A Quasipolynomial Reduction for Generalized Selective Decryption on Trees

Georg Fuchsbauer[1,*], Zahra Jafargholi[2], and Krzysztof Pietrzak[1,*]

[1] Institute of Science and Technology Austria
{gfuchsbauer|pietrzak}@ist.ac.at
[2] Northeastern University, USA
zahra@ccs.neu.com

Abstract. Generalized Selective Decryption (GSD), introduced by Panjwani [TCC'07], is a game for a symmetric encryption scheme Enc that captures the difficulty of proving adaptive security of certain protocols, most notably the Logical Key Hierarchy (LKH) multicast encryption protocol. In the GSD game there are n keys $k_1, \ldots, k_n$, which the adversary may adaptively corrupt (learn); moreover, it can ask for encryptions $\mathsf{Enc}_{k_i}(k_j)$ of keys under other keys. The adversary's task is to distinguish keys (which it cannot trivially compute) from random. Proving the hardness of GSD assuming only IND-CPA security of Enc is surprisingly hard. Using "complexity leveraging" loses a factor exponential in n, which makes the proof practically meaningless.

We can think of the GSD game as building a graph on n vertices, where we add an edge $i \to j$ when the adversary asks for an encryption of k_j under k_i. If restricted to graphs of depth ℓ, Panjwani gave a reduction that loses only a factor exponential in ℓ (not n). To date, this is the only non-trivial result known for GSD.

In this paper we give almost-polynomial reductions for large classes of graphs. Most importantly, we prove the security of the GSD game restricted to trees losing only a quasi-polynomial factor $n^{3\log n + 5}$. Trees are an important special case capturing real-world protocols like the LKH protocol. Our new bound improves upon Panjwani's on some LKH variants proposed in the literature where the underlying tree is not balanced. Our proof builds on ideas from the "nested hybrids" technique recently introduced by Fuchsbauer et al. [Asiacrypt'14] for proving the adaptive security of constrained PRFs.

1 Introduction

Proving security of protocols where an adversary can make queries and/or corrupt players *adaptively* is a notoriously hard problem. Selective security, where the adversary must commit to its queries before the protocol starts, often allows for an easy proof, but in general does not imply (the practically relevant) adaptive security notion [CFGN96].

Panjwani [Pan07] argues that the two common approaches to achieving adaptive security, namely requiring that all parties erase past data [BH93], or using *non-committing* encryption [CFGN96] are not satisfactory. He introduces the *generalized selective decryption* (GSD) problem and uses it as an abstraction of security requirements of multicast encryption protocols [WGL00,MP06]. GSD is defined by a very simple game that captures the difficulty of proving adaptive security of some interesting protocols.

The generalized selective decryption (GSD) game. In the GSD game we consider a symmetric encryption scheme Enc and a parameter $n \in \mathbb{N}$. Initially, we sample n random keys $k_1, \ldots, k_n$ and a bit $b \in \{0, 1\}$. During the game the adversary A can make two types of queries. Encryption query: on input (i, j) she receives $c = \mathsf{Enc}_{k_i}(k_j)$; corruption query: on input i, she receives k_i. At some point, A chooses some i to be challenged on. If $b = 0$, she gets the key k_i; if $b = 1$, she gets a uniformly random r_i.[3] Finally, A outputs a guess bit b'. The goal is prove

---

[3] Below, we will consider a (seemingly) different experiment and output k_i in both cases ($b = 0$ and $b = 1$), but if $b = 1$, then on any query (j, i), we will encrypt $\mathsf{Enc}_{k_j}(r_i)$ and not $\mathsf{Enc}_{k_j}(k_i)$. This is

that for any efficient A, $|\Pr[b = b'] - \frac{1}{2}|$ is negligible (or, equivalently, k_i is pseudorandom) assuming only that Enc is a secure encryption scheme. We only allow one challenge query, but this notion is equivalent to allowing any number of challenge queries by a standard hybrid argument (losing a factor that is only the number of challenge queries).

It is convenient to think of the GSD game as dynamically building a graph, which we call key graph. We start with a graph with n vertices labeled $1, \ldots, n$, where we associate vertex i with key k_i. On an encryption query $\mathsf{Enc}_{k_i}(k_j)$ we add a directed edge $i \to j$. On a corruption query i we label the vertex i as corrupted. Note that if i is corrupted then A also learns all keys k_j for which there is a path from i to j in the key graph by simply decrypting the keys along that path. To make the game non-trivial, challenge queries are thus only allowed for keys that are not reachable from any corrupted key. Another restriction we must make is to disallow encryption cycles, i.e., loops in the graph. Otherwise we cannot hope to prove security assuming only standard security (in our case IND-CPA) of the underlying encryption scheme, as this would require circular (or key-dependent-message) security [BRS03], which is stronger than IND-CPA [ABBC10]. Finally, we require that the challenge query is a leaf in the graph; this restriction too is necessary unless we make additional assumptions on the underlying encryption scheme (cf. Footnote 11).

SELECTIVE SECURITY OF GSD. In order to prove security of the GSD game, one must turn an adversary A that breaks the GSD game with some advantage $\epsilon = |\Pr[b = b'] - \frac{1}{2}|$ into an adversary B that breaks the security of Enc with some advantage $\epsilon' = \epsilon'(\epsilon)$. The security notion we consider is the standard notion of indistinguishability under chosen plaintext attacks (IND-CPA). Recall that in the IND-CPA game an adversary B is given access to an encryption oracle $\mathsf{Enc}_k(\cdot)$. At some point B chooses a pair of messages (m_0, m_1), then gets a challenge ciphertext $c = \mathsf{Enc}_k(m_b)$ for a random bit b, and must output a guess b'. The advantage of B is $|\Pr[b = b'] - \frac{1}{2}|$.

It is not at all clear how to construct an adversary B that breaks IND-CPA from an A that breaks GSD. This problem becomes much easier if we assume that A breaks the *selective* security of GSD, where A must choose all its encryption, corruption and challenge queries before the experiment starts.

In fact, it is sufficient to know the topology of the connected component in the key graph that contains the challenge node. Let α denote the number of edges in this component. One can now define a sequence of 2α hybrid games $H_0, \ldots, H_{2\alpha-1}$, where the first game is the real game (i.e., the GSD game with $b = 0$ where the adversary gets the key), the last hybrid is the random game ($b = 1$), and moreover, from any adversary that distinguishes H_i from H_{i+1} with some advantage ϵ', we get an adversary against the IND-CPA security of Enc with the same advantage. Thus, given an A breaking GSD with advantage ϵ, we can break the IND-CPA security with advantage $\epsilon' \geq \epsilon/(2\alpha - 1) \geq \epsilon/n^2$ (as an n vertex graph has $\leq n^2$ edges). We illustrate this reduction in Fig. 1.

ADAPTIVE SECURITY OF GSD. In the selective security proof for GSD we crucially relied on the fact that we knew the topology of the underlying key graph. Proving adaptive security, where the adversary decides what queries to ask adaptively during the experiment, is much more difficult. A generic trick to prove adaptive security is "complexity leveraging", where one simply turns an adaptive adversary into a selective one by initially guessing the adaptive adversary's choices and committing to those (as required by the selective security game). If during the security game the adaptive choices by the adversary disagree with the guessed ones, we simply abort. The problem with this approach is that assuming the adaptive adversary has advantage ϵ, the constructed selective adversary only has advantage ϵ/P where $1/P$ is the probability of that our guess is correct, which is typically exponentially small. Concretely, in the GSD game

---

just a semantic change assuming the following: during the experiment we always answer encryption queries of the form (a, b) with $\mathsf{Enc}_{k_a}(k_b)$ (note that we don't know if we're encrypting the challenge at this point), and once the adversary chooses a challenge i, if $b = 1$, we simply switch the values of r_i and k_i (this trick is already used in [Pan07]).

we need to guess the nodes in the connected component containing the challenge, and as the number of such choices is exponential in the number of keys n, this probability is $2^{-\Theta(n)}$.

No proofs for the adaptive security of GSD with a subexponential (in n) security loss are known in general. But remember that the GSD problem abstracts problems we encounter in proving adaptive security of many real-world applications where the underlying key graph is typically not completely arbitrary, but often has some special structure. Motivated by this, Panjwani [Pan07] investigated better reductions assuming some special structure of the key graph. He gives a proof where the security degradation is only exponential in the *depth* of the key graph, as opposed to its size. Concretely, he proves that if the encryption scheme is ϵ-IND-CPA secure then the adaptive GSD game with n keys where the adversary is restricted to key graphs of depth ℓ is ϵ'-secure where

$$\epsilon' = \epsilon \cdot O(n \cdot (2n)^\ell) \ .$$

Until today, Panjawain's bound is the only non-trivial improvement over the $2^{\Theta(n)}$ loss for GSD.

Our result. The main result of this paper is Theorem 2, which states that GSD restricted to trees can be proven secure with only a quasi-polynomial loss

$$\epsilon' = \epsilon \cdot n^{3\log(n)+5} \ .$$

Our bound is actually even stronger as the entire key graph need not be a tree; it is sufficient that the subgraph containing only the nodes from which the challenge node can be reached is a tree (when ignoring edge directions).

