

Grafting Key Trees: Efficient Key Management for Overlapping Groups

Joël Alwen¹, Benedikt Auerbach²[0000–0002–7553–6606], Mirza Ahad Baig²[0000–0003–3650–7893], Miguel Cueto Noval², Karen Klein³, Guillermo Pascual-Perez²[0000–0001–8630–415X], Krzysztof Pietrzak², and Michael Walter⁴[0000–0003–3186–2482] *

¹ AWS Wickr

alwenjo@amazon.com

² IST Austria, Klosterneuburg, Austria

{bauerbac, mbaig, mcuetono, gpascual, pietrzak}@ist.ac.at

³ ETH Zurich, Switzerland

karen.h.klein@protonmail.com

⁴ Zama, Paris

michael.walter@zama.ai

Abstract. Key trees are often the best solution in terms of transmission cost and storage requirements for managing keys in a setting where a group needs to share a secret key, while being able to efficiently rotate the key material of users (in order to recover from a potential compromise, or to add or remove users). Applications include multicast encryption protocols like LKH (Logical Key Hierarchies) or group messaging like the current IETF proposal TreeKEM.

A key tree is a (typically balanced) binary tree, where each node is identified with a key: leaf nodes hold users' secret keys while the root is the shared group key. For a group of size N , each user just holds $\log(N)$ keys (the keys on the path from its leaf to the root) and its entire key material can be rotated by broadcasting $2 \log(N)$ ciphertexts (encrypting each fresh key on the path under the keys of its parents).

In this work we consider the natural setting where we have many groups with partially overlapping sets of users, and ask if we can find solutions where the cost of rotating a key is better than in the trivial one where we have a separate key tree for each group.

* Benedikt Auerbach, Mirza Ahad Baig, and Krzysztof Pietrzak have received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (682815 - TOCNeT); Karen Klein was supported in part by ERC CoG grant 724307 and conducted part of this work at IST Austria, funded by the ERC under the European Union's Horizon 2020 research and innovation programme (682815 - TOCNeT); Guillermo Pascual-Perez was funded by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No.665385; Michael Walter conducted part of this work at IST Austria, funded by the ERC under the European Union's Horizon 2020 research and innovation programme (682815 - TOCNeT).

We show that in an asymptotic setting (where the number m of groups is fixed while the number N of users grows) there exist more general key graphs whose cost converges to the cost of a single group, thus saving a factor linear in the number of groups over the trivial solution.

As our asymptotic “solution” converges very slowly and performs poorly on concrete examples, we propose an algorithm that uses a natural heuristic to compute a key graph for any given group structure. Our algorithm combines two greedy algorithms, and is thus very efficient: it first converts the group structure into a “lattice graph”, which is then turned into a key graph by repeatedly applying the algorithm for constructing a Huffman code.

To better understand how far our proposal is from an optimal solution, we prove lower bounds on the update cost of continuous group-key agreement and multicast encryption in a symbolic model admitting (asymmetric) encryption, pseudorandom generators, and secret sharing as building blocks.

1 Introduction

Key trees. In various group communication settings, including multicast encryption [15, 16, 7] or group messaging protocols [4, 8], the most efficient constructions use a binary tree structure to manage keys. The general idea is to consider a balanced binary tree with edges directed from leaves to the root. One then identifies each node v with a key k_v (of a symmetric encryption scheme for multicast encryption and a public-key encryption scheme for group messaging). Each edge (u, v) corresponds to a ciphertext $\text{Enc}_{k_u}(k_v)$ and each leaf node v with a user u_v . A user u_v will know the (secret) key k_v , and from the ciphertexts can then retrieve all the keys on the path from its leaf to the root ε . The root key k_ε is thus known to all users, and can be used for secure communication to or among the group members.

What makes this tree structure so appealing is the fact that in a group of size N , the key material of a user u can be completely rotated by replacing only the keys on the path from u to ε , which in a balanced tree has length at most $d = \lceil \log(N) \rceil$. Moreover, as the nodes in a tree all have indegree two, one only needs to compute two fresh ciphertexts for each new key (in practice just one as the new keys can be derived via a hash-chain).

These aspects are important as the number of keys a user requires basically defines the communication and computational efficiency of a key rotation, which is the main operation performed to add or remove users, or for a user to update their keys in order to recover from a potential compromise.

Groups. In this work we consider an extension of this setting to multiple groups. We are given a base set $[N] = \{1, \dots, N\}$ of users with a set system $\mathcal{S} = \{S_1, \dots, S_k\}$ (each $S_i \subseteq [N]$), and we ask for a key managing structure such that for any set $S_i \in \mathcal{S}$, the users in S_i share a group key. This is a natural and well motivated setting; consider for example a university, where one might want to have a shared key for all students attending particular lectures.

A trivial solution to this problem is to simply use a different key-tree for every group S_i , in this work we explore more efficient solutions.

Key-graphs beyond trees. For a set system \mathcal{S} as above, instead of using disjoint trees, any directed acyclic graph (DAG) $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with the following properties is sufficient to maintain group keys:

1. Every user $i \in [N]$ corresponds to a source v_i (a node of indegree 0).
2. Every group $S_i \in \mathcal{S}$ corresponds to a sink v_{S_i} (a node of outdegree 0).
3. For every $S_i \in \mathcal{S}$ and $j \in [N]$, there is a directed path from v_j to v_{S_i} if and only if $j \in S_i$.
4. The indegree of any node is at most 2.

The first three properties ensure that any user $j \in [N]$ can learn the keys associated with the nodes of groups they are in. The last property is not really necessary, but it is without loss of generality in the sense that any graph can be turned into a graph with at most as large update cost (as we show in Section 3), and where every node other than the leaves has indegree at most 2. We call this a key-derivation graph for \mathcal{S} .

Update cost. If we rotate the keys of a user i we need to replace all keys that can be reached from v_i , which we denote by $\mathcal{D}(v_i)$, and encrypt each new key under the keys of its co-path. We thus define the update cost of a user $i \in [N]$ as $\sum_{v \in \mathcal{D}(v_i)} (\text{indeg}(v) - 1)$, which with item 4 above roughly simplifies to the number of v_i 's descendants $|\mathcal{D}(v_i)|$. The update cost $\text{Upd}(\mathcal{G})$ of a DAG \mathcal{G} is the sum over the update cost of all its leaves, which is proportional to the average update cost of users.

Towards constructing more efficient key-derivation schemes when we have multiple overlapping groups, we thus address the problem of determining how small the update cost of a key-derivation for a given set system $\mathcal{S} = \{S_1, \dots, S_k\}$ over $[N]$ can be, and how to find graphs which achieve, or at least come close to, this minimum.

Our contributions. We look at this problem from two perspectives. To get an insight on how much can be saved compared to the trivial solution, we first adapt a qualitative, asymptotic perspective, where we assume a fixed set system, but the number of users N goes to infinity while the relative size of the sets and intersections remains the same. We prove a lower bound on the update cost in this setting and give an algorithm computing graphs matching this bound.

As this solution turns out to be far from optimal for certain concrete set systems, we then also look at a quantitative non-asymptotic setting, where we consider concrete bounds and care about things like additive constants. We propose an algorithm that seems better equipped to handle such systems and prove upper and lower bounds on the update costs of graphs generated by it. Finally, we prove lower bounds on the update cost of any continuous group-key agreement scheme and multicast encryption scheme in a symbolic model.

1.1 The asymptotic setting

Given a set system $\mathcal{S} = (S_1, \dots, S_k)$ over some base set $[n]$, we let $\mathcal{S}(N)$ denote the system with base set $[N]$ we get by considering each element in S with multiplicity N/n . E.g. if $\mathcal{S} = (\{1, 2\}, \{2, 3\})$ then $\mathcal{S}(6) = (\{1, 2, 4, 5\}, \{2, 3, 5, 6\})$.⁵ Thus, as the number of users N grows the relative sizes of the groups and their intersections remain fixed.

Let $s_i := |S_i|/n$ denote the relative size of S_i and $s = \sum_{i=1}^m s_i$ be the average number of groups users are in. We assume wlog. that each user is in at least one group, implying $s \geq 1$. Let $\text{Opt}(\mathcal{S})$ denote the update cost of the best key-graph for a set system \mathcal{S} and $\text{Triv}(\mathcal{S})$ the update cost of the Trivial algorithm (which makes a key-tree for every $S_i \in \mathcal{S}$). We will show that (the hidden constants in the big-Oh notation all depend on k , the number of groups).

$$\text{Opt}(\mathcal{S}(N)) = N \log(N) + \Theta(N) \quad (1)$$

$$\text{Triv}(\mathcal{S}(N)) = s \cdot N \log(N) - \Theta(N) \quad (2)$$

$$\text{thus } \frac{\text{Triv}(\mathcal{S}(N))}{\text{Opt}(\mathcal{S}(N))} = s - o(1) \quad (3)$$

As s is the average number of groups users are in, this shows that

asymptotically (for a fixed set system \mathcal{S} but with increasing number N of users) the update cost of an optimal key-derivation graph depends only on N (but not on \mathcal{S}). In this regime, the gain we get by using more cleverly chosen key-derivation graphs (as opposed to using a key-tree for every group) can be up to linear in s , the number of groups an average user is in, but not, say, the number of groups $|\mathcal{S}|$.

While we do not know how to efficiently find the best key graph for a given set system \mathcal{S} , in Section 4 we define a family $\mathcal{G}_{\text{ao}}(\mathcal{S}(N))$ which is asymptotically optimal, i.e., matches Equation 1. Intuitively, it first partitions the universe of users $[N]$ into the sets of users that are members of exactly the same groups. More precisely, for $I \subseteq [k]$ let P_I be the set of users that are members of the groups specified by I . Then, the asymptotically optimal algorithm builds a balanced binary tree for every P_I , and in a second step connects the roots of these trees to the appropriate group keys by another layer of binary trees. For an illustration of the trivial and asymptotically optimal algorithms see Figure 1.

1.2 The non-asymptotic setting

Asymptotics can kick in slowly. The asymptotic setting gives a good idea about the efficiency we can expect once the number of users N is large compared to the number $k = |\mathcal{S}|$ of groups. Nevertheless, it should be noted that this asymptotic effect can kick in only slowly: assume the artificial example where for some small

⁵ $\mathcal{S}(N)$ is only well defined if N/n is an integer, we ignore this technicality as we will be interested in the case $N \rightarrow \infty$.

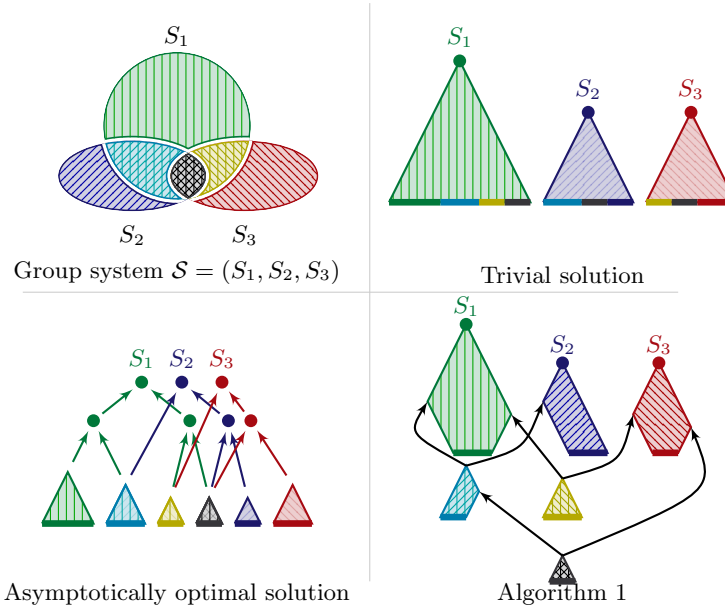


Fig. 1. Key graphs for group systems. Top left; Venn diagram of the considered group system. Top right; trivial key graph using one balanced binary tree per group. Bottom left; Asymptotically optimal key graph using one balanced binary tree per partition P_I . Bottom right; asymptotically optimal key graph obtained using Algorithm 1. In the depictions of key trees the horizontal thick lines indicates the users' personal keys.

base set $[n]$ we have a set system $\mathcal{S} = \{S_1, \dots, S_k\}$ with $k = 2^n - 1$ groups where for every non-empty subset of users we have a group. Then each user is in 2^{n-1} groups and thus needs at least that many keys, and so the $\Theta(1)$ term in the asymptotic update cost $\log(N) + \Theta(1)$ of a single user is also at least 2^{n-1} . For the $\log(N)$ term to dominate we need $\log(N) \gg 2^{n-1}$, or $N \gg 2^{2^{n-1}}$, so the number of users needs to grow doubly exponential in the base set $[n]$.

