \leq i \leq d, 1 \leq j \leq m}$ be a maximin matrix for set system \mathcal{Y} , where d is the cardinality of its base set.

Keygen : Run a G_{tRSA} to obtain a description of $f : \mathcal{U}_f \times \mathcal{E}_f \rightarrow \mathcal{U}_f$ and trapdoor td . For $1 \leq j \leq m$, pick an element $p_j \in \mathcal{E}_f$. Let perm and perm^{-1} be a publicly available permutation mapping $\mathcal{U}_f \rightarrow \mathcal{U}_f$. Let pub consist of all p_j 's

and $\{a_{ij}\}$. Pick a random $t \in U_f$. Define keys recursively in a top-down fashion in the tree \mathcal{H} :

[Top]. At the subgraph $G_{\text{root}} \in \mathcal{G}$, where root is the root of \mathcal{H} , by definition we have $N = M_{G_{\text{root}}}$. We let $k(N) = k(M_{G_{\text{root}}}) = t$.

[Intermediate]. At each atomic subgraph $G \in \mathcal{G}$, the key $k(M_G)$ is previously determined. Define the key for the dummy node: $k(D_G) = \text{perm}^{-1}(k(M_G))$. By using the trapdoor td and $k(D_G)$, we solve Eq.(11) by setting $S = D_G$ (thus $\gamma(S) = B_G$) to determine the secret s_G , i.e.,

$$s_G = f_{\text{td}}(k(D_G), \{(\max_{i \in B_G} a_{ij}) \triangleleft p_j : 1 \leq j \leq m\}^{-1}). \quad (9)$$

Then we define the key at each element in this subgraph, $S \in V(G)$, by:

$$k(S) = f(s_G, \{a_{\gamma(S),j} \triangleleft p_j : 1 \leq j \leq m\}) \quad (\text{for } S \in V_{\min}(G)), \quad (10)$$

$$k(S) = f(s_G, \{(\max_{i \in \gamma(S)} a_{ij}) \triangleleft p_j : 1 \leq j \leq m\}) \quad (\text{for } S \in V(G)). \quad (11)$$

[Bottom]. For each $u \in N$, we let $p(u) = k(\{u\})$.

Derive : Compute from the relations given in Eq.(9),(10),(11) but in the *bottom-up* fashion by using applications of $f(\cdot, \cdot)$, $\text{perm}(\cdot)$ starting from $f(p(u), \cdot)$. Note that td is not required to do this.

Theorem 11. $(X)^{\text{tacc}}$ is KINT-secure in the random permutation model (perm as an ideal random permutation), assuming the strong RSA assumption.

Characterizing Efficiency. If the set system X of n users is tree-stratifiable homogeneously into a set system Y of d users with the tree \mathcal{H} then

$$\text{KeySize}_{(X)^{\text{tacc}}}(u) = 1, \quad \text{CompCost}_{(X)^{\text{tacc}}} = O(\text{MC}_X^{\text{tacc}} + \text{PC}_X^{\text{tacc}}),$$

where the cost from modular exponentiation and prime generation are depended solely on both \mathcal{H}, Y and only Y respectively, and can be characterized as:

$$\text{MC}_X^{\text{tacc}} = h_{\mathcal{H}} \cdot \text{MC}_{\mathcal{C}_Y}^{\text{acc}}, \quad \text{PC}_X^{\text{tacc}} = \text{PC}_{\mathcal{C}_Y}^{\text{acc}}, \quad (12)$$

where $h_{\mathcal{H}}$ is the deepest depth of \mathcal{H} . The first claim follows from the fact that a user has to compute Eq.(11) for at most $h_{\mathcal{H}}$ times. The second claim is from the fact that we reuse the same set of primes across sub-posets. There is also the cost due to applications of perm , which is $O(h_{\mathcal{H}})$, but this is suppressed by MC .