The bound above is derived from a more fine-grained bound: assuming that the longest path in the key graph is of length ℓ, the in-degree of every node is at most d and the challenge node can be reached from at most s sources (i.e., nodes with in-degree 0) we get

$$\epsilon' = \epsilon \cdot dn((2d+1)n)^{\lceil \log s \rceil} (3n)^{\lceil \log \ell \rceil} \ .$$

Note that ℓ, d and s are at most n and the previous bound was derived from this by setting $\ell = d = s = n$. Panjwani [Pan07] uses his bound to give a quasi-polynomial reduction of the
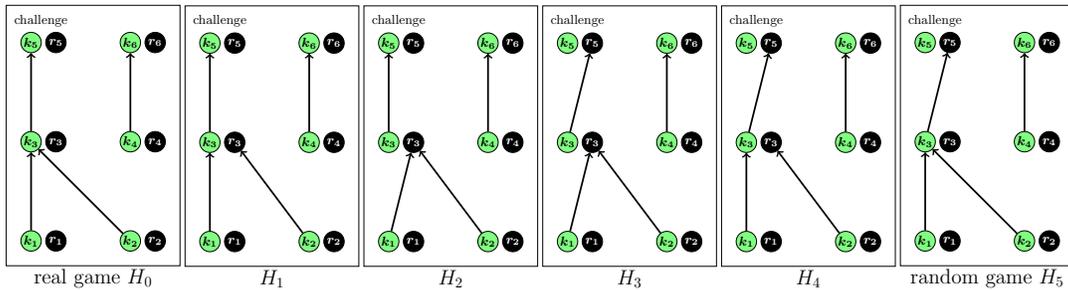

Fig. 1: Hybrids for the selective security proof. Green nodes correspond to keys, dark nodes are random values. The adversary A commits to encryption queries $(1,3), (2,3), (3,5)$ and challenge 5 (Encryption query $(4,6)$ is outside the connected component containing the challenge and thus not relevant for the hybrids. A could also corrupt keys 4 and 6, which are also outside.) Hybrid H_0 is the real game, hybrid H_5 is the random game, where instead of an encryption of the challenge key $\mathsf{Enc}_{k_3}(k_5)$, the adversary gets an encryption of the random value $\mathsf{Enc}_{k_3}(r_5)$. If an adversary A can distinguish any two consecutive hybrids H_i and H_{i+1} with some advantage δ, we can use A to construct B which breaks the IND-CPA security of Enc with the same advantage δ: E.g., assume B is given an IND-CPA challenge $C = \mathsf{Enc}_k(z)$ where z is one of two messages (which we call k_5 and r_5). Now B can simulate game H_2 for A, but when A makes the encryption query $(3,5)$, B answers with C. If $z = k_5$ then B simulates game H_2; but if $z = r_5$, it simulates game H_3. Note that B can simulate the games because k_3, which in the simulation is B's challenger's key, is not used anywhere else. Thus, B has the same advantage in the IND-CPA game as A has in distinguishing H_3 from H_4.

3

Logical Key Hierarchy (LKH) protocol [WGL00]. Panjwani first fixes a flaw in LKH, and calls the new protocol rLKH with "r" for repaired. rLKH is basically the GSD game restricted to a binary tree.[4]

The users correspond to the leaves of this tree, and their keys consists of all the nodes from the root to their leaf. Thus, if the tree is almost full and balanced, then it has only depth $\ell \approx \log n$ and Panjwani's bound loses only a quasi-polynomial factor $n^{\log(n)+2}$ (if $\ell = \log n$). As here $d = 2, \ell = \log n, s = n$, our bound gives a slightly worse bound $n^{\log(n)+\log\log(n)+4}$ for this particular problem, but this is only the case if a large fraction of the keys are actually used, and the adversary gets to see almost all of them. If ℓ is significantly larger than $\log n$ (e.g., because only few of the keys are active, or the tree is constructed in an unbalanced way like e.g. proposed in [SS00]), our bounds decrease only marginally, as opposed to exponentially fast in ℓ in [Pan07].

Graphs with small cut-width. The reason our result is restricted to trees is that in the process of generating the hybrids, we have to guess nodes such that removing this node splits the tree in a "nice" way (this has to be done $\log n$ times, losing a factor n in the distinguishing advantage every time).

One can generalize this technique (but we do not work out the details in this paper) to graphs with small "cut-width", where we say that a graph has cut-width w if for any two vertices u, v that are not connected by an edge, there exists a set of at most w vertices such that removing those disconnects u from v (a tree has cut-width $w = 1$). For graphs with cut-width w we get

$$\epsilon' = \epsilon \cdot n^{(2w+1)\log(n)+4} \ ,$$

which is subexponential in n, and thus beats the existing exponential bound whenever $w = o(n/\log^2(n))$. Whether there exists a subexponential reduction which works for any graph is an intriguing open problem.

Shorter keys from better reduction. An exponential security loss (as via complexity leveraging) means that, even when assuming exponential hardness of Enc (which is a typical assumption for symmetric encryption schemes like AES), one needs to use keys for Enc whose length is at least linear in n to get any security guarantee for the hardness of GSD at all. Whereas our bound for trees means that a key of length $polylog(n)$ is sufficient to get asymptotically overwhelming security (again assuming Enc is exponentially hard).

Nested hybrids. In a classical paper [GGM86] Goldreich, Goldwasser and Micali constructed a pseudorandom function (PRF) from a pseudorandom generator (PRG). More recently, three papers independently [BW13,KPTZ13,BGI14] observed that this construction is also a so-called *constrained* PRF, where for every string x one can compute a constrained key k_x that allows evaluation of the PRF on all inputs with prefix x. Informally, the security requirement is that an adversary that can ask for constrained keys cannot distinguish the output of the PRF on some challenge input from random.

All three papers [BW13,KPTZ13,BGI14] only prove *selective* security of this constrained PRF, where before any queries the adversary must commit to the input on which it wants to be challenged. This proof is a hybrid argument losing a factor $2m$ in the distinguishing advantage, where m is the PRF input length. One can then get adaptive security losing a huge exponential factor 2^m via complexity leveraging. Subsequently, Fuchsbauer et al. [FKPR14] gave a reduction that only loses a quasi-polynomial factor $(3q)^{\log m}$, where q denotes the number of queries made by the adversary. Our proofs borrows ideas from their work.

Very informally, the idea behind their proof is the following. In the standard proof for adaptive security using leveraging one first guesses the challenge query (losing a huge factor

---

[4] Let us stress that the graph obtained when just adding an edge for every encryption query in rLKH is not a tree after a rekeying operation. But for every node v, the subgraph we get when only keeping the nodes from which v can be reached is a tree, and as explained above, this is sufficient.
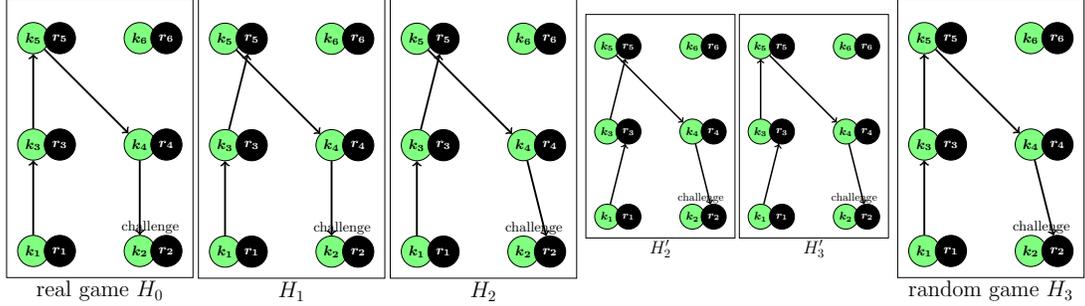
Fig. 2: Illustration of our adaptive security proof for paths.

2^m), which basically turns the adaptive attacker into a selective one, followed by a simple hybrid argument (losing a small factor $2m$) to prove selective security. The proof from [FKPR14] also first makes a guessing step, but a much simpler one, namely which of the q queries made by the adversary is the first to coincide with the challenge query on the first $m/2$ bits. This is followed by a hybrid argument losing a factor 3, so both steps together lose a factor $3q$. At this point the reduction is not finished yet, but intuitively the problem was reduced to itself but on inputs of only half the size $m/2$. These two steps can be iterated $\log m$ times (losing a total factor of $(3q)^{\log m}$) to get a reduction to the security of the underlying PRG.

Proof outline for paths. Our proof for GSD uses an approach similar to the one just explained, iterating fairly simple guessing steps with hybrid arguments, but the analogy ends here, as the actual steps are very different.

We first outline the proof for the adaptive security of the GSD game for a special case where the adversary is restricted in the sense that the connected component in the key graph containing the challenge must be a path. Even for this very special case, currently the best reduction [Pan07] loses an *exponential* factor $2^{\Theta(n)}$. We will now outline a reduction losing only a *quasi-polynomial* $n^{\log n}$ factor.[5] Recall that the standard way to prove adaptive security is to first guess the entire connected component containing the challenge, and then prove selective security as illustrated in Fig. 1.

Our approach is not to guess the entire path, but in a first step only the node in the middle of the path (as we make a uniform guess, it will be correct with probability $1/n$). This reduces the adaptive security game to a "slightly selective" game where the adversary must commit initially to this middle node, at the price of losing a factor n in the distinguishing advantage.[6]

Let H_0 and H_3 denote these "slightly selective" real and random GSD games (we also assume that the adversary initially commits to the challenge query, which costs another factor

---

[5] Let us mention that it is trivial to prove security of GSD restricted to paths if we additionally assume that for random keys k, k' the ciphertext $\mathsf{Enc}_k(k')$ is uniform given k' (this is e.g. the case for one-time pad encryption $\mathsf{Enc}_k(k') = k \oplus k'$): then the real and random challenge have the same distribution (they're uniform) and thus even a computationally unbounded adversary has zero advantage. (This is because in the path case, every key is used only once to encrypt.) The proof we outline here does not require this special property of Enc, and this will be crucial to later generalize it to more interesting graphs.