Moving on to the non-asymptotic setting, consider a group system \mathcal{S} for a fixed set of users $[N]$. The discussion above indicates that for \mathcal{S} the asymptotic update cost per user of $\log(N)$ could be very far off the truth unless N becomes fairly large compared to the number of groups. This leaves the possibility that for concrete group systems where N is not huge relative to \mathcal{S} , already the trivial key-graph performs fairly well in practice. This, however, turns out to not be the case.

First, let us observe that the gap in update cost can never be larger than $\log(N)$, for any \mathcal{S} over $[N]$

$$\text{Triv}(\mathcal{S}) \leq \log(N) \cdot \text{Opt}(\mathcal{S}) \tag{4}$$

To see this we observe that the update cost for every user $i \in [N]$ is at most a factor $\log(N)$ larger in the trivial solution: a user i that is in $s_i = |\{S \in \mathcal{S} : i \in S\}|$ groups has an update cost of at least s_i in any key graph, in particular in $\text{Opt}(\mathcal{S})$, and at most $\sum_{S \in \mathcal{S}, i \in S} \log(|S|) \leq s_i \cdot \log(N)$ in the trivial key graph.

In Section 4.2 we will show that this is not merely a theoretical gap by giving an example of a natural system \mathcal{S} for which the update costs of both the trivial and the asymptotically optimal algorithms match the gap of $\log(N)$.

A greedy algorithm based on Huffman codes. The discussion above indicates that for set systems mapping groups that we might encounter in practice, one should not simply use an asymptotically optimal solution, but aim for a solution that is optimal, or at least close to optimal, for all instances.

Algorithm 1 that we propose in Section 5 is an algorithm for computing a key-graph given a set system \mathcal{S} . In a first step, the algorithm computes a “Boolean-lattice graph” for \mathcal{S} , and in a second iteratively runs the algorithm to compute Huffman Codes to compute the key graph. As the algorithm is basically a composition of greedy algorithms, it is very efficient. We leave it as an open question whether it really is optimal, and if not, whether there’s an efficient (polynomial time) algorithm to compute $\text{Opt}(\mathcal{S})$ and find the corresponding key graph for a given \mathcal{S} in general.⁶

We present Algorithm 1 in Section 5 and discuss its connection to Boolean lattices. Then, we derive concrete lower and upper bounds on its update cost, that can serve as a good estimate on how much it saves compared to the trivial algorithm and the asymptotically optimal algorithm of Section 1.1. We further show that Algorithm 1 and a class of algorithms generalizing the approach taken are optimal in the asymptotic setting. While the same is true for the algorithm discussed in Section 1.1, Algorithm 1 seems better suited for practical applications as key-derivation graphs constructed by it reflect the hierarchical structure inherent to such systems. An example of a key graph generated by it is in Figure 1.

Our analysis concerns static group systems, but in the full version of this work [3] we show how known techniques that allow adding and removing users from groups in the settings of continuous group-key agreement and multicast encryption for a single group, can be adapted to key-derivation graphs generated by the greedy algorithm.

Lower bounds. To get a feeling for how close to optimal our approach is, we prove a lower bound on the average update cost for arbitrary schemes for continuous group-key agreement (in Section 6) and multicast encryption (in the full version of this work [3]) that are based only on simple primitives such as encryption, pseudorandom generators, and secret sharing in a *symbolic security model*. This closely follows ideas from Micciancio and Panjwani [14], who considered such a symbolic model to analyze the worst-case update cost of multicast encryption schemes. We improve on their results by considering the setting of *multiple*

⁶ The question whether a polynomial time algorithm for computing $\text{Opt}(\mathcal{S})$ exists can be naturally asked in various ways. We discuss it in more detail in Section 7.

potentially overlapping groups and proving a lower bound on the *average* communication complexity.

Our bound essentially shows that on average the cost of a user in any CGKA scheme or multicast encryption scheme for group system S_1, \dots, S_k constructed from the considered primitives satisfies

$$\text{Upd}(\mathcal{G}) \geq \frac{1}{N} \cdot \sum_{\emptyset \neq I \subseteq [k]} |P_I| \cdot \log(|P_I|) ,$$

where $P_I \subseteq [N]$ is the set of users *exactly* in the groups specified by index set $I \subseteq [k]$. We consider it an interesting open question to either improve on this bound or to construct an algorithm matching it.

1.3 Related Work

In the setting of a single group, key graphs have been used to construct secure multicast encryption, e.g. [15, 16, 7], and continuous group-key agreement (CGKA), e.g. [4, 8]. In the setting of multiple groups, the approach to use binary trees for every set of users that are members of exactly the same groups similarly to the asymptotically optimal algorithm, has been suggested in [13, 17]. However, the trees are then combined in a way that induces an overhead that is linear in the number of trees.

In [9], Cremers et al. consider the post-compromise security guarantees of CGKA protocols for multiple groups. They show that in certain update scenarios, protocols based on pairwise channels have better healing properties than protocols based on key trees, as updates in a single group also benefit all subgroups of it. We stress that these issues do not arise in our approach, as updates in our setting are global and thus affect all groups the updating user is a member of.

The symbolic security model was first introduced by Dolev and Yao [10] and used by Micciancio and Panjwani [14] to prove worst case bounds on the update cost of multicast encryption schemes for a single group. In the context of CGKA schemes it was recently used by Bienstock et al. [6], who analyze the communication cost of concurrent updates in CGKA schemes for a single group.

2 Preliminaries

2.1 Notation

Throughout the paper \log denotes the logarithm with respect to base 2.

Graph notation. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a directed acyclic graph (DAG). To node $v \in \mathcal{V}$ we associate the sets $\mathcal{A}(v) = \{v' \in \mathcal{V} \mid \exists \text{ path from } v' \text{ to } v\}$ of *ancestors* of v , and $\mathcal{D}(v) = \{v' \in \mathcal{V} \mid \exists \text{ path from } v \text{ to } v'\}$ of *descendants* of v . Here, we allow paths of length 0 and hence $v \in \mathcal{A}(v)$ and $v \in \mathcal{D}(v)$. Let $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$ be a subgraph of \mathcal{G} and $v \in \mathcal{V}'$. We denote the set of parents of v by $\mathcal{P}(v)$. The set of *co-parents* $\mathcal{CP}(v, \mathcal{G}') \subseteq \mathcal{V}$ of v with respect to \mathcal{G}' in \mathcal{G} is the set of vertices that are parents of v in \mathcal{G} but not in \mathcal{G}' .

Probability distributions. Let X be a random variable with outcomes x_1, \dots, x_ℓ with probability p_1, \dots, p_ℓ . Then we denote by $\mathbb{E}[X]$ its expectation and by $H(X) = -\sum_{i=1}^{\ell} p_i \log(p_i)$ its Shannon entropy.

2.2 Huffman Codes

Given a collection v_1, \dots, v_ℓ of disconnected leaves of weight $w_1, \dots, w_\ell \in \mathbb{N}$ a *Huffman Tree* is constructed as follows. From the set $\{v_1, \dots, v_\ell\}$ two nodes v_{i_1}, v_{i_2} with the smallest weights are picked. Then a node v and edges $(v_{i_1}, v), (v_{i_2}, v)$ are added to the graph. v 's weight is set to $w_{i_1} + w_{i_2}$ and the set of nodes to be considered updated to $\{v_1, \dots, v_\ell\} \cup \{v\} \setminus \{v_{i_1}, v_{i_2}\}$. This step is repeated until all leaves are collected under a single root.

Since all nodes have indegree 2 the Huffman tree defines a prefix-free binary code for (v_1, \dots, v_ℓ) . We will make use of the following property of Huffman Codes.

Lemma 1 (Optimality of Huffman Codes [11]). *Consider a Huffman tree \mathcal{T} over leaves v_1, \dots, v_ℓ of weight $w_1, \dots, w_\ell \in \mathbb{N}$. Let $w = \sum_{i=1}^{\ell} w_i$ and let $U_{\mathcal{T}}$ denote the probability distribution that picks leaf v_i with probability w_i/w proportional to its weight. Then, if $\text{len}(U_{\mathcal{T}})$ denotes the random variable measuring the length of the path from a leaf picked according to $U_{\mathcal{T}}$ to the root, we have that the average length of such paths is bounded by*

$$H(U_{\mathcal{T}}) \leq \mathbb{E}[\text{len}(U_{\mathcal{T}})] \leq H(U_{\mathcal{T}}) + 1 .$$

3 Key-derivation Graphs for Multiple Groups

In this section we discuss key-derivation graphs for systems consisting of multiple groups. In Section 3.1 we briefly recall two applications of such graphs; continuous group-key agreement and multicast encryption. In Section 3.2 we formally define key-derivation graphs, discuss how key material in a graph is refreshed, and define its update cost.

3.1 Continuous Group-key Agreement and Multicast Encryption

Continuous group-key agreement. *Continuous group-key agreement (CGKA)* schemes [1] are an important building block in the construction of secure asynchronous group messaging schemes. As the name indicates, the goal of a CGKA scheme is to establish a common key that is to be used to secure the communication between members of a group. As groups can typically be long-lived, users need to also be able to frequently update the key material known to them, to on one hand, recover from a potential compromise and, on the other hand, ensure forward-secrecy of messages sent in the past.

In this work we are interested in the more general setting in which users $n \in [N]$ want to agree on keys for a system of groups $S_1, \dots, S_k \subseteq 2^{[N]}$. After

the groups have been established in a setup phase user n can use the procedure $\text{Upd}(n)$ to produce an update message that rotates the key material known to them, thus eliminating any keys that may have leaked during a compromise. This update message is broadcast to the other users using the untrusted delivery server. Given their own secret keys, users are then able to retrieve the refreshed keys that should be known to them. A natural goal to aim for is to minimize the communication cost incurred by such update messages.

Naturally, one would like to additionally support dynamic operations, i.e., allow users to add and remove other users from groups in the system. While in this work we focus on the update costs of schemes for a system of static groups, in the full version of this work [3] we show that the known techniques of blanking and unmerged leaves used in the MLS protocol [4] can be adapted to schemes obtained from our approach.

Efficient CGKA protocols [4, 8] (in the single group setting) establish a key-derivation graph in the setup phase that, in turn, allows user to update at a cost that is logarithmic in the number of group members.⁷ In Section 3.2 we formally define key-derivation graphs and discuss how the updating process works.

Multicast encryption. The goal of a *multicast encryption* scheme [15, 16, 7] is to establish a key for a group of users to enable them to decrypt ciphertexts broadcast to the group. Every user holds a personal long-term key, but opposed to CGKA there also exists a central authority that has access to all secret key material. After a setup phase, the central authority is able to add and remove users from the group by refreshing key material and broadcasting messages to the group. The central goal in the construction of multicast schemes is to minimize the communication complexity incurred by such operations. Typically, multicast encryption schemes also rely on key-derivation graphs.

As in the case of CGKA, we are interested in the more general setting of a system of potentially overlapping groups of users.

3.2 Key-derivation Graphs

We now discuss key-derivation graphs. In our exposition we will focus on graphs for continuous group-key agreement. At the end of the section we discuss the differences to graphs for multicast encryption.

Consider a set of parties $[N]$ and a collection $\mathcal{S} \subseteq 2^{[N]}$ of subgroups of $[N]$. A key-derivation graph (kdg) for $[N]$ and \mathcal{S} organizes key pairs in a way that allows the members of a particular subgroup to agree on a key, and further enables parties to refresh the key material known to them. Every node v in the graph is associated to a key pair $(\mathbf{pk}_v, \mathbf{sk}_v)$ of a public-key encryption scheme $(\text{KGen}, \text{Enc}, \text{Dec})$, and edges (v, v') indicate that parties with access to

⁷ In order to ensure authenticity of update messages and to prevent the server from sending users inconsistent update messages these protocols employ additional techniques. We leave the question how to adapt these to key-derivation graphs for multiple groups to future work (See Section 7).

sk_v also possesses $\text{sk}_{v'}$. The personal keys of users correspond to sources and every group is represented by a node that holds the corresponding secret group key. We formalize the structural requirements on the graph in the multi-group setting as follows.

Definition 1. Let $N \in \mathbb{N}$, $\mathcal{S} \subseteq 2^{[N]}$, and $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ a DAG. We say that \mathcal{G} is a key-derivation graph for universe of elements $[N]$ and groups \mathcal{S} if

1. For every $n \in [N]$ there exists a source $v_n \in \mathcal{V}$ and for every $S \in \mathcal{S}$ there exists a node $v_S \in \mathcal{V}$. We further require that $v_n \neq v'_n$ for $n \neq n'$.
2. For $n \in [N]$ and $S \in \mathcal{S}$ we have $v_S \in \mathcal{D}(v_n)$ exactly if $n \in S$.