Generic Application. We now confine our interest to the case where \mathcal{H} is the balanced completed $n^{1/k}$ -ary tree of depth $h_{\mathcal{H}} = k$. This forces the base sets of Y and X to have cardinality $n^{1/k}$ and n respectively. In this case we say $X = \text{hier}_k(Y)$. The operation hier_k is well-defined and can be thought as the converse direction of tree-stratification; thus, from any poset Z one can construct a tree-stratifiable poset, namely $\text{hier}_k(Z)$, by first scaling down the cardinality of the

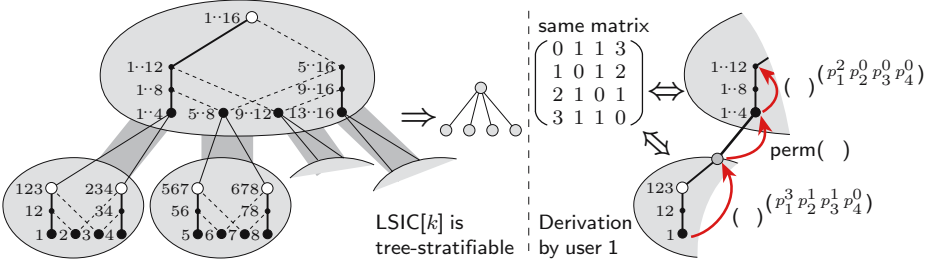


Fig. 7. Instantiating $LSIC[k]$ ($n = 16, k = 2$, see Fig.3) by trapdoor RSA based framework

base set of Z to $n^{1/k}$. (Since usually any set system is originally defined in term of n). We write $Z(n^{1/k})$ to emphasize the cardinality of base set. The point is that when k is a constant, Eq.(12) allows one to construct a full scheme of n users but with exactly the same asymptotic performances as those of $(Z(n^{1/k}))^{acc}$, which is a “scaled-down” scheme, in both parameters MC, PC! Moreover, if $c_{Z(n)}(n, r) = O(r)$ then we can show that $c_{hier_k(Z(n^{1/k}))}(n, r) = O(kr) = O(r)$ (by exactly the same proof as that of Theorem 2); therefore, HeaderSize is also unaffected.

6.2 Trapdoor RSA Based Instantiation for LSIC

It is easy to see that $LSIC[k]$ is tree-stratifiable since $LSIC[k] = hier_k(SIC(n^{1/k}))$. (We could have define $LSIC$ via $hier$ operation rather than directly in Sec.3.2). An example is shown in Fig.7. From the efficiency characterization we have:

Corollary 12. (i) $MC_{LSIC[k]}^{tacc} = O(n^{1/k}(\log^2 n)/k)$, $PC_{LSIC[k]}^{tacc} = O((\log^5 n)/k^5)$.
 (ii) $MC_{LSIC[\log_a n]}^{tacc} = O(a \log a \log n)$, $PC_{LSIC[\log_a n]}^{tacc} = O(1)$.

Proof. See that $MC_{SIC(n^{1/k})}^{acc} = O((n^{1/k} \log^2 n)/k^2)$, $PC_{SIC(n^{1/k})}^{acc} = O((\log^5 n)/k^5)$; and $MC_{SIC(a)}^{acc} = O(a \log^2 a)$, $PC_{SIC(a)}^{acc} = O(1)$. (In fact, for the case $SIC(a)$, the maximum number of primes used per user is $\log a + 1$, a small constant). \square

7 Concluding Remarks

We presented three generic frameworks for constructing broadcast encryption and give some efficient instantiations. Almost all subset-cover broadcast encryption schemes based on PRSG (or one-way function) or RSA accumulator in the literature can be rewritten as instantiations in our paradigms. In fact, [18,14,17,4,12,20,15] can be viewed as PRSG-instantiated schemes and [2,3,11] are non-trapdoor-RSA-instantiated schemes from our frameworks.