[6] We never actually construct this "slightly selective" adversary, but (as in complexity leveraging) we simply commit to a random guess, then run the adaptive adversary, and if its queries are not consistent with our guess, we abort outputting a random value. (We could also output a constant value; the point is that the advantage of the adversary, conditioned on our guess being wrong, is zero; whereas, conditioned on the guess being correct, it is the same as the advantage of the adaptive adversary). However, instead of this experiment it is easier to follow our proof outline by thinking of the adversary actually committing to its choices initially, but the reduction paying a factor (in the distinguishing advantage of the adversary that is allowed to make this choice adaptively) that corresponds to the size of the sample space of this guess.

of n). We illustrate this with a small example featuring a path of length 4 in Fig. 2. The correct guess for the middle node for the particular run of the experiment illustrated in the figure is $i = 5$. As now we know the middle vertex is $i = 5$, we can define new games H_1 and H_2 which are derived from H_0 and H_3, respectively, by replacing the ciphertext $\mathsf{Enc}_{k_j}(k_i)$ with an encryption $\mathsf{Enc}_{k_j}(r_i)$ of a random value (in the figure this is illustrated by replacing the edge $k_j \to k_i$ with $k_j \to r_i$).

So, what have we gained? If our adaptive adversary has advantage ϵ in distinguishing the real and random games then she has advantage at least ϵ/n to distinguish the "slightly selective" real and random games H_0 and H_3, and thus for some $i \in \{0, 1, 2\}$ she can distinguish the games H_i and H_{i+1} with advantage $\epsilon/3n$. Looking at two consecutive games H_i and H_{i+1}, we see that they only differ in one edge (e.g., in H_2 we answer the query $(3, 5)$ with $\mathsf{Enc}_{k_3}(r_5)$, in H_3 with $\mathsf{Enc}_{k_3}(k_5)$), and moreover this edge will be at the end of a path that now has only length 2, that is, half the length of the path in our original real and random games.

We can now continue this process, constructing new games where the path length is halved, paying a factor $3n$ in distinguishing advantage. For example, as illustrated in Fig. 2, we can guess the node that halves the path leading to the differing query in games H_2 and H_3 (for the illustrated path this would be $i = 3$), then define new games where we assume the adversary commits to this node (paying a factor n), and then define two new games H_2' and H_3', which are derived from games H_2 and H_3 (which now are augmented by our new guess), respectively, by answering the query (j, i) that asks for an encryption of this node (in the figure $(j, i) = (1, 3)$) with an encryption $\mathsf{Enc}_{k_1}(r_3)$ instead of $\mathsf{Enc}_{k_1}(k_3)$.

If we start with a path of length $\ell \leq n$ then after $\log \ell \leq \log n$ iterations of this process we proved the existence of two consecutive games (call them G_0 and G_1) that differ only in a single edge $j \to i$ and the vertex j has in-degree 0. That is, both games are identical, except that in one game the encryption query (j, i) is answered with $\mathsf{Enc}_{k_j}(k_i)$ and in the other with $\mathsf{Enc}_{k_j}(r_i)$. Moreover, the key k_j is not used anywhere else in the experiment and we know exactly when this query is made during the experiment (as the adversary committed to i).

Given a distinguisher A for G_0 and G_1, we can now construct an attacker B that breaks the IND-CPA security of the underlying encryption scheme with the same advantage: in the IND-CPA game B chooses two random messages m_0, m_1 and asks to be challenged on them.[7] The game samples a random bit b and returns the challenge $C = \mathsf{Enc}_k(m_b)$ to B, which must then output a guess b' for b. At this point, B invokes A and simulates the game G_0 for it, choosing all keys at random, except that it uses C to answer the encryption query (j, i).[8] Finally, B forwards A's guess b'. Identifying (k, m_0, m_1) with (k_j, k_i, r_i), we see that depending on whether $b = 0$ or $b = 1$, B simulates either G_0 or G_1. Thus, whatever advantage A has in distinguishing G_0 from G_1, B will break the IND-CPA security of Enc with the same advantage.

Proof outline for trees. We will now outline our reduction of the adaptive security of GSD to the IND-CPA security of Enc for a more general case. Namely, the adversary is only restricted in that the key graph resulting from its queries is such that the connected component containing the challenge is a tree. (Recall that we already disallowed cycles in the key graph as this would require circular security. Being a tree means that we also have no cycles in the key graph when ignoring edge directions). Note that paths as discussed in the previous section are very special trees. The GSD problem on trees is particularly interesting, as it captures some multicast encryption protocols like the Logical Key Hierarchy (LKH) protocol [WGL00]. We refer the reader to [Pan07] for details.

TREES WITH IN-DEGREES ≤ 1. Let us first consider the case where the connected component containing the challenge is a tree, and moreover all its vertices have in-degree 0 or 1. It turns out that the proof outlined for paths goes through with only minor changes for such trees. Note that such a tree has exactly one vertex with in-degree 0, which we call the *root*, and there is

---

[7] Note that B makes no encryption queries at all (which are allowed by the IND-CPA experiment).

[8] Note that since node j has in-degree 0, we can identify k_j with the key k used by the IND-CPA experiment, as we never have to encrypt k_j.
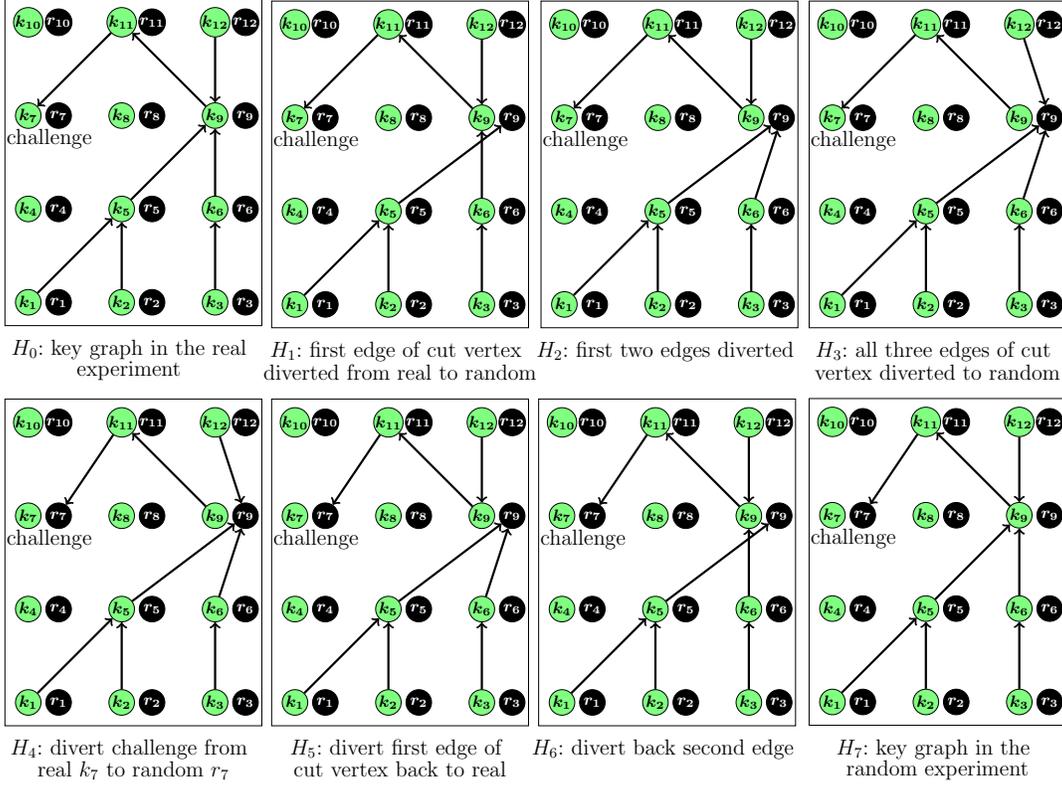
Fig. 3: Illustration of our adaptive security proof for general trees.

a unique path from the root to the challenge node. We can basically ignore all the edges not on this path and do a reduction as the one outlined above. The only difference is that now, when simulating the game G_b (where b is 0 or 1 depending on the whether the challenge C with which we answer the encryption query (j, i) is $\mathsf{Enc}_{k_j}(k_i)$ or $\mathsf{Enc}_{k_j}(r_i)$), the adversary can also ask for encryption queries (j, x) for any x. This might seem like a problem as we do not know k_j (we identified k_j with the key used by the IND-CPA challenger). But recall that in the IND-CPA game there is an encryption oracle $\mathsf{Enc}_{k_j}(\cdot)$, which we can query for the answer $\mathsf{Enc}_{k_j}(k_x)$ to such encryption queries.

GENERAL TREES. For general trees, where nodes can have in-degree greater than 1, we need to work more. The proof for paths does not directly generalize, as now nodes (in particular, the challenge) can be reached from more than one node with in-degree 0. We call these the *sources* of this node; for example in the tree H_0 in Fig. 3, the (challenge) node k_7 has 4 sources k_1, k_2, k_3 and k_{12}.

On a high level, our proof strategy will be to start with a tree where the challenge node c has s sources (more precisely, we have two games that differ in one edge that points to k_i in one game, and to r_i in the other, like games H_0 and H_7 in Fig. 3). We then guess a node v that "splits" the tree in a nice way, by which we mean the following: Assume v has in-degree d and we divert every edge going into v to a freshly generated node; let's call them $v_1, \ldots, v_d$. Then this splits the tree into a forest consisting of $d + 1$ trees (the component containing the challenge and one component for every v_i). The node v "well-divides" the tree if after the split the node c and all of $v_1, \ldots, v_d$ have at most $\lceil s/2 \rceil$ sources.

As an example, consider again the tree H_0 in Fig. 3, where the challenge node k_7 has 4 sources. The node k_9 would be a good guess, as it well-divides the tree: consider the forest after splitting at this node as described above (creating new nodes v_1, v_2, v_3 and diverting the edges

going into k_9 to them, i.e., replacing $k_5 \rightarrow k_9$ by $k_5 \rightarrow v_1$, $k_6 \rightarrow k_9$ by $k_6 \rightarrow v_2$, and $k_{12} \rightarrow k_9$ by $k_{12} \rightarrow v_3$). Then we obtain 4 trees, where now $c = k_7$ has only one source (k_9) and the new nodes v_1, v_2, v_3 have $2, 1$ and 1 sources, respectively.