In the definition above node v_n correspond to user n 's personal key, and nodes v_S to group keys. The second property encodes correctness and security, intuitively saying that n is able to derive the group key of S exactly if $n \in S$.

Updates. Let \mathcal{G} be a key-derivation graph for $[N]$ and \mathcal{S} . If party n wants to perform an update she has to refresh all key-material corresponding the sub-graph $\mathcal{D}(v_n)$ known to her and communicate the change to the other parties. To this end she picks a spanning tree $\mathcal{T}_n = (\mathcal{V}', \mathcal{E}')$ of $\mathcal{D}(v_n)$, as well as a random seed Δ_{v_n} . Then starting from the source v_n , if v' is the i th child of node v she defines the seed of v' as $\Delta_{v'} = \text{H}(\Delta_v, i)$, where H is a hash function. $\Delta_{v'}$ is then used to derive a new key-pair $(\text{pk}_{v'}, \text{sk}_{v'}) \leftarrow \text{KGen}(\Delta_{v'})$ for v' . Finally, for every $v \in \mathcal{V}'$ and every co-parent $v' \in \mathcal{CP}(v, \mathcal{T}_n)$, n computes the ciphertext $c_{v,v'} = \text{Enc}(\text{pk}_{v'}, \Delta_v)$. The set of all ciphertexts together with the set of new public keys forms the update message. Finally, n deletes all seeds Δ_v .

We now show that the construction preserves correctness, i.e., users $n' \neq n$ are able to deduce all new secret keys in $\mathcal{D}(v_{n'})$ from the update message and thus in particular the group keys of all groups they are a member of. To this end, let $v \in \mathcal{D}(v_n) \cap \mathcal{D}(v_{n'})$. Then there exists a path $(v_{n'} = v_1, \dots, v_\ell = v)$ in $\mathcal{D}(v_{n'})$. Let i be maximal with $v_i \notin \mathcal{D}(v_n)$ (Note that such i must exist as $v_{n'}$ is a source). By maximality of i the node v_i must be a coparent of v_{i+1} with respect to $\mathcal{D}(v_n)$. Thus, the update message contains an encryption of $\Delta_{v_{i+1}}$ to pk_{v_i} . As sk_{v_i} was not replaced by the update and is known to n' the user can recover $\Delta_{v_{i+1}}$ and in turn $\text{sk}_{v_{i+1}}$. Now, n' can recover the remaining $\Delta_{v_{i+2}}, \dots, \Delta_{v_\ell}$ and the corresponding secret keys as the seeds were either derived by hashing or, in the case that v_{j+1} is a coparent of v_j with respect to $\mathcal{D}(v_n)$, encrypted to the new key pk_{v_j} , the secret key of which was already recovered by n' .

Update cost. Using the size of ciphertexts as a unit, the update cost of n is given by $\text{Upd}(n) = \sum_{v \in \mathcal{T}_n} |\mathcal{CP}(v, \mathcal{T}_n)| = \sum_{v \in \mathcal{T}_n} (|\mathcal{P}(v)| - 1)$. Note that this quantity is independent of the particular choice of spanning tree \mathcal{T}_n . In this work we are interested in minimizing the *average* update cost, assuming that every user updates with the same probability. We define the *total update cost* $\text{Upd}(\mathcal{G}) = \sum_{n \in [N]} \text{Upd}(n)$ of \mathcal{G} . Note that $\text{Upd}(\mathcal{G})/N$ is the average update cost of a user, and we can thus focus on trying to minimize $\text{Upd}(\mathcal{G})$, which will allow for easier exposition. The following lemma shows that we can restrict our view to graphs

in which every non-source has indegree 2. Note, that for graphs \mathcal{G} with this property we have $|\mathcal{CP}(v, \mathcal{T}_n)| = 1$ for every n, \mathcal{T}_n , and $v \in \mathcal{T}_n$ that is not a source and thus in this case we can compute the update cost as

$$\text{Upd}(\mathcal{G}) = \sum_{n \in [N]} (|\mathcal{T}_n| - 1) = \sum_{n \in [N]} (|\mathcal{D}(n)| - 1) = \sum_{n \in [N]} |\mathcal{D}(n)| - N . \quad (5)$$

Lemma 2. *Let $n \in \mathbb{N}$, $\mathcal{S} \subseteq 2^{[N]}$, and \mathcal{G} a key-derivation graph for $[N]$ and \mathcal{S} . Then there exists a key-derivation graph \mathcal{G}' for $[N]$ and \mathcal{S} satisfying $\text{Upd}(\mathcal{G}') \leq \text{Upd}(\mathcal{G})$ such that for every non-source $v \in \mathcal{V}'$ we have $\text{indeg}(v) = 2$.*

Due to space constraints we defer the proof to the full version of this work [3].

Key-derivation graphs for multicast encryption. Opposed to kdgs for CGKA key-derivation graphs for multicast encryption rely on symmetric encryption. Let (\mathbf{E}, \mathbf{D}) be a symmetric encryption scheme. Every node v in a kdg \mathcal{G} for $[N]$ and \mathcal{S} is associated to a key k_v , and an edge (v, v') indicates that a party with access to k_v knows $k_{v'}$. We require structural requirements on \mathcal{G} that are analogous to Definition 1. Updates with respect to leaf v_n , which for multicast encryption are computed by the central authority, and their update cost, are defined analogous to the setting of CGKA as well.

While the main goal of multicast encryption is not to recover from compromise of keys by updating, but instead to allow the central authority to dynamically change the structure of the groups S_1, \dots, S_k , the notion of an update with respect to a leaf v_n still turns out to be useful. Assume that the central authority performed an update for v_n starting with seed Δ . We can distinguish two cases. If Δ is not known to the owner n of leaf v_n then n lost access to all keys corresponding to $\mathcal{D}(v_n)$. Thus, by updating, the central authority can remove a party from *all* groups they are a member of. Assume on the other hand that the leaf was previously unpopulated and that Δ can be derived from n 's long term key. Then n gained access to all group keys that can be reached from v_n . In the full version of this work [3] we discuss how updates can be used as the basic building block of implementing more fine grained operations, i.e., adding or removing a user from particular group S_i . The efficiency of these operations is significantly determined by the update cost as defined in this section.

3.3 Security

The main focus of this work is to investigate the communication complexity of key-derivation graphs for group systems. We do not give formal security proofs in this work. The structural requirements on kdgs and definition of update procedures are chosen with the goal of the resulting CGKA to achieve *post-compromise forward-secrecy* (PCFS) [2] roughly corresponding to *post-compromise security* (PCS) and *forward-secrecy* (FS) simultaneously. In the following paragraphs we provide an intuition on the security properties of kdgs. For ease of exposition we will discuss PCS and FS separately instead of PCFS.

Note that CGKA schemes constructed from kdgs employ further mechanisms to ensure authenticity and prevent a malicious server to send users inconsistent update messages. We consider the construction of such mechanisms as well as a formal security analysis of kdgs to be important open questions for future work.

Preserving the graph invariant. We first discuss how updates preserve the invariant, that users n know exactly the secret keys corresponding to $\mathcal{D}(v_n)$, which by Condition 2 of Definition 1 implies that n will never be able to derive a group key for some group they are not a member of. Note that if n is the updating user then they will only replace keys in $\mathcal{D}(v_n)$. If n receives an update message, on the other hand, then they will only be able to recover a key sk_v if either the corresponding seed Δ_v was encrypted to a key known to n or if Δ_v was derived by hashing from a seed $\Delta_{v'}$ recoverable by n . By iteratively applying this argument to $\Delta_{v'}$ we obtain that there must exist some $\Delta_{v''}$ that was encrypted to a key known to n such that v'' has a path to v . Thus, it must hold that $v \in \mathcal{D}(v_n)$. (Note that the one-wayness of the used hash function ensures that seeds derived by hashing can only be recovered from each other in the correct direction.)

Post-compromise security. The goal of PCS is to allow users whose secret state has been exposed to recover from this exposure by performing an update. Using the example of a single compromised user we now discuss how kdgs for group systems achieve this property. Assume that an adversary knows exactly the secret state of user n , i.e., all keys sk_v for $v \in \mathcal{D}(v_n)$, and that n then performs an update. Then the adversary is not able to deduce any of the replaced keys: Note that the initial random seed Δ_{v_n} is not encrypted to any key and thus cannot be leaked to the adversary. Thus, all other seeds Δ_v can only be derived by the adversary if Δ_v itself, or a seed from which Δ_v was derived by iterated hashing was encrypted to a key known to the adversary. However, the adversary only knows the keys corresponding to $\mathcal{D}(v_n)$ before the update, and those keys were replaced by freshly sampled ones before computing the ciphertexts. Thus, seeds are encrypted to either “old” keys not known to the adversary or new keys, and so after the update all keys are secure again.

Forward secrecy. Forward secrecy requires that compromising a user’s secret state does not allow the adversary to recover previous group keys. In key-derivation graphs old keys get deleted over time providing a limited form of forward-secrecy. Concretely, if a user n is corrupted all group keys before their last update remain secure. This holds, since seeds that were generated before this point in time and can be used to recover group keys were encrypted to keys no longer in n ’s memory. Note however, that group keys generated in between n ’s last update and the time of n ’s corruption might leak to the adversary. For example, a seed from which such keys can be derived might have been encrypted to the key sk_{v_n} , which remained unchanged until the corruption.

Improved forward secrecy using supergroups. CGKA constructions relying on kdgs like TreeKEM [5] rely on an additional mechanism to improve their forward-secrecy guarantees. Instead of directly using group keys sk_{v_S} to communicate

within the group these keys are used to derive a so called *application secret* K that serves as the symmetric key for group communication. Whenever an update occurs, the new application secret of S is computed as $H_2(\text{sk}_{v_S}, K)$ the output of a hash function on input of the new group key and the previous application secret. Then, the old application secret is deleted from memory. The effect of this is that when a user’s state leaks (including the current application secret K_t), no old application secret K_i can be recomputed from old update messages, unless K_{i-1} was already known to the adversary by former corruptions. In short, users gain the advantage of forward secrecy not only by issuing but also by processing updates of other users in S .

In the setting of a group system \mathcal{S} we can further improve on this: Consider some group $S \in \mathcal{S}$ and let S_1, \dots, S_ℓ be the maximal (with respect to inclusion) groups in \mathcal{S} that contain S . We denote the application secrets for S and the S_i by K_S and K_{S_i} respectively. Now, whenever a member of any of the S_i issues an update the application secret of S is updated to $K_S \leftarrow H_2(\text{sk}_{v_S}, K_{S_1}, \dots, K_{S_\ell})$.⁸ Note that for every i since $S \subseteq S_i$ all members of S do indeed have access to K_{S_i} and thus are able to compute K_S , and that an update by users in S implies that all S_i are updated as well. The effect of this modification is that even updates by users *outside* of S —more precisely in any of the sets $S_i \setminus S$ — imply forward secrecy of users in S . Note that this is in particular helpful in the case where $|S| \ll |S_i|$ and updates in the large group occur much more frequently than in the small group, for example in the case of two members of a large group having a private conversation.

3.4 The Trivial Algorithm

To construct a key-derivation graph for a single group S the parties $n \in S$ are typically arranged as the leaves of a balanced binary tree \mathcal{T} . The tree’s root serves as the group key. In this case the length of paths from leaf to root is at most $\lceil \log(|S|) \rceil$ and in turn $\text{Upd}(\mathcal{T}) \leq |S| \cdot \lceil \log(|S|) \rceil$. On the other hand, \mathcal{T} defines a prefix-free binary code for the set S . Thus, by Shannon’s source coding theorem the average length of paths from leaf to root is at least $\log(|S|)$ which implies $\text{Upd}(\mathcal{T}) \geq |S| \cdot \log(|S|)$.

An algorithm for multiple groups. A trivial approach to construct a key derivation graph for parties $[N]$ and group system $\mathcal{S} = \{S_1, \dots, S_k\}$ is to simply apply the method described above to all S_i in parallel. That is, for $i \in [k]$ construct a balanced binary tree \mathcal{T}_i with $|S_i|$ leaves such that for $n \in [N]$ the node v_n is a leaf of exactly the trees \mathcal{T}_i with $n \in S_i$. Let \mathcal{G} denote the resulting graph. The conditions of Definition 1 clearly hold and we can bound the total update cost of \mathcal{G} by

$$\sum_{i \in [k]} |S_i| \cdot \log(|S_i|) \leq \text{Upd}(\mathcal{G}) \leq \sum_{i \in [k]} |S_i| \cdot \lceil \log(|S_i|) \rceil .$$

⁸ Regarding PCFS it might even be advantageous to include $K_{S'}$ for all $S' \supseteq S$.