The whole paradigm abstracts away the computational security issues and reduces the problem to only pure combinatorics. We leave as an open problem the question of showing any combinatorial bound from the efficiency characterization in each sub-framework. Note that the previous bounds for broadcast encryption [16] are done in the setting where no key derivation is involved.

Acknowledgements. We would like to thank Kazukuni Kobara, Ryo Nojima, and also anonymous reviewers for their valuable comments on earlier versions.

References

1. S. G. Akl, P. D. Taylor, "Cryptographic Solution to a Problem of Access Control in a Hierarchy," *ACM Transactions on Computer Systems*, Vol. 1, No. 3 (1983), pp. 239-248.
2. T. Asano, "A Revocation Scheme with Minimal Storage at Receivers," *ASIACRYPT 2002*, LNCS 2501, pp. 433-450.
3. N. Attrapadung, K. Kobara, H. Imai, "Broadcast Encryption with Short Keys and Transmissions," *ACM Workshop on Digital Rights Management*, 2003.
4. N. Attrapadung, K. Kobara, H. Imai, "Sequential Key Derivation Patterns for Broadcast Encryption and Key Predistribution Schemes," *ASIACRYPT 2003*, LNCS 2894, pp. 374-391.
5. J. Benaloh, M. de Mare, "One-way accumulators: A decentralized alternative to digital signatures," *EUROCRYPT 1993*, LNCS 765, pp. 274-285.
6. D. Boneh, A. Silverberg, "Applications of Multilinear Forms to Cryptography," *Contemporary Mathematics*, 324, pp. 7190, (2003).
7. D. Boneh, C. Gentry, B. Waters, "Collusion Resistant Broadcast Encryption With Short Ciphertexts and Private Keys," *CRYPTO 2005*, LNCS 3621 (to appear).
8. G. C. Chick, S. E. Tavares, "Flexible Access Control with Master Keys," *CRYPTO89*, LNCS 435, pp. 316-322.
9. R. Diestel, "Graph theory," 2nd ed., *Graduate texts in mathematics* 173, (2000).
10. Y. Dodis, J. Katz, "Chosen-Ciphertext Security of Multiple Encryption," *TCC 2005*, LNCS 3378, pp. 188-209.
11. C. Gentry, Z. Ramzan, "RSA Accumulator Based Broadcast Encryption," *ISC 2004*, LNCS 3225, pp. 73-86.
12. M. T. Goodrich, J.Z. Sun, R. Tamassia, "Efficient Tree-Based Revocation in Groups of Low-State Devices," *CRYPTO 2004*, LNCS 3152, pp. 511-527.
13. A. Fiat, M. Naor, "Broadcast Encryption," *CRYPTO 1993*, LNCS 0773, pp. 480-491.
14. D. Halevy, A. Shamir, "The LSD Broadcast Encryption Scheme," *CRYPTO 2002*, LNCS 2442, pp. 47-60.
15. N. Jho, J.Y. Hwang, J.H. Cheon, M.H. Kim, D.H. Lee, E.S. Yoo, "One-Way Chain Based Broadcast Encryption Schemes," *EUROCRYPT 2005*. LNCS 3494, pp. 559-574.
16. M. Luby, J. Staddon, "Combinatorial Bounds for Broadcast Encryption," *EUROCRYPT 1998*, LNCS 1403, pp. 512-526.
17. M.J. Mihaljevic, "Key Management Schemes for Stateless Receivers Based on Time Varying Heterogeneous Logical Key Hierarchy," *ASIACRYPT 2003*, LNCS 2894, pp. 137-154.
18. D. Naor, M. Naor and J. Lotspiech, "Revocation and Tracing Schemes for Stateless Receivers," *CRYPTO 2001*, LNCS 2139, pp. 41-62.
19. Z. Star, "An Asymptotic Formula in the Theory of Compositions," *Aequationes Math* (1976).
20. P. Wang, P. Ning, D.S. Reeves, "Storage-Efficient Stateless Group Key Revocation," *ISC 2004*, LNCS 3225, pp. 25-38.