Once we have guessed a well-dividing node v (or equivalently, the adversary has committed to such a node), we define $2d$ hybrid games (where d is the degree of the well-dividing node) between the two initial games, which we call H_0 and H_{2d+1}, as follows. H_1 is derived from H_0 by diverting the first encryption query that asks for an encryption of v (i.e., that is of the form (j, v) for some j) from real to random; that is, we answer with $\mathsf{Enc}_{k_j}(r_v)$ instead of $\mathsf{Enc}_{k_j}(k_v)$. For $i \leq d$, H_i is derived from H_0 by diverting the first i encryption queries. H_{d+1} is derived from H_d by diverting the encryption query that asks for an encryption of the challenge c from real to random. The final $d - 1$ hybrids games are used to switch the encryption of v back from random to real, one edge at a time. This process is illustrated in the games H_0 to H_7 in Fig. 3.

Because v was well-dividing (and we show in Lemma 2 that such a node always exists), we can prove the following property for any two consecutive games H_i and H_{i+1}: they differ in exactly one edge, which for some j, v in one game is $k_j \rightarrow k_v$ and $k_j \rightarrow r_v$ in the other, and moreover, k_j has at most $\lceil s/2 \rceil$ sources.

If an adversary can distinguish H_0 and H_{2d+1} with advantage ϵ then it must distinguish two hybrids H_i and H_{i+1} with advantage $\epsilon/((2d+1)n)$ (where n accounts for guessing the well-dividing node). But any such two hybrids now only have at most $\lceil s/2 \rceil$ sources. If we repeat this guessing/hybrid steps $\log s$ times, we end up with two games G_0 and G_1 which differ in one edge that has only one source. At this point we can then use our reduction for trees with only one source outlined above.

ANALYZING THE SECURITY LOSS. To halve the number of sources, we guess a well-dividing vertex (which costs a factor n in the reduction), and then must add up to $2d$ intermediate hybrids (where d is the maximum in-degree of any node), costing another factor $2d + 1$. Assuming that the number of sources is bounded by s, we have to iterate the process at most $\log s$ times. Finally, we lose another factor d (but only once) because our final node can have more than one ingoing edge. Overall, assuming the adversary breaks the GSD game with advantage ϵ on trees with at most s sources and in-degree at most d, our reduction yields an attacker against the IND-CPA security of Enc with advantage

$$\epsilon/\, dn((2d+1)n)^{\lceil \log s \rceil} (3n)^{\lceil \log \ell \rceil} \ .$$

For general trees, since $s, d \leq n$, we have $\epsilon/\, n^{3 \log n + 5}$.

2 Preliminaries

For $a \in \mathbb{N}$, we let $[a] = \{1, 2, \ldots, a\}$ and $[a]_0 = [a] \cup \{0\}$. We say adversary (or distinguisher) D is t-bounded if D runs in time t.

Definition 1. (Indistinguishability) *Two distributions X and Y are (ϵ, t)-indistinguishable, denoted $Y \sim_{(\epsilon, t)} X$ or $\Delta_t(Y, X) \leq \epsilon$, if no t-bounded distinguisher D can distinguish them with advantage greater than ϵ, i.e.,*

$$\Delta_t(Y, X) \leq \epsilon \iff \forall \mathsf{D}_t : \big| \Pr[\mathsf{D}_t(X) = 1] - \Pr[\mathsf{D}_t(Y) = 1] \big| \leq \epsilon \ .$$

Symmetric encryption. A pair of algorithms $(\mathsf{Enc}, \mathsf{Dec})$ with input $k \in \{0,1\}^\lambda$, where λ is the security parameter, and a message m (or a ciphertext) from $\{0,1\}^*$ is a symmetric-key encryption scheme if for all k, m we have $\mathsf{Dec}_k(\mathsf{Enc}_k(m)) = m$. Consider the game $\mathbf{Exp}_{\mathsf{Enc}, \mathsf{D}}^{\text{IND-CPA-}b}$ between a challenger C and a distinguisher D: C chooses a uniformly random key $k \in \{0,1\}^\lambda$ and a bit $b \in \{0,1\}$; D can make encryption queries for messages m and receives $\mathsf{Enc}_k(m)$; finally, D outputs a pair (m_0, m_1), is given $\mathsf{Enc}_k(m_b)$ and outputs a bit $b' \in \{0,1\}$, which is also the output of $\mathbf{Exp}_{\mathsf{Enc}, \mathsf{D}}^{\text{IND-CPA-}b}$.[9]

---

[9] For this notion to be satisfied, Enc must be probabilistic. In this paper one may also consider deterministic encryption, in which case the security definition must explicitly require that the challenge messages are fresh in the sense that D has not asked for encryptions of them already.

Definition 2. *Let $t \in \mathbb{N}^+$ and $0 < \epsilon < 1$. An encryption scheme $(\mathsf{Enc}, \mathsf{Dec})$ is (t, ϵ)-IND-CPA secure if for any t-bounded distinguisher* D*, we have*

$$\left| \Pr\left[\mathbf{Exp}_{\mathsf{Enc}, \mathsf{D}}^{\text{IND-CPA-1}} = 1 \right] - \Pr\left[\mathbf{Exp}_{\mathsf{Enc}, \mathsf{D}}^{\text{IND-CPA-0}} = 1 \right] \right| \leq \epsilon \ .$$

3   The GSD Game

In this section we describe the generalized selective decryption game as defined in [Pan07] and give our main theorem. Consider the following game, $\mathbf{Exp}_{\mathsf{Enc}, \mathsf{A}}^{\text{GSD-}(n, b)}$ called the generalized selective decryption (GSD) game, parameterized by an encryption scheme Enc,[10] an integer n and a bit b. It is played by the adversary A and the challenger B. First B samples n keys $k_1, k_2, \ldots, k_n$ uniformly at random from $\{0,1\}^\lambda$. A can make three types of queries during the game:

- encrypt: A query of the form $\mathsf{encrypt}(i, j)$ is answered with $c \leftarrow \mathsf{Enc}_{k_i}(k_j)$.
- corrupt: A query of the form $\mathsf{corrupt}(i)$ is answered with k_i.
- challenge: The response to $\mathsf{challenge}(i)$ depends on the bit b: if $b = 0$, the answer is k_i; if $b = 1$, the answer is a random value $r_i \in \{0,1\}^\lambda$.

A can make multiple queries of each type, adaptively and in any order. It can also make several challenge queries at any point in the in the game. Allowing multiple challenge queries models the fact that the respective keys are jointly pseudorandom (as opposed to individual keys being pseudorandom by themselves). Allowing to interleave challenges with other queries models that they remain pseudorandom even after corrupting more keys or seeing further ciphertexts.

We can think of the n keys that B creates as n vertices, labeled $1, 2, \ldots, n$, in a graph. In the beginning of the game there are no edges, but every time A queries $\mathsf{encrypt}(i, j)$, we add the edge $i \rightarrow j$ to the graph. When A queries $\mathsf{corrupt}(i)$ for some $i \in [n]$, we mark i as a corrupt vertex; when A queries $\mathsf{challenge}(i)$, we mark it as a challenge vertex. For an adversary A we call this graph the *key graph*, denoted $G(\mathsf{A})$ and we write $V^{\text{corr}}(\mathsf{A})$ and $V^{\text{chal}}(\mathsf{A})$ for the sets of corrupt and challenge nodes, respectively. (Note that $G(\mathsf{A})$ is a random variable depending on the randomness used by A and its challenger.)

Legitimate adversaries.   Consider an adversary that corrupts a node i in $G(\mathsf{A})$ and queries $\mathsf{challenge}(j)$ for some j which is reachable from i. Then A can successively decrypt the keys on the path from i to j, in particular k_j, and thus deduce the bit b. We only consider non-trivial breaks and require that no challenge node is reachable from a corrupt node in $G(\mathsf{A})$.

Two more restrictions must be imposed on $G(\mathsf{A})$ if we only want to assume that Enc satisfies IND-CPA. First, we do not allow key cycles, that is, queries yielding

$$\mathsf{Enc}_{k_{i_1}}(k_{i_2}), \mathsf{Enc}_{k_{i_2}}(k_{i_3}), \ldots, \mathsf{Enc}_{k_{i_{s-1}}}(k_{i_s}), \mathsf{Enc}_{k_s}(k_{i_1}) \ ,$$

as this would require the scheme to satisfy key-dependent-message (a.k.a. circular) security [BRS03,CL01].

Second, IND-CPA security does not imply that keys under which one has seen encryptions of random messages remain pseudorandom.[11] Pseudorandomness of keys (assuming only IND-CPA security of the underlying scheme) can thus only hold if their corresponding node does not have any outgoing edges. We thus require that all challenge nodes in the key graph are sinks (i.e., their out-degree is 0). The requirements (as formalized also in [Pan07]) are summarized in the following.

---

[10] We will never actually use the decryption algorithm Dec in the game, and thus will not mention it explicitly.

[11] Consider any IND-CPA-secure scheme $(\mathsf{Enc}, \mathsf{Dec})$ and define a new scheme as follows: keys are doubled in length and encryption under $k = k_1 || k_2$ is defined as $\mathsf{Enc}_k(m) = \mathsf{Enc}_{k_1}(m) || k_2$. This scheme is still IND-CPA, but given a ciphertext $C = \mathsf{Enc}_k(m)$ one can easily distinguish k from a random value even if m is random and unknown.