Further, the update cost of a single user $n \in [N]$ is bounded by $\text{Upd}(n) \leq \sum_{i:n \in S_i} \lceil \log(|S_i|) \rceil$.

4 Key-derivation Graphs in the Asymptotic Setting

In this section we investigate the update cost of key-derivation graphs for multiple groups in an asymptotic setting. More precisely, for a system consisting of a fixed number of groups, we consider the setting in which the number of users tends to infinity while the relative size of the groups stays constant. In Section 4.1 we first compute the asymptotically optimal update cost of key-derivation graphs and then show that the trivial algorithm does not achieve it. We then present an algorithm achieving the optimal update cost. In Section 4.2 we show that both approaches can perform badly for *concrete* group systems.

4.1 Key-derivation Graphs in the Asymptotic Setting

We investigate the update cost of key derivation graphs in an asymptotic setting. That is, we consider N parties that form a subgroup system $\mathcal{S} = \{S_1, \dots, S_k\}$ and fix values $p_I \in [0, 1]$ for $I \subseteq [k]$ that indicate the fraction of users that are members of exactly the groups specified by I .

More precisely, let $k \in \mathbb{N}_{\geq 2}$ be fixed and let $\{p_I\}_{I \subseteq [k]}$ be such that $\sum_{I \subseteq [k]} p_I = 1$. For $N \in \mathbb{N}$ let $\mathcal{S}(N) = \{S_1(N), \dots, S_k(N)\} \subseteq 2^{[N]}$ be a subgroup system that satisfies $|P_I(N)| = N \cdot p_I$ for all I , where $P_I(N) = \bigcap_{i \in I} S_i(N) \setminus \bigcup_{j \in [k] \setminus I} S_j(N)$ is the set of users exactly in the groups specified by I .⁹ Throughout this section we assume that $p_\emptyset = 0$, i.e., every user is in at least one group, and that at least two groups are non-empty. We are interested in the update cost of key-derivation graphs for $\mathcal{S}(N)$ when N tends to infinity.

Lower bound in the asymptotic setting. We first compute a lower bound on the update cost of kdgs in the asymptotic setting. The bound follows from the following combinatorial result on *concrete* graphs, that will also turn out to be useful for our symbolic lower bound of Section 6. Recall that for graphs $\mathcal{G}' \subseteq \mathcal{G}$ and a vertex v the set $\mathcal{CP}(v, \mathcal{G}')$ is the set of co-parents of v with respect to \mathcal{G}' in \mathcal{G} . Due to space constraints, we defer its proof to the full version of this work [3].

Lemma 3. *Let $M \in \mathbb{N}$ be fixed, $S_1, \dots, S_k \subseteq [M]$, and let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a DAG such that there exist pairwise disjoint sets of sources V_n , $n \in [M]$, and nodes v_{S_i} , $i \in \{1, \dots, k\}$ such that*

$$n \in S_i \quad \Rightarrow \quad \exists v_n \in V_n \text{ such that there is a path from } v_n \text{ to } v_{S_i} .$$

Further let \mathcal{T}_n be a spanning forest of $\mathcal{D}(V_n) = \bigcup_{v_n \in V_n} \mathcal{D}(v_n)$. Then

$$M \cdot \mathbb{E} \left[\sum_{v \in \mathcal{T}_n} |\mathcal{CP}(v, \mathcal{T}_n)| \right] \geq \sum_{\emptyset \neq I \subseteq [k]} |P_I| \cdot \log(|P_I|) ,$$

⁹ $\mathcal{S}(N)$ is only well defined if $N \cdot p_I$ is an integer for all I , we ignore this technicality as we are interested in the case $N \rightarrow \infty$.

where the expectation is to be understood with respect to the uniform distribution on $[N]$.

Note that Lemma 3 in the case $|V_n| = 1$ for all n can be seen as a lower bound on the total update cost of key-derivation graphs as defined in Section 3 since $M \cdot \mathbb{E}[\sum_{v \in \mathcal{T}_n} |\mathcal{CP}(v, \mathcal{T}_n)|] = \sum_{v \in \mathcal{T}_n} |\mathcal{CP}(v, \mathcal{T}_n)|$.

Turning to the asymptotic setting we have

$$\begin{aligned} \sum_{I \subseteq [k]} N \cdot p_I \cdot \log(N \cdot p_I) &= N \cdot \sum_{I \subseteq [k]} p_I \log(N) + N \cdot \sum_{I \subseteq [k]} p_I \log(p_I) \\ &= N \log(N) + N \cdot \sum_{I \subseteq [k]} \log(p_I) = N \log(N) + \Theta(N) , \end{aligned}$$

where we used that $\sum_I p_I = 1$. As we will show below, there exist key-derivation graphs matching this bound. We conclude that the optimal update cost in the asymptotic setting only depends on the overall number of users but not the particular set system:

$$\text{Opt}(\mathcal{S}(N)) = N \log(N) + \Theta(N) .$$

Note, however, that the term $\Theta(N)$ hides a constant (with respect to N), that can be *exponential in k* .

Asymptotic update cost of the trivial algorithm. The trivial algorithm constructs a separate balanced binary tree for every group $S_i(N)$. For $i \in [k]$ let s_i be such that $N \cdot s_i = |S_i(N)|$ and further let $s = \sum_{i=1}^k s_i$ be the average number of groups a user are member of. Then, we can bound the update cost $\text{Triv}(\mathcal{S}(N))$ of the trivial algorithm in the asymptotic setting as follows, showing that it does not match the optimal cost.

Claim. For $I \subseteq [k]$ let $p_I \in [0, 1]$ be such that $\sum_{I \subseteq [k]} p_I = 1$ and $p_\emptyset = 0$. Let $\mathcal{S}(N)$ be the corresponding group system and s_i, s as defined above. Then

$$\text{Triv}(\mathcal{S}(N)) = s \cdot N \log(N) + \Theta(N) .$$

Due to space constraints we defer the proof of this claim to the full version of this work [3].

An asymptotically optimal graph. We will sketch how to construct an asymptotically optimal key graph $\mathcal{G}_{\text{ao}}(N)$ for a given set system \mathcal{S} over $[n]$. In a first step, for every I with $P_I(N) \neq \emptyset$, the algorithm constructs a balanced binary tree with root v_I using as leafs the elements of $P_I(N)$. Then, in a second step, for every group $S_i(N)$ it builds a balanced binary tree with root v_{S_i} using as leafs the nodes $\{v_I \mid I : i \in I\}$. An illustration of the algorithm's working principle is in Figure 1. Correctness of the construction follows by inspection.

To bound the update cost $\text{Upd}(\mathcal{G}_{\text{oa}}(N))$ we split it in two parts; the first accounts for the contribution of the nodes generated during the first step, the




Fig. 2. Illustration of $\text{Triv}(\mathcal{S}_N^\dagger)$ (left) and $\text{Opt}(\mathcal{S}_N^\dagger)$ for $N = 5$. For each user, the update cost (i.e., the indegree 2 nodes reachable) is indicated.

second for the contribution of the second step. As $\sum_I p_I = 1$, the first part contributes at most $\sum_{I \subseteq [k]} p_I \cdot N \cdot \log(N \cdot p_I) \leq N \cdot \log N$, while the contribution of the second part for every single user is constant as $\{v_I\}$ is independent of N , implying that with respect to the total update cost it is $\Theta(N)$. Thus, overall we get $\text{Upd}(\mathcal{G}_{\text{oa}}(N)) \leq N \cdot \log N + \Theta(N)$, matching the optimal update cost.

4.2 Update Cost for Concrete Group Systems

Now consider a concrete group system $\mathcal{S} = \{S_1, \dots, S_k\}$ for a fixed set of users $[N]$. As already discussed in Section 1.2, it is possible that the number k of groups can be as large as $2^N - 1$. Thus, for concrete group systems the asymptotic update cost per user of $\log(N)$ (that contains hidden constants dependent on k) derived in Section 4.1 could be very far off the truth unless N becomes fairly large compared to the number of groups. This leaves the possibility that in the case where N is not huge relative to k , already the trivial key-graph performs fairly well in practice. In this section we show that this is not the case by giving an example where not only the trivial key-graph (which has a balanced tree for every set), but also our asymptotically optimal \mathcal{G}_{oa} , perform poorly.

Recall that by Equation 4 the update costs of the trivial and optimal solutions always satisfy $\text{Triv}(\mathcal{S}) \leq \log(N) \cdot \text{Opt}(\mathcal{S})$. The above argument seems very loose, but we show an example where we indeed have a gap of $\approx \log(N) - 1$ and thus almost match this seemingly loose $\log(N)$ bound. Define the “hierarchical” set system \mathcal{S}_N^\dagger over $[N]$ as

$$\mathcal{S}_N^\dagger := \{S_1, \dots, S_N\} \quad \text{where} \quad S_i = \{i, i+1, \dots, N\} .$$

Note that while \mathcal{S}_N^\dagger is defined for all N , it is not asymptotic in the sense discussed in Section 4.1, as the number of groups grows with the number of users N . Further, for this group system the key derivation graphs output by the trivial and asymptotically optimal algorithms coincide, as for every P_I with $P_I \neq \emptyset$ we have $|P_I| = 1$. As the optimal solution for \mathcal{S} is just a path, as illustrated in

Figure 2, we obtain update costs of $\text{Triv}(\mathcal{S}_N^\uparrow) = \sum_{i=1}^N i \log(i) \approx \frac{N^2}{2} \log(N)$ and $\text{Opt}(\mathcal{S}_N^\uparrow) = \sum_{i=1}^N i = \frac{N(N+1)}{2} \approx \frac{N^2}{2}$.

Thus $\text{Triv}(\mathcal{S}_N^\uparrow)/\text{Opt}(\mathcal{S}_N^\uparrow) \approx \log(N)$ matching the (4) bound. An interesting observation is the fact that an optimal solution can have much larger *depth* than the trivial one: for \mathcal{S}_N^\uparrow the depth of the optimal solution is N , while in the trivial solution it is just $\log(N)$. The discussion above indicates that neither the trivial nor the asymptotically optimal algorithm are well-equipped to handle certain group systems. In the following section we propose an algorithm that is not only asymptotically optimal, but also generates key-derivation graphs better reflecting the hierarchical nature of group systems, and, in particular, recovers the optimal solution for the example above.

5 A Greedy Algorithm Based on Huffman Codes

In this section we propose an algorithm to compute key-derivation graphs for group systems. Its formal description is in Section 5.1. In Section 5.2 we compute bounds on its total update cost and compare it to the trivial algorithm and the asymptotically optimal algorithm of Section 4.1. Finally, in Section 5.3 we show that the algorithm, as well as a class generalizing it, are asymptotically optimal.

5.1 Algorithm Description

We now describe Algorithm 1 that, on input parties $[N]$ and a set of groups $\mathcal{S} \subseteq 2^{[N]}$, constructs a key-derivation graph. Its formal description is in Figure 3.

Conceptually, the algorithm proceeds in two phases. The first phase (lines 1 to 11) determines the macro structure of the key-derivation graph. For reasons explained below we will refer to the graph generated in this phase as the *lattice graph*. In the second phase (lines 12 to 20), sources for the individual users are added at the correct position in the lattice graph, which afterwards is binarized to reduce the update size.

More precisely, at the beginning of the first phase the algorithm initializes a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ consisting of isolated nodes $v_{S'}$ with $S' \in \mathcal{S}$ that, looking ahead, will represent the group keys. Every node $v_{S'}$ is associated to a set $\mathcal{S}(v_{S'})$ that is initialized to group S' . The algorithm then determines nodes v_1, v_2 such that the intersection of their associated sets is maximal and adds a node v_3 as well as the edges $(v_3, v_1), (v_3, v_2)$ to the graph. The associated set of v_3 is set to $\mathcal{S}(v_1) \cap \mathcal{S}(v_2)$ and the associated sets of v_1 and v_2 are updated to $\mathcal{S}(v_1) \setminus \mathcal{S}(v_3)$ and $\mathcal{S}(v_2) \setminus \mathcal{S}(v_3)$ respectively. This step is repeated until the associated sets of all nodes are pairwise disjoint.

Let $\mathcal{G}_{\text{lat}} = (\mathcal{V}_{\text{lat}}, \mathcal{E}_{\text{lat}})$ denote the resulting lattice graph. In the second phase, for every node $v \in \mathcal{V}_{\text{lat}}$ for all $n \in \mathcal{S}(v)$, a source v_n representing user n together with edge (v_n, v) is added to the graph. Finally, for every node v with $\text{indeg}(v) \geq 3$, a Huffman tree from the parents to the node is built. Here, the weight of a source is 1, and the weight of non-sources is given as the number of sources below it.