Definition 3. *An adversary* A *is* legitimate *if in any execution of* A *in the GSD game the values of* $G(\mathsf{A})$, $V^{corr}(\mathsf{A})$ *and* $V^{chal}(\mathsf{A})$ *are such that:*

- *For all* $i \in V^{corr}(\mathsf{A})$ *and* $j \in V^{chal}(\mathsf{A})$: *$j$ is unreachable from i in $G(\mathsf{A})$.*
- *$G(\mathsf{A})$ is a directed acyclic graph (DAG) and every node in $V^{chal}(\mathsf{A})$ is a sink.*

Let $n \in \mathbb{N}^+$ and $\mathcal{G}$ be a class of DAGs with n vertices. We say that a legitimate adversary A *is a $\mathcal{G}$-adversary if in any execution the key graph belongs to $\mathcal{G}$, i.e., $G(\mathsf{A}) \in \mathcal{G}$.*

Definition 4. *Let $t \in \mathbb{N}^+$, $0 < \epsilon < 1$. An encryption scheme* Enc *is called $(n, t, \epsilon, \mathcal{G})$-GSD secure if for every $\mathcal{G}$-adversary* A *running in time t, we have*

$$\left| \Pr\left[\mathbf{Exp}_{\mathsf{Enc, A}}^{\text{GSD-}(n,\,1)} = 1\right] - \Pr\left[\mathbf{Exp}_{\mathsf{Enc, A}}^{\text{GSD-}(n,\,0)} = 1\right] \right| \leq \epsilon \ .$$

Assuming one challenge query is enough. Although the definition of GSD allows the adversary to make any number of corruption queries, Panjwani [Pan07] observes that by a standard hybrid argument one can turn any adversary with advantage ϵ (which makes at most $q \leq n$ challenge queries) into an adversary that makes only one challenge query, but still has advantage at least ϵ/q. From now on we therefore only consider adversaries that make exactly one challenge query (keeping in mind that we have to pay an extra factor n in the final distinguishing advantage for statements about general adversaries).

4 Single Source

In this section we will analyze the GSD game for key graphs in which the challenge node is only reachable from one source node. That is, for some $q \leq n$ there is a path $p_1 \to p_2 \to \ldots \to p_q$ where p_1 has in-degree 0, all nodes p_i, $2 \leq i \leq q$ have in-degree 1 (but arbitrary out-degree) and the (single) challenge query is $\mathsf{challenge}(p_q)$ (recall that the challenge has out-degree 0). Let $\mathcal{G}_1$ be the set of all such graphs, and $\mathcal{G}_1^\ell \subseteq \mathcal{G}_1$ be the subset where this path has length at most ℓ.

Theorem 1 (GSD on trees with one path to challenge). *Let $t \in \mathbb{N}$, $0 < \epsilon < 1$ and $\mathcal{G}_1$ be the class of key graphs just defined. If an encryption scheme is (t, ϵ)-IND-CPA secure then it is also $(n, t', \epsilon', \mathcal{G}_1)$-GSD secure for*

$$\epsilon' = \epsilon \cdot n \, (3n)^{\lceil \log n \rceil} \qquad and \qquad t' = t - Q_{\mathsf{Adv}} T_{\mathsf{Enc}} - \tilde{O}(Q_{\mathsf{Adv}}) \ ,$$

where T_{Enc} denotes the time required to encrypt a key, and Q_{Adv} denotes an upper bound on the number of queries made by the adversary.[12] More generally, if we replace $\mathcal{G}_1$ with $\mathcal{G}_1^\ell$, we get

$$\epsilon' = \epsilon \cdot n \, (3n)^{\lceil \log \ell \rceil} \qquad and \qquad t' = t - Q_{\mathsf{Adv}} T_{\mathsf{Enc}} - \tilde{O}(Q_{\mathsf{Adv}}) \ .$$

GSD on single-source graphs. For $b \in \{0, 1\}$, we consider the GSD game $\mathbf{Exp}_{\mathsf{Enc}}^{\text{GSD-}(n,\,b)}$ on $\mathcal{G}_1$ between B and an adversary A. Challenger B first samples n random keys $k_1, k_2, \ldots, k_n$ and we assume that already at this point B samples fake keys $r_1, \ldots, r_n$. On all $\mathsf{encrypt}(i, j)$ queries B returns real responses $\mathsf{Enc}_{k_i}(k_j)$. If $b = 0$, the response to $\mathsf{challenge}(z)$ is k_z; if $b = 1$, the response is r_z.

We require that the key graph is in $\mathcal{G}_1$, that is the connected component of the key graph which contains the challenge z has a path $p_1 \to p_2 \to \ldots \to p_q = z$ with p_1 having in-degree 0, all other p_i having in-degree 1 and $p_q = z$ having out-degree 0 (this means A made queries $\mathsf{encrypt}(p_{i-1}, p_i)$, but no queries $\mathsf{encrypt}(x, p_i)$ for $x \neq p_{i-1}$).

---

[12] If Enc is deterministic then w.l.o.g. we can assume $Q_{\mathsf{Adv}} \leq n^2$ as there are at most $n(n-1)/2$ possible encryption queries (plus $\leq n$ corruption and challenge queries). If Enc is probabilistic then A is allowed any number of encryption queries.

Eventually, A outputs a bit $b' \in \{0, 1\}$, which is also the output of the game. If the encryption scheme Enc is not $(t', \epsilon', \mathcal{G}_1)$-GSD secure then there exists a $\mathcal{G}_1$-adversary A running in time t' such that

$$\left| \Pr\left[\mathbf{Exp}_{\mathsf{Enc, A}}^{\text{GSD-}(n, 0)} = 1 \right] - \Pr\left[\mathbf{Exp}_{\mathsf{Enc, A}}^{\text{GSD-}(n, 1)} = 1 \right] \right| > \epsilon' \ . \tag{1}$$

Our goal. Suppose we knew that our GSD adversary A wants to be challenged on a fixed node z^* and that it will make a query $\mathsf{encrypt}(y, z^*)$ for some y which it will not use in any other query. Then we could use A directly to construct a distinguisher D as in Definition 2: D sets up all keys k_x, $x \in [n]$, samples a value r_{z^*} and runs A, answering A's queries using its keys; except when $\mathsf{encrypt}(y, z^*)$ is queried for any $y \in [q]$, D queries its own challenger on (k_{z^*}, r_{z^*}) and forwards the answer to A. Moreover, $\mathsf{challenge}(z^*)$ is answered with k_{z^*}. If D's challenger, C, chose $b = 0$, this perfectly simulates the real game for A. If $b = 1$ then A gets an encryption of r_{z^*} and the challenge query is answered with k_{z^*}, although in the random GSD game A expects an encryption of k_{z^*} and $\mathsf{challenge}(z^*)$ to be answered with r_{z^*}. However, these two games are distributed identically, since both k_{z^*} and r_{z^*} are uniformly random values that do not occur anywhere else in the game. Thus D simulates the real game when $b = 0$ and the random game when $b = 1$. Note that D implicitly set k_y to the key that C chose, but that's fine, since we assumed that k_y is not used anywhere else in the game and thus not needed by D for the simulation.

Finally, suppose that, in addition to the challenge z^*, we knew y^* for which A will query $\mathsf{encrypt}(y^*, z^*)$. Then we could also allow A to issue queries of the form $\mathsf{encrypt}(y^*, x)$, for x other than z^*. D could easily simulate any such query by querying k_x to its encryption oracle.

Unfortunately, general GSD adversaries can decide adaptively on which node they want to be challenged, and worse, they can make queries $\mathsf{encrypt}(x, y)$, where y is a key that encrypts the challenge.

We will construct a series of hybrids where any two consecutive games **Game** and **Game**$'$ are such that from a distinguisher A for them, we can construct an adversary D against the encryption scheme with the same advantage. For this, the two games should only differ in the response of one encryption query on the path to the challenge, say $\mathsf{encrypt}(y, z)$, which is responded to with a real ciphertext $\mathsf{Enc}_{k_y}(k_z)$ in **Game** and with a fake ciphertext $\mathsf{Enc}_{k_y}(r_z)$ in **Game**$'$.

Moreover, the key k_y must not be encrypted anywhere else in the game, as our distinguisher D will implicitly set k_y to be the key of its IND-CPA challenger C. Thus, in **Game** and **Game**$'$ all queries $\mathsf{encrypt}(x, y)$, for any x, are responded to with a fake ciphertext $\mathsf{Enc}_{k_x}(r_y)$. Summing up, we need the two games to have the following properties for some y:

- Property 1. **Game** and **Game**$'$ are identical except for the response to one query $\mathsf{encrypt}(y, z)$, which is replied to with a real ciphertext in **Game** and a fake one in **Game**$'$.
- Property 2. Queries $\mathsf{encrypt}(x, y)$ are replied to with a fake response in both games.

If we knew the entire key graph $G(\mathsf{A})$ before answering A's queries then we could define a series of $2q - 1$ games as in Fig. 1 where we consecutively replace edges from the source to the challenge by fake nodes and then go back replacing fake edges with real ones starting with $p_{q-2} \to p_{q-1}$. Any two consecutive games in such a sequence would satisfy the two properties, so we could use them to break IND-CPA.

The problem is that in general the probability of guessing the connected component containing the challenge is exponentially small in n and consequently from a GSD adversary's advantage ϵ' we will obtain a distinguisher D with advantage $\epsilon = \epsilon'/O(n!)$. To avoid an exponential loss, we thus must avoid guessing the entire component at once.

The first step. Our first step is to define two new games $\mathsf{Game}_{\emptyset}^{\{q\}}$ and $\mathsf{Game}_{\{q\}}^{\{q\}}$, which are modifications of $\mathbf{Exp}^{\text{GSD-0}}$ and $\mathbf{Exp}^{\text{GSD-1}}$, respectively. Both new games have an extra step at the beginning of the game: B guesses which key is going to be the challenge key and at the end of the game only if its guess was correct, the output of the game is A's output and

otherwise it is 0. Clearly B's guess is correct with probability $^1/_n$. Aside from this guessing step, $\mathsf{Game}_\emptyset^{\{q\}}$ is identical to $\mathbf{Exp}^{\text{GSD-0}}$; all responses are real. We therefore have $\Pr[\mathsf{Game}_\emptyset^{\{q\}} = 1] = {}^1/_n \cdot \Pr[\mathbf{Exp}^{\text{GSD-0}} = 1]$.

Analogously, we define an auxiliary game, $\mathsf{Game}_1^{\{q\}}$, which is identical to $\mathbf{Exp}^{\text{GSD-1}}$, except for the guessing step. Again we have $\Pr[\mathsf{Game}_1^{\{q\}} = 1] = {}^1/_n \cdot \Pr[\mathbf{Exp}^{\text{GSD-1}} = 1]$. We then define $\mathsf{Game}_{\{q\}}^{\{q\}}$ exactly as $\mathsf{Game}_1^{\{q\}}$, except for a syntactical change: Let z be the guessed value for the challenge node. Then any query $\mathsf{encrypt}(x, z)$ is replied to with $\mathsf{Enc}_{k_x}(r_z)$, that is, an encryption of the *fake* key r_z. (Note that this game can be simulated, since we "know" z when guessing correctly.) On the other hand, the query $\mathsf{challenge}(z)$ is answered with k_z (rather than r_z in $\mathbf{Exp}^{\text{GSD-1}}$). Since the difference between $\mathsf{Game}_1^{\{q\}}$ and $\mathsf{Game}_{\{q\}}^{\{q\}}$ is that we have replaced all occurrences of k_z by r_z and all occurrences of r_z by k_z, which are distributed identically (thus we've merely swapped the names of k_z and r_z), we have $\Pr[\mathsf{Game}_{\{q\}}^{\{q\}} = 1] = \Pr[\mathsf{Game}_1^{\{q\}} = 1] = {}^1/_n \cdot \Pr[\mathbf{Exp}^{\text{GSD-1}} = 1]$.

Together with Equation (1), we have thus

$$\left| \Pr\left[\mathsf{Game}_\emptyset^{\{q\}} = 1\right] - \Pr\left[\mathsf{Game}_{\{q\}}^{\{q\}} = 1\right]\right|$$
$$= {}^1/_n \cdot \left| \Pr\left[\mathbf{Exp}^{\text{GSD-0}} = 1\right] - \Pr\left[\mathbf{Exp}^{\text{GSD-1}} = 1\right]\right| > {}^1/_n \cdot \epsilon' \ .$$

We continue to use the notational convention that for sets $I \subseteq P \subseteq [n]$, the game Game_I^P is derived from the real game by additionally guessing the nodes corresponding to P and answering encryptions of the nodes in I with fake keys. This is made formal in Fig. 4 below.

The second step. Assume q is a power of 2 and consider $\mathsf{Game}_\emptyset^{\{q/2, q\}}$, which is identical to $\mathsf{Game}_\emptyset^{\{q\}}$, except that in addition to the challenge node, B also guesses which node $x \in [n]$ is going to be the node in the middle of the path to the challenge, i.e. $p_{q/2} = x$. The output of $\mathsf{Game}_\emptyset^{\{q/2, q\}}$ is A's output if the guess was correct and 0 otherwise. Since B guesses correctly with probability $^1/_n$, we have

$$\Pr\left[\mathsf{Game}_\emptyset^{\{q/2, q\}} = 1\right] = {}^1/_n \cdot \Pr\left[\mathsf{Game}_\emptyset^{\{q/2\}} = 1\right] \ .$$

By guessing the middle node, we can assume the middle node is *known* and this will enable us to define a hybrid game, $\mathsf{Game}_{\{q/2\}}^{\{q/2, q\}}$, in which the query for the encryption of $k_{p_{q/2}}$ is responded to with a fake answer. In addition, we consider games $\mathsf{Game}_{\{q\}}^{\{q/2, q\}}$ and $\mathsf{Game}_{\{q/2, q\}}^{\{q/2, q\}}$ which are similarly defined by making the same changes to game $\mathsf{Game}_{\{q\}}^{\{q\}}$, i.e. guessing the middle node and replying to the encryption query of the guessed key with a fake and a real ciphertext respectively. Again, we have $\Pr[\mathsf{Game}_{\{q\}}^{\{q/2, q\}} = 1] = {}^1/_n \cdot \Pr[\mathsf{Game}_{\{q\}}^{\{q\}} = 1]$. Therefore $(t', \epsilon'/n)$-distinguishability of $\mathsf{Game}_\emptyset^{\{q\}}$ and $\mathsf{Game}_{\{q\}}^{\{q\}}$ implies that $\mathsf{Game}_\emptyset^{\{q/2, q\}}$ and $\mathsf{Game}_{\{q\}}^{\{q/2, q\}}$ are $(t', \epsilon'/n^2)$-distinguishable, i.e. $\Delta_t\left(\mathsf{Game}_\emptyset^{\{q/2, q\}}, \mathsf{Game}_{\{q\}}^{\{q/2, q\}}\right) > \epsilon'/n^2$, and therefore by the triangle inequality

$$\Delta_t\left(\mathsf{Game}_\emptyset^{\{q/2, q\}}, \mathsf{Game}_{\{q/2\}}^{\{q/2, q\}}\right) + \Delta_t\left(\mathsf{Game}_{\{q/2\}}^{\{q/2, q\}}, \mathsf{Game}_{\{q/2, q\}}^{\{q/2, q\}}\right)$$
$$+ \Delta_t\left(\mathsf{Game}_{\{q/2, q\}}^{\{q/2, q\}}, \mathsf{Game}_{\{q\}}^{\{q/2, q\}}\right) \geq \Delta_t\left(\mathsf{Game}_\emptyset^{\{q/2, q\}}, \mathsf{Game}_{\{q\}}^{\{q/2, q\}}\right)$$
$$> {}^1/_{n^2} \cdot \epsilon' \ . \tag{2}$$

By Line 2, at least one of the pairs of games on the left-hand side must be $(t', \epsilon'/3n^2)$-distinguishable. The two games of every pair differ in exactly one point, as determined by the subscript of each game. For instance, the difference between the last pair $\mathsf{Game}_{\{q/2, q\}}^{\{q/2, q\}}$ and $\mathsf{Game}_{\{q\}}^{\{q/2, q\}}$ is the encryption of node $q/2$.

Recall that our goal is to construct a pair of hybrids where the differing query $\mathsf{encrypt}(y, z)$ is such that all queries $\mathsf{encrypt}(x, y)$ are replied to with $\mathsf{Enc}_{k_x}(r_y)$, as formalized by Property 2. Games $\mathsf{Game}_\emptyset^{\{q\}}$ and $\mathsf{Game}_{\{q\}}^{\{q\}}$ differed in the last query on the path and the only key above it that is not encrypted anywhere is the start of the path. What we have achieved with our

Game_I^P, with $I \subseteq P \subseteq [n]$ is defined as follows:

- For every $i \in P$, B chooses $v_i \leftarrow [n]$, which is B's guess for the node at position i in the final path.
- B chooses $2n$ keys $k_1, r_1, k_2, r_2 \ldots, k_n, r_n \leftarrow \{0,1\}^\lambda$ and runs A.
- Whenever A makes a query $\mathsf{encrypt}(x, y)$, B does the following: If $y = v_i$ for some $i \in I$ then reply with $\mathsf{Enc}_{k_x}(r_{v_i})$; otherwise reply with $\mathsf{Enc}_{k_x}(k_y)$.
- When A makes the query $\mathsf{challenge}(z)$, return k_z.
- Let $b' \in \{0,1\}$ be A's output. At then end of the game, consider the longest path $p_0 \rightarrow p_1 \rightarrow \cdots \rightarrow p_q$ in $G(\mathsf{A})$, with p_q being the argument of A's $\mathsf{challenge}$ query. If for all $i \in P$: $v_i = p_i$ then B returns b'; otherwise, B returns 0.

Fig. 4: Definition of Game_I^P for the single-source case.

games above is to *halve* that distance: the first pair, $(\mathsf{Game}_\emptyset^{\{q/2,\, q\}}, \mathsf{Game}_{\{q/2\}}^{\{q/2,\, n\}})$, and the last pair, $(\mathsf{Game}_{\{q/2,\, q\}}^{\{q/2,\, q\}}, \mathsf{Game}_{\{q\}}^{\{q/2,\, q\}})$, differ in a node that is only half way down the path; and the middle pair, $(\mathsf{Game}_{\{q/2\}}^{\{q/2,\, q\}}, \mathsf{Game}_{\{q/2,\, q\}}^{\{q/2,\, q\}})$, differ in the last node, but half way up the path there is a key, namely $k_{q/2}$, which is not encrypted anywhere, as all queries $\mathsf{encrypt}(x, q/2)$ are answered with $\mathsf{Enc}_{k_x}(r_{q/2})$.

The remaining steps.   For any of the three pairs that is $(t', \epsilon'/3n^2)$-distinguishable (and by Line 2 there must exist one), we can repeat the same process on the half of the path which ends with the query that is different in the two games. For example, assume this holds for the last pair, that is

$$\Delta_t \left(\mathsf{Game}_{\{q/2,\, q\}}^{\{q/2,\, q\}}, \mathsf{Game}_{\{q\}}^{\{q/2,\, q\}} \right) \; > \; \frac{\epsilon'}{3n^2} \; . \tag{3}$$

We repeat the process of guessing the middle node between the differing node and the random node above (in this case the root of the path), which is thus node $q/4$, and obtain a new pair which satisfies

$$\Delta_t \left(\mathsf{Game}_{\{q/2,\, q\}}^{\{q/4,\, q/2,\, q\}}, \mathsf{Game}_{\{q\}}^{\{q/4,\, q/2,\, q\}} \right) \; > \; \frac{\epsilon'}{3n^3} \; , \tag{4}$$

by Line 3 and the fact that the guess is correct with probability $1/n$. We can now define two intermediate games

$$\mathsf{Game}_{\{q/4,\, q/2,\, q\}}^{\{q/4,\, q/2,\, q\}} \quad \text{and} \quad \mathsf{Game}_{\{q/4,\, q\}}^{\{q/4,\, q/2,\, q\}} \tag{5}$$

where we replaced the encryption of $k_{p_{q/4}}$ by one of $r_{p_{q/4}}$. As in Line 2, we can again define a sequence of games by putting the games in Line 5 between the ones in Line 4 and argue that by Line 4, two consecutive hybrids must be $(t', \epsilon'/(3^2 n^3))$-distinguishable. What we have gained is that any pair in this sequence differs by exactly one edge and the closest fake answer above is only a fourth of the path length away.