Input: (N, \mathcal{S})

```

1 :  $\mathcal{G} = (\mathcal{V}, \mathcal{E}) \leftarrow (\emptyset, \emptyset)$ 
2 : for  $S' \in \mathcal{S}$ 
3 :    $\mathcal{V} \leftarrow \mathcal{V} \cup \{v_{S'}\}$ 
4 :    $\mathcal{S}(v_{S'}) \leftarrow S'$ 
5 : while the sets associated to  $\mathcal{V}$  are not disjoint
6 :    $v_1, v_2 \leftarrow \arg \max_{v_1, v_2 \in \mathcal{V}} (|\mathcal{S}(v_1) \cap \mathcal{S}(v_2)|)$ 
7 :   add the node  $v_3$ 
8 :    $\mathcal{S}(v_3) \leftarrow \mathcal{S}(v_1) \cap \mathcal{S}(v_2)$ 
9 :    $\mathcal{S}(v_1) \leftarrow \mathcal{S}(v_1) \setminus \mathcal{S}(v_3)$ 
10 :   $\mathcal{S}(v_2) \leftarrow \mathcal{S}(v_2) \setminus \mathcal{S}(v_3)$ 
11 :  add the edges  $(v_3, v_1), (v_3, v_2)$ 
12 : for  $v \in \mathcal{V}$ 
13 :   for  $n \in \mathcal{S}(v)$ 
14 :     add the node  $v_n$ 
15 :     add the edge  $(v_n, v)$ 
16 :      $\mathcal{S}(v) \leftarrow \mathcal{S}(v) \setminus \{n\}$ 
17 : compute the weight of each node as the number of sources below it
18 : for every node with indegree  $> 1$ 
19 :   build a Huffman tree from the parents to the node
20 : return  $\mathcal{G}$ 

```

Fig. 3. Algorithm 1

Properties of the lattice graph. We now derive several properties of the lattice graph, which will be used to prove correctness and compute bounds on the total update cost of the generated key-derivation graph. Thus, let $\mathcal{G}_{\text{lat}} = (\mathcal{V}_{\text{lat}}, \mathcal{E}_{\text{lat}})$ be the lattice graph generated on input of $[N]$ and set of k groups $\mathcal{S} = \{S_1, \dots, S_k\} \subseteq 2^{[N]}$. For index set $I' \subseteq [k]$ we denote by

$$P_{I'} := \bigcap_{i \in I'} S_i \setminus \bigcup_{j \in [k] \setminus I'} S_j ,$$

the set of parties that are members of exactly the groups specified by I' . Further, for $v \in \mathcal{V}_{\text{lat}}$ we define

$$I(v) := \{i \in [k] \mid \text{exists path from } v \text{ to } v_{S_i}\} ,$$

the index set of group nodes that can be reached from v . Finally, for a collection $\mathcal{V}' \subseteq \mathcal{V}$ of nodes we generalize the notation for associated sets to $\mathcal{S}(\mathcal{V}') := \bigcup_{v \in \mathcal{V}'} \mathcal{S}(v)$. We obtain the following.

Lemma 4. Let $N, k \in \mathbb{N}$, $\mathcal{S} = \{S_1, \dots, S_k\} \subseteq 2^{[N]}$, and let $\mathcal{G}_{\text{lat}} = (\mathcal{V}_{\text{lat}}, \mathcal{E}_{\text{lat}})$ be the lattice graph generated on input $([N], \mathcal{S})$. Then the following holds.

1. Let $v, v' \in \mathcal{V}_{\text{lat}}$ be such that $I(v) = I(v')$. Then $v = v'$.
2. $I(v) \neq \emptyset$ for all $v \in \mathcal{V}_{\text{lat}}$.
3. For every $v \in \mathcal{V}_{\text{lat}}$ and every $i \in I(v)$ there is exactly one path from v to v_{S_i} .
4. Consider the ancestor graph $\mathcal{A}(v)$ for $v \in \mathcal{V}_{\text{lat}}$. Then

$$\bigcup_{v' \in \mathcal{A}(v)} S(v') \subseteq \bigcap_{i \in I(v)} S_i .$$

If $|I(v)| = 1$ then the equation holds with equality, i.e., $\bigcup_{v' \in \mathcal{A}(v_S)} S(v') = S$ for all $S \in \mathcal{S}$.

5. Consider some $v \in \mathcal{V}_{\text{lat}}$. Then we have $S(v) = P_{I(v)}$.

Due to space constraints, we defer the proof to the full version of this work [3]. We briefly discuss how Lemma 4 allows us to interpret the lattice graph as a subgraph of the Boolean lattice with respect to the power set of $[k]$, i.e., the graph $\mathcal{G}_B = (\mathcal{V}_B, \mathcal{E}_B)$ with $\mathcal{V}_B = \{v_I \mid I \subseteq [k]\}$ and edges $\mathcal{E}_B = \{(v_I, v_{I'}) \mid I, I' \subseteq [k] : I' \subseteq I\}$. Indeed, Properties 1 and 2 allow us to map every $v \in \mathcal{V}_{\text{lat}}$ to a unique index set $I \subseteq [k]$. Since the existence of an edge $(v, v') \in \mathcal{E}_{\text{lat}}$ implies that $I(v) \supseteq I(v')$ all edges adhere to the structure of \mathcal{G}_B . Summing up, the map $\mathcal{G} \rightarrow \mathcal{G}_B; v \mapsto v_{I(v)}$ is an injective graph homomorphism. This allows us to identify nodes of the lattice graph with nodes of \mathcal{G}_B and sometimes write $v_{I'}$ for a unique node $v \in \mathcal{V}_{\text{lat}}$ with $I(v) = I' \in \mathcal{P}([k])$. By Property 5 the associated set of v is $P_{I'}$, the set of users exactly in the groups specified by I' . Figure 4 depicts an example execution of Algorithm 1.

Correctness. We show that key-derivation graph \mathcal{G} output by Algorithm 1 satisfies the correctness properties of Definition 1. Note that the first property holds by construction.

To see that the second property holds as well, consider the lattice graph. By Lemma 4, Property 4 for every group $S' \in \mathcal{S}$ the associated sets of the ancestors of $v_{S'}$ form a partition of S' . In the second phase of the algorithm a source v_n is added for every user and connected to corresponding node in the lattice graph. Thus, after this step the set of users with a path to $v_{S'}$ is exactly S' . As this property remains unaffected by the binarization step of line 19 the final key-derivation graph is indeed correct.

5.2 Total Update Cost

In this section we analyze the total update cost $\text{Upd}(\mathcal{G}) = \sum_{n \in [N]} \text{Upd}(n)$ of key-derivation graphs \mathcal{G} generated by Algorithm 1. To this end, we will split $\text{Upd}(\mathcal{G})$ into the contribution made by the constituting Huffman trees \mathcal{T} . Tree \mathcal{T} has a single root and all non-sources in \mathcal{T} have indegree 2. Let $\mathcal{L}(\mathcal{T})$ denote the set of leaves of \mathcal{T} . As argued in Lemma 2, the update cost of a leaf u with respect to \mathcal{T} corresponds to the length $\text{len}(u)$ of its path to the root. Note, however,

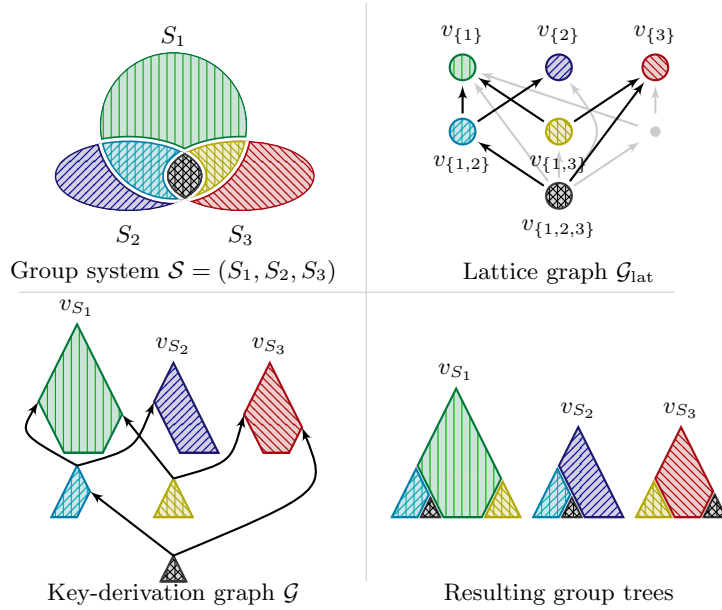


Fig. 4. Working principle of the algorithm. Top left; Venn diagram of the considered group system. Top right; resulting lattice graph after the first phase. Node v_I has associated set $\mathcal{S}(v_I) = P_I$, the set of users in exactly the groups indicated by I . Nodes and edges of the Boolean lattice that are not part of \mathcal{G}_{lat} are depicted in gray. Bottom left; final key derivation graph. Bottom right; resulting trees corresponding to groups S_1 , S_2 , S_3 . Note that components of the same color are shared among different trees.

that leaves of \mathcal{T} may represent more than one user in the key-derivation graph. Indeed, by construction of the algorithm, the weight w_u of u counts the number of leaves in \mathcal{G} below u . Thus, the contribution of Huffman tree \mathcal{T} towards the total update cost of \mathcal{G} is given by $\text{Upd}(\mathcal{T}) = \sum_{u \in \mathcal{L}(\mathcal{T})} w_u \text{len}(u)$. If $U_{\mathcal{T}}$ is the probability distribution that picks $u \in \mathcal{L}(\mathcal{T})$ with probability proportional to its weight w_u , we can express the update cost of \mathcal{T} in terms of the expected length from leaves to the root as

$$\text{Upd}(\mathcal{T}) = \mathbb{E}[\text{len}(U_{\mathcal{T}})] \cdot \sum_{u \in \mathcal{L}(\mathcal{T})} w_u . \quad (6)$$

We first consider Algorithm 1 for the simplest case of two subgroups and compare it to the trivial algorithm.

Example 1. Let $N \in \mathbb{N}$ and let \mathcal{S} consist of two subgroups S_1 , S_2 of sizes N_1 and N_2 respectively. Further assume that $|S_1 \cap S_2| = K$. Consider the key derivation graphs generated by the trivial algorithm and Algorithm 1, which in both cases decompose into several Huffman trees. The trivial algorithm essentially generates




Fig. 5. Key-derivation graphs of the trivial algorithm (left) and Algorithm 1 (right) for two subgroups. Users that are members of both subgroups are marked in thick.

two trees \mathcal{T}'_1 and \mathcal{T}'_2 , the first containing all members of S_1 , the other all members of S_2 . Algorithm 1 first collects the K parties that are members of both groups in a tree $\mathcal{T}_{1,2}$. The remaining $(N_1 - K)$ members of S_1 and the root of $\mathcal{T}_{1,2}$ are collected in a tree \mathcal{T}_1 , the remaining $(N_2 - K)$ members of S_2 and the root of $\mathcal{T}_{1,2}$ in a tree \mathcal{T}_2 (See Figure 5).