Repeating these two steps a maximum number of $\lceil \log q \rceil$ times, we arrive at two consecutive games, where the distance from the differing node to the closest "fake" node above is 1. We have thus found two games that satisfy Properties 1 and 2, meaning we can use a distinguisher A to construct an adversary D against the encryption scheme.

Since a path has at most n nodes, after at most $\log n$ steps we end up with two games that are $(t', \epsilon'/n(3n)^{\lceil \log n \rceil})$-distinguishable and which can be used to break the encryption scheme. If the adversary is restricted to paths of length ℓ (i.e., graphs in $\mathcal{G}_1^\ell$), this improves to $(t', \epsilon'/n(3n)^{\lceil \log \ell \rceil})$.

Proof of Theorem 1.   We formalize our method to give a proof of the theorem. In Fig. 4 we describe game Game_I^P, which is defined by the nodes on the path that are guessed (represented by the set P) and the nodes where an encryption of a key is replaced with an encryption of a value r (represented by $I \subseteq P$).

13

Lemma 1. *Let $I \subseteq P \subseteq [n]$ and $z \in P \setminus I$. Also let y be the largest number in I such that $y < z$, and $y = 0$ if z is smaller than all elements in I. If Game_I^P and $\mathsf{Game}_{I \cup \{z\}}^P$ are (t, ϵ)-distinguishable then the following holds.*

- *If $z = y + 1$ then Enc is not $(t + Q_{\mathsf{Adv}}T_{\mathsf{Enc}} + \tilde{O}(Q_{\mathsf{Adv}}))$, $\epsilon)$-IND-CPA-secure.*
- *If $z > y + 1$, define $z' = y + \lfloor (z - y)/2 \rfloor$, $P' = P \cup \{z'\}$ and*

$$ I_1 = I \ , \qquad I_2 = I \cup \{z'\} \ , \qquad I_3 = I \cup \{z', z\} \ , \qquad I_4 = I \cup \{z\} \ . $$

Then for some $i \in \{1, 2, 3\}$, games $\mathsf{Game}_{I_i}^{P'}$ and $\mathsf{Game}_{I_{i+1}}^{P'}$ are $(t, \epsilon/3n)$-distinguishable.

Applying Lemma 1 repeatedly $\lceil \log n \rceil$ times (or $\lceil \log \ell \rceil$ if we know an upper bound on the path length ℓ), we obtain the proof of Theorem 1.

Proof. [of Lemma 1] If $z = y + 1$ then the pair $(\mathsf{Game}_I^P, \mathsf{Game}_{I \cup \{z\}}^P)$ satisfies Properties 1 and 2. As we'll explain below, this allows us to construct distinguisher D that breaks the IND-CPA security of the encryption scheme with the same advantage, in time t that it takes to run A plus the time that it takes to simulate the game Game_I^P, which consists of at most Q_{Adv} encryption or challenge queries, each taking time T_{Enc}, plus some $\tilde{O}(Q_{\mathsf{Adv}})$ bookkeeping overhead. Thus we get a total running time of of $t + Q_{\mathsf{Adv}}T_{\mathsf{Enc}} + \tilde{O}(Q_{\mathsf{Adv}})$.

In both games the query $\mathsf{encrypt}(p_{y-1}, p_y)$ is answered with $\mathsf{Enc}_{k_{p_{y-1}}}(r_{p_y})$, meaning that k_{p_y} is never encrypted (as there cannot be other queries $\mathsf{encrypt}(x, p_y)$, since A is a $\mathcal{G}_1$ adversary). Moreover, note that p_y cannot be queried to $\mathsf{corrupt}$ either. The query $\mathsf{encrypt}(p_y, p_z)$ is responded to with a real answer $\mathsf{Enc}_{k_{p_y}}(k_{p_z})$ in Game_I^P and with a fake answer $\mathsf{Enc}_{k_{p_y}}(r_{p_z})$ in $\mathsf{Game}_{I \cup \{z\}}^P$.

We now let our distinguisher D against Enc simulate the game Game_I^P implicitly replacing k_{p_y} with its challenger's key (but sampling all remaining real and random keys itself). D answers queries $\mathsf{encrypt}(p_y, x)$, for $x \neq p_z$ by using the encryption oracle. When A queries $\mathsf{encrypt}(p_y, p_z)$, D outputs the challenge (k_{p_z}, r_{p_z}) and forwards the answer to A. Depending on D's challenger's bit, D either simulates Game_I^P or $\mathsf{Game}_{I \cup \{z\}}^P$, and thus has the same advantage as A.

For the case $z > y + 1$ we have that the node z' is roughly in the middle of y and z on the path. If Game_I^P and $\mathsf{Game}_{I \cup \{z\}}^P$ are (t, ϵ)-distinguishable then games $\mathsf{Game}_I^{P \cup \{z'\}}$ and $\mathsf{Game}_{I \cup z}^{P \cup \{z'\}}$ are $(t, \epsilon/n)$-distinguishable: the guess of $z' \leftarrow [n]$ does not influence A's behavior (it only influences B's output after A has stopped), and the probability of guessing correctly is $1/n$. We therefore have $\Pr[\mathsf{Game}_J^P = 1] = 1/n \cdot \Pr[\mathsf{Game}_J^{P \cup \{z'\}} = 1]$ for any $J \subseteq P$ and thus

$$ \Delta_t \left(\mathsf{Game}_I^{P \cup \{z'\}}, \mathsf{Game}_{I \cup \{z\}}^{P \cup \{z'\}} \right) = \frac{1}{n} \cdot \Delta_t \left(\mathsf{Game}_I^P, \mathsf{Game}_{I \cup \{z\}}^P \right) > \frac{\epsilon}{n} \ . $$

By the triangle inequality we have (recall that $P' = P \cup \{z'\}$):

$$ \Delta_t \left(\mathsf{Game}_I^{P'}, \mathsf{Game}_{I \cup \{z'\}}^{P'} \right) + \Delta_t \left(\mathsf{Game}_{I \cup \{z'\}}^{P'}, \mathsf{Game}_{I \cup \{z', z\}}^{P'} \right) $$
$$ + \Delta_t \left(\mathsf{Game}_{I \cup \{z', z\}}^{P'}, \mathsf{Game}_{I \cup \{z\}}^{P'} \right) \geq $$
$$ \Delta_t \left(\mathsf{Game}_I^{P'}, \mathsf{Game}_{I \cup \{z\}}^{P'} \right) > \frac{\epsilon}{n} \ . $$

Thus at least one of the pairs must be $(t, \epsilon/3n)$-distinguishable.

5   General Trees

For a node v in a directed graph G let T_v denote the subgraph of G we get when only keeping the edges on paths that lead to v. In this section we prove bounds for GSD if the underlying key graph is a tree. Concretely, let $\mathcal{G}_\tau$ be the class of key graphs that contain one designated "challenge node" z and where the graph T_z is a tree (when ignoring edge directions).

To give more fine-grained bounds we define a subset $\mathcal{G}_\tau^{s,d,\ell} \subseteq \mathcal{G}_\tau$ as follows. For $G \in \mathcal{G}_\tau$, let z be the challenge node and T_z as above. Then $G \in \mathcal{G}_\tau^{s,d,\ell}$ if the challenge node has at most s sources (i.e., there are at most s nodes u of in-degree 0 s.t. there is a directed path from u to z), every node in T_z has in-degree at most d and the longest path in T_z has length at most ℓ. Note that as $d < n, s < n$ and $\ell \leq n$ any $G \in \mathcal{G}_\tau$ with n nodes is trivially in $\mathcal{G}_\tau^{n-1,n-1,n}$.

Theorem 2 (Security of GSD on trees). *Let $n, t \in \mathbb{N}$, $0 < \epsilon < 1$ and $\mathcal{G}_\tau$ be the class of key graphs just defined. If an encryption scheme is (t, ϵ)-IND-CPA secure then it is also $(n, t', \epsilon', \mathcal{G}_\tau)$-GSD secure for*

$$\epsilon' = \epsilon \cdot n^2 (6n^3)^{\lceil \log n \rceil} \leq \epsilon \cdot n^{3\lceil \log n \rceil + 5} \qquad and \qquad t' = t - Q_{\mathsf{Adv}} T_{\mathsf{Enc}} - \tilde{O}(Q_{\mathsf{Adv}})$$

(with $Q_{\mathsf{Adv}}, T_{\mathsf{Enc}}$ as in Theorem 1). If we replace $\mathcal{G}_\tau$ with $\mathcal{G}_\tau^{s,d,\ell}$ then

$$\epsilon' = \epsilon \cdot dn((2d+1)n)^{\lceil \log s \rceil}(3n)^{\lceil \log \ell \rceil} \qquad and \qquad t' = t - Q_{\mathsf{Adv}} T_{\mathsf{Enc}} - \tilde{O}(Q_{\mathsf{Adv}}) \ .$$

5.1 Proof of Theorem 2

Notation. As all our graphs will have vertex set $[n]$, with "graph" we also refer to the set of edges of a graph. Thus for two graphs (set of edges) T, T', by $\mathsf{T} \setminus \mathsf{T}'$ we mean the graph T (with vertex set $[n]$) after removing all edges of T' and $\mathsf{T} \cup \mathsf{T}'$ is the graph of all edges of T and T'.