By Equation 6 we have

$$\text{Upd}(\mathcal{G}_{\text{triv}}) = \text{Upd}(\mathcal{T}'_1) + \text{Upd}(\mathcal{T}'_2) = N_1 \mathbb{E}[\text{len}(U_{\mathcal{T}'_1})] + N_2 \mathbb{E}[\text{len}(U_{\mathcal{T}'_2})]$$

and

$$\begin{aligned} \text{Upd}(\mathcal{G}_{\text{a1}}) &= \text{Upd}(\mathcal{T}_1) + \text{Upd}(\mathcal{T}_2) + \text{Upd}(\mathcal{T}_{1,2}) \\ &= N_1 \mathbb{E}[\text{len}(U_{\mathcal{T}_1})] + N_2 \mathbb{E}[\text{len}(U_{\mathcal{T}_2})] + K \mathbb{E}[\text{len}(U_{\mathcal{T}_{1,2}})] . \end{aligned}$$

By optimality of Huffman codes (Lemma 1) we have that

$$H(U_{\mathcal{T}}) \leq \mathbb{E}[\text{len}(U_{\mathcal{T}})] \leq H(U_{\mathcal{T}}) + 1$$

for $\mathcal{T} \in \{\mathcal{T}'_1, \mathcal{T}'_2, \mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_{1,2}\}$, where $H(U_{\mathcal{T}})$ is the Shannon entropy of $U_{\mathcal{T}}$. For \mathcal{T}'_1 , \mathcal{T}'_2 , and $\mathcal{T}_{1,2}$ the leaves are distributed uniformly and we have $H(\mathcal{T}'_1) = \log(N_1)$, $H(\mathcal{T}'_2) = \log(N_2)$, $H(\mathcal{T}_{1,2}) = \log(K)$. Let $i \in \{1, 2\}$ and consider \mathcal{T}_i . Then the first $N_i - K$ leaves have probability $1/N_i$ and the last leaf K/N_i . Thus $H(U_{\mathcal{T}_i}) = (N_i - K)/N_i \log(N_i) + K/N_i \log(N_i/K) = \log(N_i) - K/N_i \log(K)$. Summing up we obtain

$$\begin{aligned} &\text{Upd}(\mathcal{G}_{\text{triv}}) - \text{Upd}(\mathcal{G}_{\text{a1}}) \\ &\geq N_1 \log(N_1) + N_2 \log(N_2) - N_1(\log(N_1) - K/N_1 \log(K) + 1) \\ &\quad - N_2(\log(N_2) - K/N_2 \log(K) + 1) - K(\log(K) + 1) \\ &= K(\log(K) - 1) - (N_1 + N_2) . \end{aligned}$$

Note that for $K \geq 2$ the first term is non-negative (For $K = 1$ it is easy to see that Algorithm 1 performs better than the trivial algorithm.).

Before turning to arbitrary group systems, we derive a generalized statement on the update cost $\text{Upd}(\mathcal{T})$ contributed by Huffman trees as defined above. Its proof is in the full version of this work [3].

Lemma 5. Let \mathcal{T} be a Huffman tree over leaves v_1, \dots, v_ℓ of weight $w_1, \dots, w_\ell \in \mathbb{N}$. Let $w = \sum_{i=1}^\ell w_i$. Then \mathcal{T} 's update cost is bounded by

$$w \log(w) - \sum_{i=1}^\ell w_i \log(w_i) \leq \text{Upd}(\mathcal{T}) \leq w(\log(w) + 1) - \sum_{i=1}^\ell w_i \log(w_i) .$$

Regarding general systems of subgroups we obtain the following result. Due to space constraints, we defer the proof to the full version of this work [3].

Theorem 1. Let $N \in \mathbb{N}$, $S_1, \dots, S_k \subseteq [N]$, and \mathcal{G} the key-derivation graph output by Algorithm 1. Let $\mathcal{G}_{\text{lat}} = (\mathcal{V}_{\text{lat}}, \mathcal{E}_{\text{lat}})$ be the corresponding lattice graph. Then

$$\begin{aligned} & \sum_{i=1}^k |S_i| \cdot \log(|S_i|) - \sum_{v \in \mathcal{V}_{\text{lat}} : |I(v)| \geq 2} \left| \bigcup_{v' \in \mathcal{A}(v)} P_{I(v')} \right| \cdot \log \left(\left| \bigcup_{v' \in \mathcal{A}(v)} P_{I(v')} \right| \right) \quad (7) \\ & \leq \text{Upd}(\mathcal{G}) \leq \sum_{i=1}^k |S_i| \cdot (\log(|S_i|) + 1) \\ & \quad - \sum_{v \in \mathcal{V}_{\text{lat}} : |I(v)| \geq 2} \left| \bigcup_{v' \in \mathcal{A}(v)} P_{I(v')} \right| \cdot \left(\log \left(\left| \bigcup_{v' \in \mathcal{A}(v)} P_{I(v')} \right| \right) - 1 \right), \quad (8) \end{aligned}$$

where $\mathcal{A}(v)$ denotes the set of ancestors of v in \mathcal{G}_{lat} , $I(v)$ denotes the set $\{i \in [k] : \exists \text{ path from } v \text{ to } v_{S_i}\}$, and for $I' \subseteq [N]$ the set $P_{I'} := \bigcap_{i \in I'} S_i \setminus \bigcup_{j \in [k] \setminus I'} S_j$ indicates the users exactly in the subgroups corresponding to I' .

The bounds of Theorem 1 depend on the structure of the lattice graph generated by the algorithm. Using Properties 4 and 5 of Lemma 4 to bound $|\mathcal{S}(\mathcal{A}(v'))|$ it is possible to obtain a weaker bound on $\text{Upd}(\mathcal{G})$ that only depends on $[N]$ and \mathcal{S} .

We conclude the section by comparing the update cost of Algorithm 1 to that of the trivial algorithm and the asymptotically optimal algorithm of Section 4.1.

Comparison to the trivial algorithm. Note that the terms $\sum_{i=1}^k |S_i| \cdot \log(|S_i|)$ and $\sum_{i=1}^k |S_i| \cdot (\log(|S_i|) + 1)$ in Theorem 1 match the bounds on the update cost of the trivial algorithm derived in Section 3.4. Thus the second term of

$$\sum_{v \in \mathcal{V}_{\text{lat}} : |I(v)| \geq 2} \left| \bigcup_{v' \in \mathcal{A}(v)} P_{I(v')} \right| \cdot \left(\log \left(\left| \bigcup_{v' \in \mathcal{A}(v)} P_{I(v')} \right| \right) - 1 \right)$$

provides a good estimate on how much Algorithm 1 saves compared to the trivial one. For the group system depicted in Figure 4, for example, this would amount to $|S_1 \cap S_2| \cdot \log(|S_1 \cap S_2|) + |S_1 \cap S_3 \setminus S_2| \cdot \log(|S_1 \cap S_3 \setminus S_2|) + |S_1 \cap S_2 \setminus S_3| \cdot \log(|S_1 \cap S_2 \setminus S_3|)$. Due to the ‘‘rounding error’’ of $+1$ in $\sum_{i=1}^k |S_i| \cdot (\log(|S_i|) + 1)$, Theorem 1 unfortunately does not allow us to conclude that the update cost of Algorithm 1 always improves on the one of the trivial algorithm. In the full version of this work [3], we provide an alternative analysis of $\text{Upd}(\mathcal{G})$ that directly compares the two algorithms and gives conditions that imply Algorithm 1 outperforming the trivial one.

Comparison to the asymptotically optimal algorithm of Section 4.1. The algorithm of Section 4.1 in a first step constructs a binary tree for every non-empty partition $P_{I'}$ and then, in a second step, builds a binary tree for every group using the roots of the “partition trees” as leaves. We can interpret this as an algorithm that, similarly to Algorithm 1, in the first phase chooses a lattice graph \mathcal{G}_{lat} , concretely the graph that connects every node $v_{I'}$ directly with edges to all corresponding group nodes $\{v_{\{i\}} \mid i \in I'\}$, and in the second phase builds Huffman trees for every lattice node.¹⁰

Thus, by Lemma 5, we can lower bound the update cost of key graphs $\mathcal{G}_{\text{asopt}}$ generated by it by $\text{Upd}(\mathcal{G}_{\text{asopt}}) \geq \sum_{i=1}^k \left(|S_i| \cdot \log(|S_i|) - \sum_{I' \subseteq [N]: i \in I' \wedge |I'| \geq 2} |P_{I'}| \cdot \log(|P_{I'}|) \right) + \sum_{I' \subseteq [N]: |I'| \geq 2} |P_{I'}| \cdot \log(|P_{I'}|)$, which, taking into account that every I' with $|I'| = \ell$ corresponds to exactly ℓ groups, simplifies to

$$\text{Upd}(\mathcal{G}_{\text{asopt}}) \geq \sum_{i=1}^k |S_i| \cdot \log(|S_i|) - \sum_{I' \subseteq [N]: |I'| \geq 2} (|I'| - 1) |P_{I'}| \cdot \log(|P_{I'}|) . \quad (9)$$

For a comparison to Algorithm 1, consider a key derivation graph \mathcal{G}_{a1} output by it. We now compute a lower bound on $\text{Upd}(\mathcal{G}_{\text{asopt}}) - \text{Upd}(\mathcal{G}_{\text{a1}})$. Let $\mathcal{G}'_{\text{lat}}$ be the lattice graph of \mathcal{G}_{a1} and $v_{I'} \in \mathcal{G}'_{\text{lat}}$ such that $|I'| \geq 2$. Every non-sink in $\mathcal{G}'_{\text{lat}}$ has outdegree 2 and $v_{I'}$ is connected to all $v_{\{i\}}$ with $i \in I'$ by exactly one path. Thus, the subgraph of $\mathcal{G}'_{\text{lat}}$ induced by these paths is a binary tree with root $v_{I'}$ and $|I'|$ leaves, and thus consists of exactly $2|I'| - 1$ nodes, $|I'|$ of which have an index set of size 1. This implies that there exists $|I'| - 1$ many nodes $v_{I''}$ in $\mathcal{G}'_{\text{lat}}$ with $|I''| \geq 2$ such that $v_{I'} \in \mathcal{A}(v_{I''})$.

Using f as shorthand for the function $f : N \mapsto N \log(N)$ and $p_{I'} = |P_{I'}|$, we now can distribute the expressions $|P_{I'}| \cdot \log(|P_{I'}|)$ of Equation 9 on the negative summands of Equation 8 and obtain

$$\text{Upd}(\mathcal{G}_{\text{asopt}}) - \text{Upd}(\mathcal{G}_{\text{a1}}) \geq \sum_{v \in \mathcal{V}'_{\text{lat}}: |I(v)| \geq 2} \left(f\left(\sum_{v' \in \mathcal{A}(v)} p_{I(v')}\right) - \sum_{v' \in \mathcal{A}(v)} f(p_{I(v')} - 1) \right) .$$

Since f grows super-linearly, the terms $f(\sum_{v' \in \mathcal{A}(v)} p_{I(v')}) - \sum_{v' \in \mathcal{A}(v)} f(p_{I(v')} - 1)$ are non-negative, and can even be of order N as for example $f(2N/2) - 2f(N/2) = N$. While, again due to the terms -1 , we are unfortunately not able to conclude that Algorithm 1 is always more efficient, this shows that it still can save substantially in terms of update cost, in particular if the $p_{I'}$ are large.

In this section we were considering the total update cost of key-derivation graphs generated by Algorithm 1, which relates to the average update cost of parties. As we have shown, this metric will typically improve compared to the trivial algorithm. However, it might still be possible, that the update cost of

¹⁰ Formally, the algorithm as described in Section 4.1 collects all users that are *only* in group S_i in a tree before computing the tree for S_i , while in the lattice-graph variant these users are directly included in the tree for S_i . Note, however, that the latter approach can only improve the total update cost.

particular, fixed users increases. In the full version of this work [3] we show that while this may indeed happen, the increase is essentially bounded by a small constant.

5.3 Asymptotic Optimality of Boolean-lattice based Graphs

As discussed in Section 5.1, we can interpret our algorithm as follows. On input $([N], \mathcal{S} = \{S_1, \dots, S_k\})$, in the first phase the algorithm picks a subgraph of the Boolean lattice $\mathcal{G}_B = (\mathcal{V}_B, \mathcal{E}_B)$ with respect to the power set of $[k]$, where

$$\mathcal{V}_B = \{v_I \mid I \subseteq [k]\} \quad \text{and} \quad \mathcal{E}_B = \{(v_I, v_{I'}) \mid I, I' \subseteq [k] : I' \subseteq I\} .$$

We refer to this subgraph as the lattice graph. In the second phase, for $I \subseteq [k]$, a source for every party in P_I , i.e., the set of parties belonging exactly to the groups specified by I , is added and connected to node v_I . Each node in the graph is assigned a weight; sources have weight 1 and the weight of all other nodes is the sum of the weights of their parents. Finally, for every v_I a Huffman tree to its parents according to the weight distribution is built, resulting in the key-derivation graph.

In this section we consider key-derivation graphs for general choices of the lattice graph, i.e., key derivation graphs \mathcal{G} obtained by executing the second phase of the algorithm as described above with respect to a lattice-graph $\mathcal{G}_{\text{lat}} = (\mathcal{V}_{\text{lat}}, \mathcal{E}_{\text{lat}}) \subseteq \mathcal{G}_B$.¹¹ We say \mathcal{G} is the key-derivation graph associated to \mathcal{G}_{lat} , $[N]$ and \mathcal{S} . The following theorem shows that the update cost of every lattice-based key derivation graphs, and in particular graphs generated by Algorithm 1, is optimal in the asymptotic setting of Section 4. Due to space constraints, we defer its proof to the full version of this work [3].

Theorem 2. *Let $k \in \mathbb{N}$ be fixed, and for $I \subseteq [k]$ let $p_I \in [0, 1]$ be such that $\sum_{I \subseteq [k]} p_I = 1$ and $p_\emptyset = 0$. For $N \in \mathbb{N}$ let $\mathcal{S}(N)$ be the subgroup system associated to the p_I .*

Let $\mathcal{G}_{\text{lat}} = (\mathcal{V}_{\text{lat}}, \mathcal{E}_{\text{lat}})$ be a subgraph of the Boolean-lattice graph with respect to $[k]$ satisfying that $v_I \in \mathcal{V}_{\text{lat}}$ for all I with $p_I > 0$, and let $\mathcal{G}(N)$ be the key-derivation graph associated to \mathcal{G}_{lat} and $\mathcal{S}(N)$. Then

$$\text{Upd}(\mathcal{G}(N)) \xrightarrow{N \rightarrow \infty} \sum_{I \subseteq [k]} |N \cdot p_I| \cdot \log(|N \cdot p_I|) + \Theta(N) = N \log(N) + \Theta(N) .$$

6 Lower Bound on the Update Cost of CGKA

In this section we prove a lower bound on the average update cost of continuous group-key agreement schemes for multiple groups. As an intermediate step we will further prove a bound on the update cost of key-derivation graphs. To this

¹¹ Naturally, one would require that the resulting key-derivation graph satisfies correctness. However, this is not necessary for our analysis of its update cost.

aim, we follow the approach of Micciancio and Panjwani [14], who analyzed the worst-case communication complexity of multicast key distribution in a *symbolic* security model, where cryptographic primitives are considered as abstract data types. We will first recall their security model, adapt it to CGKA, and then prove how to extend their results to our setting. In the full version of this work [3], using a similar approach, we prove a lower bound for multicast encryption.

6.1 Symbolic Model

We first define a symbolic model in the style of Dolev and Yao [10] for CGKA schemes. It follows the approach of Micciancio and Panjwani [14], but as it admits the uses of public-key encryption also includes elements of the model of Bienstock et al. [6], who analyze the communication cost of concurrent updates in CGKA schemes.

Building blocks. We restrict the analysis to schemes that are constructed from the following three primitives. Note that our construction is a special case of the constructions analyzed in this section.

- *Public-key Encryption:* Let $(\text{KGen}, \text{Enc}, \text{Dec})$ denote a public-key encryption scheme, where
 - KGen on input of secret key sk returns the corresponding public key pk .
 - Enc takes as input a public key pk and a message m , and outputs a ciphertext $c \leftarrow \text{Enc}(\text{pk}, \text{m})$.
 - Dec takes as input a secret key sk and a ciphertext c , and outputs a message $\text{m} = \text{Dec}(\text{sk}, c)$. We assume perfect correctness: $\text{Dec}(\text{sk}, \text{Enc}(\text{pk}, \text{m})) = \text{m}$ for all sk , $\text{pk} = \text{KGen}(\text{sk})$, and messages m .
- *Pseudorandom generator:* The algorithm G takes as input a secret key sk and expands it to a sequence of keys $\text{G}_0(\text{sk}), \dots, \text{G}_\ell(\text{sk})$.
- *Secret sharing:* Let S, R denote the sharing and recovering procedures of a secret sharing scheme: For some access structure $\Gamma \subseteq 2^{[h]}$, the algorithm S takes as input a message m and outputs a set of shares $\text{S}_1(\text{m}), \dots, \text{S}_h(\text{m})$ such that for any $I \in \Gamma$ it holds $\text{R}(I, \{\text{S}_i(\text{m})\}_{i \in I}) = \text{m}$, but for any $I \not\subseteq \Gamma$ the message m cannot be recovered from $\{\text{S}_i(\text{m})\}_{i \in I}$.

We consider the following data types that can be derived from other objects according to the following rules.

Data type	Grammar rules
Message m	$\leftarrow \text{sk}, \text{pk}, \text{Enc}(\text{pk}, \text{m}), \text{S}_1(\text{m}), \dots, \text{S}_h(\text{m})$
Public key pk	$\leftarrow \text{KGen}(\text{sk})$
Secret key sk	$\leftarrow \text{R}, \text{G}_0(\text{sk}), \dots, \text{G}_\ell(\text{sk})$

To describe the information that can be recovered from a set of messages M , the *entailment relation* is defined by the following rules:

$$\begin{aligned}
 \text{m} \in M &\Rightarrow M \vdash \text{m} \\
 M \vdash \text{sk} &\Rightarrow M \vdash \text{G}_0(\text{sk}), \dots, \text{G}_\ell(\text{sk}) \\
 M \vdash \text{Enc}(\text{pk}, \text{m}), \text{sk} : \text{pk} = \text{KGen}(\text{sk}) &\Rightarrow M \vdash \text{m} \\
 \exists I \in \Gamma : \forall i \in I : M \vdash \text{S}_i(\text{m}) &\Rightarrow M \vdash \text{m}
 \end{aligned}$$

By restricting to these relations we essentially assume *secure* encryption and secret sharing schemes. Examples and further comments (in the setting of multicast encryption) can be found in [14, Section 3.2]. The set of messages which can be recovered from M using relation \vdash is denoted by $\text{Rec}(M)$.

Continuous group-key agreement. We now define continuous group-key agreement protocols in the symbolic model. We consider the case of CGKA for a static system of users $[N]$ and groups $S_1, \dots, S_k \subseteq [N]$. Note that a lower bound for schemes in this setting in particular also excludes schemes which allow dynamic operations, i.e., adding and removing users from groups.

A CGKA scheme for $[N]$ and S_1, \dots, S_k specifies two procedures:

- Initially, **Setup** assigns each user $n \in [N]$ a personal set SK_n^0 of secret keys. Furthermore, **Setup** generates a set $\text{msgs}(0)$ of so-called *rekey* messages to establish for every group S_j a group secret key $\text{sk}_{S_j}^0$. We require that the initial sets of personal keys consist of uniformly random keys, and that for all $n' \neq n$ and $\text{sk} \in \text{SK}_{n'}^0$ we have $\text{sk} \notin \text{Rec}(\text{SK}_{n'}^0, \text{msgs}(0))$.
- In round t , the algorithm **Update** takes as input a user identity $n \in [N]$, establishes new sets $\text{SK}_{n'}^t$ for all users n' , and outputs some rekey messages $\text{msgs}(t)$ to establish for every group S_j an epoch t group key $\text{sk}_{S_j}^t$. We do *not* require the new sets and group keys to be distinct from the ones of round $t - 1$. We denote the set of new uniformly random keys that were generated during the update procedure by the updating party by F_n^t .

Note that the only party generating new keys during update t is the updating party n . For ease of notation we define $\text{F}_{n'}^t = \emptyset$ for all $n' \neq n$, and set $\text{F}_n^0 = \emptyset$ for all n .

For *correctness*, we require that, (a) at all times members of a group are able to derive the current group key from their set of personal keys and the sent messages, and (b) if some user updated in round t , then all users are able to derive their new set of personal keys from their old one, the sent messages, and in the case of the updating party the new keys generated during the update. The latter condition accounts for the fact that changes to a user's set of personal keys need to be communicated to them.

More precisely, for (a) we require that for any subgroup structure and any sequence of updating users (n_1, \dots, n_t) , for all $j \in [k]$ each member n of subgroup S_j can recover $\text{sk}_{S_j}^t$:

$$\text{sk}_{S_j}^t \in \text{Rec}\left(\text{SK}_n^t \cup \bigcup_{\iota \in [t]_0} \text{msgs}(\iota)\right) .$$

For (b) we require that for any subgroup structure and any sequence of updating users (n_1, \dots, n_t) , we have for all n that

$$\text{SK}_n^t \subseteq \text{Rec}\left(\text{SK}_n^{t-1} \cup \text{F}_n^t \cup \bigcup_{\iota \in [t]_0} \text{msgs}(\iota)\right) .$$

For security, we assume the minimal requirement of *post-compromise security* (PCS), which essentially says that users can recover from compromise, which leaks their state and the keys generated during the time period of being compromised, by updating. Note that a lower bound in this setting in particular excludes protocols achieving stronger security notions desired in practice, like post compromise forward security [2].

More precisely, we formalize PCS as the condition that no group key can be recovered from members outside the group, and/or members' personal keys and the keys generated by them before their last update. To this end, for round t and user $n \in [N]$, let $t_{\text{up}}(t, n)$ denote the round in which n performed their last update, where we set $t_{\text{up}}(t, n) = 0$ if no such update occurred. I.e., we require that for any group system, any update pattern, in every round t we have that

$$\text{sk}_{S_j}^t \notin \text{Rec} \left(\bigcup_{\substack{n \in [N] \setminus S_j, \\ t' \in [t]_0}} (\text{SK}_n^{t'} \cup \text{F}_n^{t'}) \cup \bigcup_{n \in S_j} \bigcup_{t' \in [t_{\text{up}}(t, n) - 1]_0} (\text{SK}_n^{t'} \cup \text{F}_n^{t'}) \cup \bigcup_{t' \in [t]_0} \text{msgs}(t') \right).$$

Note that in the definition above, excluding all sets of personal secret keys since a user's last update is necessary even in the case that another user's update might have replaced them before round t , as otherwise SK_n^t and in turn $\text{sk}_{S_j}^t$ could trivially be recovered by the two correctness conditions.

Our goal is to derive a lower bound on the communication complexity of CGKA schemes achieving PCS, i.e., the number of messages $|\bigcup_{t' \in [t]_0} \text{msgs}(t')|$ sent by the protocol.

Key Graphs. The execution of any CGKA scheme can be reflected by a graph structure representing recoverability of the keys involved (cf. [14]). To define this graph, we first need to recall the definition of *useful* keys and messages.

A secret key sk is called *useless at time t* if it can be recovered from old key material, i.e., if

$$\text{sk} \in \text{Rec} \left(\bigcup_{n \in [N]} \bigcup_{t' \in [t_{\text{up}}(t, n) - 1]_0} (\text{SK}_n^{t'} \cup \text{F}_n^{t'}) \cup \bigcup_{t' \in [t]_0} \text{msgs}(t') \right),$$

otherwise sk is called *useful*. As we will show below, if a CGKA scheme satisfies correctness and post-compromise security, then for all $t \in \mathbb{N}$, $n \in [N]$, $j \in [k]$ it must hold that at least one of the user's personal keys $\text{sk}_n^t \in \text{SK}_n^t$ as well as all group keys $\text{sk}_{S_j}^t$ are useful at time t .

To decide whether a message is useful, one has to consider the information it contains, where messages can be arbitrarily nested applications of encryption Enc and secret sharing S . Thus, a message m is said to *encapsulate* a (pseudo)random key sk if $\text{m} = e_1(e_2(\dots(e_j(\text{sk}))\dots))$ where $e_i = \text{Enc}_{\text{pk}_i}$ or $e_i = \text{S}_{h_i}$ (for some public key pk_i and $h_i \in [h]$). A message is then called *useful* if it encapsulates a useful key.

Definition 2 (Key graph [14]). *The key graph $\mathcal{KG}_t = (\mathcal{V}_t, \mathcal{E}_t)$ for a CGKA scheme at time t is defined as follows. \mathcal{V}_t consists of all the keys that are useful*

at time t , and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ consists of all ordered pairs $(\text{sk}_1, \text{sk}_2)$ such that one of the following is true:

1. There exists $j \in [l]$ such that $\text{sk}_2 = G_j(\text{sk}_1)$.
2. There exists a message $\mathbf{m} \in \bigcup_{j \in [t]_0} \text{msgs}(j)$ with $\mathbf{m} = e_1(\text{Enc}(\text{pk}_1, e_2(\text{sk}_2)))$ with $\text{pk}_1 = \text{KGen}(\text{sk}_1)$. Here e_1 and e_2 are some sequences of encryption and secret sharing, and we require that e_2 does not contain any encryption under a public key that has a matching secret key that is useful at time t .

Edges of the second type are called communication edges.

One can show that for any node sk in \mathcal{KG} there is at most one edge of the first type incident to sk (the proof is analogous to [14, Proposition 1]). Note that edges of the first type do not incur any communication cost, while edges of the second type require at least one message. Thus, in the following we will be interested in the number of communication edges. To this aim, we prove the following properties of key graphs. In particular, we show that even if a CGKA scheme does not rely on the use of a fixed key-derivation graph as discussed in Section 3, after every update the key graph must still have the properties of Definition 1.

We will rely on the following Lemma that can be proved analogously to [14, Lemma 1].

Lemma 6. *Consider a secure and correct CGKA scheme for $N \in \mathbb{N}$, $S_1, \dots, S_k \subseteq [N]$. Then, for any $t \in \mathbb{N}$ and sequence of updates (n_1, \dots, n_t) , the corresponding key graph \mathcal{KG}_t satisfies the following. For every set of keys SK , and key sk_2 that is useful at time t , such that $\text{sk}_2 \in \text{Rec}\left(\text{SK} \cup \bigcup_{t' \in [t]_0} \text{msgs}(t')\right)$, there exists a useful $\text{sk}_1 \in \text{SK}$ such that there is a path from sk_1 to sk_2 in \mathcal{KG}_t that only consists of keys sk with $\text{sk} \in \text{Rec}\left(\text{SK} \cup \bigcup_{t' \in [t]_0} \text{msgs}(t')\right)$.*

Note that the converse of Lemma 6 is not true, since, for example, a message $\text{Enc}(\text{pk}_1, S_1(\text{sk}_2))$ with useful keys sk_1, sk_2 and $\text{pk}_1 = \text{KGen}(\text{sk}_1)$ incurs an edge $(\text{sk}_1, \text{sk}_2)$ while sk_2 can only be recovered from sk_1 if $\{1\} \in \Gamma$.

6.2 Lower Bound on the Average Update Cost.

The communication complexity of a CGKA scheme after t updates is given by $\left| \bigcup_{t' \in [t]_0} \text{msgs}(t') \right|$. To measure the efficiency of the scheme we will consider the *amortized communication complexity*

$$\text{Com}_A := \left| \bigcup_{t' \in [t]_0} \text{msgs}(t') \right| / t .$$

We now are ready to compute a bound on the expectation of Com_A in the scenario where, in every round, the updating party is chosen uniformly at random. In the full version of this work [3] we prove an analogous bound for multicast

encryption that improves on [14, Theorem 1] in two aspects. It generalizes the bound to the setting of several, potentially overlapping groups, and further gives a bound on the *average* communication complexity of updates, as opposed to a worst case bound.

Theorem 3. *Consider a CGKA scheme CGKA for $N \in \mathbb{N}$, $S_1, \dots, S_k \subseteq [N]$ that is secure in the symbolic model. Then the expected amortized average communication cost after t updates is bounded from below by*

$$\mathbb{E}[\text{Com}_A] \geq (1 - 1/t) \cdot \frac{1}{N} \sum_{\emptyset \neq I \subseteq [k]} |P_I| \cdot \log(|P_I|) .$$

and the asymptotic (in the number of update operations) update cost of the protocol is at least $\frac{1}{N} \sum_{\emptyset \neq I \subseteq [k]} |P_I| \cdot \log(|P_I|)$.

Due to space constraints, we defer the proof to the full version of this work [3].

7 Open problems

We conclude by discussing some open problems.

7.1 Optimal Key-derivation Graphs

Unfortunately we are not able to tell how far from optimal the solutions generated by Algorithm 1 are for concrete group systems. We consider it an interesting open question to resolve this issue.

General kdgs. We first discuss this problem in its general form. I.e., given a system $\mathcal{S} = \{S_1, \dots, S_k\}$ of subgroups of the set $[N]$ of users compute the key-derivation graph for \mathcal{S} (as defined in Definition 1) that has minimal update cost. The question of whether a polynomial time algorithm for solving this problem exists can be naturally asked in various ways. E.g., when polynomial means polynomial in the number of users N (think of N being given in unary), or polynomial in a reasonable description of the set system \mathcal{S} , say, when we are given the sizes of all non-empty intersections of sets in \mathcal{S} . Here N can be exponential in the input length, so a potential solution would need to have a very succinct description. Algorithm 1 (which we do not know whether is optimal) can be turned into one of the latter kind by using an implicit representation during the Huffman coding step.

We are thankful to one reviewer of this work, who pointed out an interesting connection of key-derivation graphs for a group system $\mathcal{S} = \{S_1, \dots, S_k\}$ to the *disjunctive complexity* of \mathcal{S} , which, given variables $x_1, \dots, x_N \in \{0, 1\}$, corresponds to the size of the smallest circuit of fanin-2 OR-gates computing

$$\bigvee_{i \in S_1} x_i, \dots, \bigvee_{i \in S_k} x_i . \tag{10}$$

Note that circuits computing (10) correspond exactly to key-derivation graphs for \mathcal{S} . So the two problems differ only by the used metric; while disjunctive complexity counts the number of non-sources in the graph, the update cost of a kdg weighs each of these nodes by the number of sources below it. As there exist upper and lower bound on the disjunctive complexity of group systems (see e.g. [12]), we consider it an interesting open questions whether these can be used to establish bounds on the update cost of kdgs. We want to point out, however, that this metric might be not fine-grained enough to capture certain properties of kdgs: E.g., for $N \in \mathbb{N}$ the systems $\mathcal{S}_1 = \{[N]\}$ and $\mathcal{S}_2 = \{[1], [2], \dots, [N]\}$ both have disjunctive complexity $N - 1$, but their total update costs as kdgs are of order $N \cdot \log(N)$ and N^2 respectively.

Lattice based kdgs. If we restrict our view to algorithms using Boolean-lattice based graphs as defined in Section 5.3, and are willing to make simplifying assumptions, the question of optimality translates to an optimization problem on graphs: we are going to (a) consider only lattice graphs \mathcal{G}_{lat} where all nodes v are connected with their descendants $v' \in \mathcal{D}(v)$ by an *unique* path, and (b) assume in our analysis of the update cost that the algorithms second step (i.e., the generation of Huffman trees) is instead implemented with an idealized code, that has average codeword length matching the entropy of the leaf distribution. This essentially corresponds to ignoring the terms of +1 in Lemma 1.

Recall that for groups system $\{S_1, \dots, S_k\}$ the nodes $v_I \in \mathcal{V}_{\text{lat}}$ of a lattice graph correspond to index sets $I \subseteq [k]$. It is easy to see that the correctness of \mathcal{G}_{lat} together with condition (a), are equivalent to requiring that the only sinks in the graph are the singleton sets $\{i\}$, and that for every $v_I \in \mathcal{V}_{\text{lat}}$

$$I = I_1 \cup \dots \cup I_\ell \tag{11}$$

holds, where $v_{I_1}, \dots, v_{I_\ell}$ are the children of v_I and disjointness enforces unique paths.

The total update cost of a graph satisfying this property can be computed as follows. To every node v_I we associate the weight $w_I = \left| \bigcap_{i \in I} S_i \setminus \bigcup_{j \in [k] \setminus I} S_j \right|$ corresponding to the number of users exactly in the groups specified by I . Further, we inductively define the total weight t_I of v_I as

$$t_I = \begin{cases} w_I & \text{if } v_I \text{ is source} \\ w_I + \sum_{I': v_{I'} \in \mathcal{P}(v_I)} t_{I'} & \text{else} \end{cases},$$

where $\mathcal{P}(v_I)$ denotes the set of parents of v_I . By assumptions (a) and (b), and Lemma 5, the update cost contributed by node v_I thus corresponds to

$$\text{Upd}(v_I) = t_I \log(t_I) - \sum_{I': v_{I'} \in \mathcal{P}(v_I)} t_{I'} \log(t_{I'}) , \tag{12}$$

and we end up with the following optimization problem on lattice graphs.

Problem 1. Let $k \in \mathbb{N}$. Given weights $\{w_I\}_{I \subseteq [k]}$ with $w_I \in \mathbb{N}$ among the subgraphs of the Boolean lattice with respect to the power set of $[k]$ that satisfy Condition 11 find the subgraph \mathcal{G}_{lat} of minimal total update cost

$$\text{Upd}(\mathcal{G}_{\text{lat}}) = \sum_{I \subseteq [k]} \text{Upd}(v_I) .$$

We consider it an interesting open question whether Algorithm 1 solves this problem and, if not, to find an efficient algorithm that does.

7.2 Security

In this work we focused on the communication complexity of key-derivation graphs and only gave an intuition on their security. Security proofs for secure group messaging are typically quite complex, and protocols rely on additional mechanisms (e.g. confirmation tag, transcript hash, and parent hash) ensuring that users of the system can not be tricked into inconsistent views of the graph. We consider it an important open question, to adapt these mechanisms to kdgs for several groups and give a formal security proof for the resulting CGKA protocols.

7.3 Efficiency of Dynamic Operations

In the full version of this work [3] we show that the techniques of blanking and unmerged leaves can be adapted to key-derivation graphs, in order to allow dynamic changes to the group membership. As is the case for singular groups, blanking and unmerged leaves decrease the efficiency of updates of a user n , since they destroy the binary structure of the graph, resulting in potentially more than a single ciphertext per node in $\mathcal{D}(v_n)$ having to be generated. However, the graph gradually recovers from this, assuming that parties with update trees overlapping $\mathcal{D}(v_n)$ update. It is an interesting open question how the decrease in efficiency compares to that of the trivial algorithm.

References

1. Alwen, J., Coretti, S., Dodis, Y., Tselekounis, Y.: Security analysis and improvements for the IETF MLS standard for group messaging. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020, Part I. LNCS, vol. 12170, pp. 248–277. Springer, Heidelberg (Aug 2020)
2. Alwen, J., Coretti, S., Jost, D., Mularczyk, M.: Continuous group key agreement with active security. In: Pass, R., Pietrzak, K. (eds.) TCC 2020, Part II. LNCS, vol. 12551, pp. 261–290. Springer, Heidelberg (Nov 2020)
3. Alwen, J., Auerbach, B., Baig, M.A., Cueto, M., Klein, K., Pascual-Perez, G., Pietrzak, K., Walter, M.: Grafting key trees: Efficient key management for overlapping groups. Cryptology ePrint Archive, Report 2021/1158 (2021), <https://ia.cr/2021/1158>

4. Barnes, R., Beurdouche, B., Millican, J., Omara, E., Cohn-Gordon, K., Robert, R.: The Messaging Layer Security (MLS) Protocol. Internet-Draft draft-ietf-mls-protocol-11, Internet Engineering Task Force (Dec 2020), <https://datatracker.ietf.org/doc/html/draft-ietf-mls-protocol-11>, work in Progress
5. Bhargavan, K., Barnes, R., Rescorla, E.: TreeKEM: Asynchronous Decentralized Key Management for Large Dynamic Groups (May 2018)
6. Bienstock, A., Dodis, Y., Rösler, P.: On the price of concurrency in group ratcheting protocols. In: Pass, R., Pietrzak, K. (eds.) TCC 2020, Part II. LNCS, vol. 12551, pp. 198–228. Springer, Heidelberg (Nov 2020)
7. Canetti, R., Garay, J.A., Itkis, G., Micciancio, D., Naor, M., Pinkas, B.: Multicast security: A taxonomy and some efficient constructions. In: IEEE INFOCOM'99. pp. 708–716. New York, NY, USA (Mar 21–25, 1999)
8. Cohn-Gordon, K., Cremers, C., Garratt, L., Millican, J., Milner, K.: On ends-to-ends encryption: Asynchronous group messaging with strong security guarantees. In: Lie, D., Mannan, M., Backes, M., Wang, X. (eds.) ACM CCS 2018. pp. 1802–1819. ACM Press (Oct 2018)
9. Cremers, C., Hale, B., Kohbrok, K.: Efficient post-compromise security beyond one group. Cryptology ePrint Archive, Report 2019/477 (2019), <https://eprint.iacr.org/2019/477>
10. Dolev, D., Yao, A.: On the security of public key protocols. IEEE Transactions on Information Theory 29(2), 198–208 (1983)
11. Huffman, D.A.: A method for the construction of minimum-redundancy codes. Proceedings of the IRE 40(9), 1098–1101 (1952)
12. Jukna, S.: Boolean function complexity: advances and frontiers, vol. 27. Springer Science & Business Media (2012)
13. Mapoka, T.T., Shepherd, S., Abd-Alhameed, R., Anoh, K.O.: Novel rekeying approach for secure multiple multicast groups over wireless mobile networks. In: 2014 International Wireless Communications and Mobile Computing Conference (IWCMC). pp. 839–844. IEEE (2014)
14. Micciancio, D., Panjwani, S.: Optimal communication complexity of generic multicast key distribution. In: Cachin, C., Camenisch, J. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 153–170. Springer, Heidelberg (May 2004)
15. Wallner, D.M., Harder, E.J., Agee, R.C.: Key management for multicast: Issues and architectures. Internet Draft (Sep 1998), <http://www.ietf.org/ID.html>
16. Wong, C.K., Gouda, M., Lam, S.S.: Secure group communications using key graphs. IEEE/ACM Transactions on Networking 8(1), 16–30 (Feb 2000)
17. Zhong, H., Luo, W., Cui, J.: Multiple multicast group key management for the internet of people. Concurrency and Computation: Practice and Experience 29(3), e3817 (2017), <https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.3817>, e3817 CPE-15-0502.R1