For a node $x \in [n]$ we denote by T_x the graph (set of edges) of all the paths in $G(\mathsf{A})$ that reach x. Recall that a node with in-degree 0 is called a source and denote by $S(v)$ the number of sources in the tree T_v. Note that $S(v)$ is also the number of paths to a node v. For a tree T with at least 2 sources we say a node v *well-divides* T if the number of sources in each subtree obtained from removing all edges of the form $* \to v$, is less than or equal to half of the sources in T. If $(x, y) \in \mathsf{T}$, we call x a parent of y and y a child of x. Furthermore, we let $x[i]$ denote the i-th parent of node x, which means that in the GSD game resulting in $G(\mathsf{A})$ there was a query $\mathsf{encrypt}(x[i], x)$ and before that there were $i - 1$ queries of the form $\mathsf{encrypt}(y, x)$ with $y \neq x[i]$.

Our approach. Generalizing the proof for single-source graphs from Section 4, we will define a sequence of hybrids such that there are two games **Game** and **Game**$'$ which have the following properties:

- the two games differ in exactly one query, $\mathsf{encrypt}(y, z)$;
- the node y has only one source, i.e., $S(y) = 1$.

Finding **Game** and **Game**$'$ will enable us to apply Theorem 1 to these two games and derive contradiction as we did before.

In the single-source game we guessed which node will be in the middle of the path to the challenge, as this allowed us to reduce the problem to smaller problems. Here we need a new metric since there are multiple paths. Whereas before we halved the *length* of the path to the challenge, we now first halve the *number* of such paths, by guessing a node v that well-divides the tree.

Let us look at the original games $\mathbf{Exp}^{\mathrm{GSD}\text{-}0}$ and $\mathbf{Exp}^{\mathrm{GSD}\text{-}1}$ and suppose they are δ-distinguishable. Unlike for the single-source case, these games differ not in one but in at most d edges: with z denoting the challenge, in $\mathbf{Exp}^{\mathrm{GSD}\text{-}0}$ all edges $* \to z$ are real, whereas they are fake in $\mathbf{Exp}^{\mathrm{GSD}\text{-}1}$ (and there are at most d of them). We start with defining two new games, where the only change is that we guess the challenge node. Other than that $\mathsf{Game}_{(0)}^{(0)}$ is defined as $\mathbf{Exp}^{\mathrm{GSD}\text{-}0}$ and $\mathsf{Game}_{(d)}^{(0)}$ as $\mathbf{Exp}^{\mathrm{GSD}\text{-}1}$ (except for the syntactical swap of k_z and r_z). Note that the subscript (d) of Game denotes that the "first d" (that is, all) queries of the form $\mathsf{encrypt}(*, z)$ are replied to with fake ciphertexts. Games $\mathsf{Game}_{(0)}^{(0)}$ and $\mathsf{Game}_{(d)}^{(0)}$ are thus (δ/n)-distinguishable.

We now define $d - 1$ intermediate games $\mathsf{Game}_{(i)}^{(0)}$, $1 \leq i \leq d - 1$, where in $\mathsf{Game}_{(i)}^{(0)}$ the first i responses to $\mathsf{encrypt}(*, z)$ are fake and the rest are real. For some $i \in [d]$, $\mathsf{Game}_{(i-1)}^{(0)}$ and

$\mathsf{Game}_{(i)}^{(0)}$ are thus $(\delta/(dn))$-distinguishable. What we have gained is that now in the sequence of hybrids, two games only differ in the response of one query.

Suppose that $\mathsf{Game}_{(j-1)}^{(0)}$ and $\mathsf{Game}_{(j)}^{(0)}$ are $(\delta/(dn))$-distinguishable and let z denote the challenge. The games differ in the edge $z[j] \to z$, which is real in $\mathsf{Game}_{(j-1)}^{(0)}$ and fake in $\mathsf{Game}_{(j)}^{(0)}$. The node $z[j]$ could itself still have s sources, so our next step is to reduce this. We thus define two new games where we guess a well-dividing node for $\mathsf{T}_{z[j]}$. We define $\mathsf{Game}_{(j-1)}^{(0,j)}$ as $\mathsf{Game}_{(j-1)}^{(0)}$, where in addition to the challenge, we guess the node v that well-divides $\mathsf{T}_{z[j]}$, that is the tree above the j-th edge that goes into z. We analogously define $\mathsf{Game}_{(j)}^{(0,j)}$ and have that $\mathsf{Game}_{(j-1)}^{(0,j)}$ and $\mathsf{Game}_{(j)}^{(0,j)}$ are $(\delta/(dn^2))$-distinguishable.

Having found a well-dividing node v, we can now define $2d$ intermediate hybrids, where we replace every edge $* \to v$, one by one, with a fake edge and call these hybrids $\mathsf{Game}_{(j-1,i)}^{(0,j)}$ for $i = 1, \ldots, d$. $\mathsf{Game}_{(j-1,i)}^{(0,j)}$ is thus defined as $\mathsf{Game}_{(j-1)}^{(0,j)}$, except that the first i answers to queries $\mathsf{encrypt}(*, v)$ are fake. In $\mathsf{Game}_{(j-1,d)}^{(0,j)}$, our splitting node v has thus only fake ingoing edges, which means that in the key graph we have removed $S(v)$ real sources from the challenge and added one source, namely v. The next hybrid is $\mathsf{Game}_{(j,d)}^{(0,j)}$ (where we changed the edge $z[j] \to z$ from real to random), and then a sequence of hybrids $\mathsf{Game}_{(j,i)}^{(0,j)}$ for $i = d, \ldots, 1$, which changes back the edges $* \to v$ from fake to real.

What have we gained by this? Since v well-divided the tree, each differing edge between two consecutive hybrids in the sequence

$$\mathsf{Game}_{(j-1)}^{(0,j)} = \mathsf{Game}_{(j-1,0)}^{(0,j)}, \ldots, \mathsf{Game}_{(j-1,d)}^{(0,j)}, \mathsf{Game}_{(j,d)}^{(0,j)}, \ldots, \mathsf{Game}_{(j,0)}^{(0,j)} = \mathsf{Game}_{(j)}^{(0,j)} \qquad (6)$$

has at most $\lceil s/2 \rceil$ sources. Moreover, there must be two consecutive hybrids that are $\delta/((2d+1)dn^2)$-distinguishable .

Continuing with this pair and guessing a well-dividing node (paying a factor $1/n$), we embed $2d$ intermediate hybrids, of which a pair must be $(\delta/((2d+1)^2 dn^3))$-distinguishable. The differing edge in this pair has now only $\lceil s/4 \rceil$ source nodes. We can now continue this recursion, e.g., guessing and then embedding between games $\mathsf{Game}_{(j_1, j_2-1)}^{(0,j_1,j_2)}$ and $\mathsf{Game}_{(j_1,j_2)}^{(0,j_1,j_2)}$ the sequence

$$\mathsf{Game}_{(j_1, j_2-1)}^{(0,j_1,j_2)}$$
$$= \mathsf{Game}_{(j_1, j_2-1,0)}^{(0,j_1,j_2)}, \ldots, \mathsf{Game}_{(j_1, j_2-1,d)}^{(0,j_1,j_2)}, \mathsf{Game}_{(j_1,j_2,d)}^{(0,j_1,j_2)}, \ldots, \mathsf{Game}_{(j_1,j_2,0)}^{(0,j_1,j_2)}$$
$$= \mathsf{Game}_{(j_1,j_2)}^{(0,j_1,j_2)} \quad .$$

But what happens if in Line 6, the middle games $\mathsf{Game}_{(j-1,d)}^{(0,j)}$ and $\mathsf{Game}_{(j,d)}^{(0,j)}$ are distinguishable? They differ in the j-th edge to the challenge, i.e., $z[j] \to z$. Since in both games all edges going into v (the first splitting node) are fake, we now need to find a node which well-divides the tree $\mathsf{T}_z \setminus \mathsf{T}_v$, as this is the tree containing all the real sources leading to z. Whereas $j \in [d]$ in the superscript denoted guessing a well-dividing node for the j-th tree above, we denote by '0' the fact that the guessed node must well-divide the tree below. We thus obtain $\mathsf{Game}_{(j-1,d)}^{(0,j,0)}$ and $\mathsf{Game}_{(j,d)}^{(0,j,0)}$ from $\mathsf{Game}_{(j-1,d)}^{(0,j)}$ and $\mathsf{Game}_{(j,d)}^{(0,j)}$, respectively, where in addition we guess a node v' well-dividing the tree $\mathsf{T}_z \setminus \mathsf{T}_v$. In any case we obtain a pair that is $(\delta/((2d+1)^3 dn^4)$-distinguishable and the differing edge has fewer than $\lceil s/8 \rceil$ sources.

Since, every time we guess and embed intermediate games, we halve the number of sources, after $\lceil \log s \rceil$ such steps we have two games that differ in one edge which has only 1 source, (meaning we have found **Game** and **Game'**). These games are thus $(\delta/d(2d+1)^{\lceil \log s \rceil} n^{\lceil \log s \rceil + 1})$-distinguishable. We now switch to the technique from Section 4, guessing the middle node of the path and embedding 2 intermediate hybrids, losing $1/3n$ in every step. After $\lceil \log \ell \rceil$ (where ℓ is the maximum length of any path), we arrive at two games from which we can construct

Game_I^P is defined as follows:

- P is a list $(a_1, \ldots, a_{m_P})$, where $a_1 = 0$ and $a_i \in [d]_0$.
- I is a list $(e_1, \ldots, e_{m_I})$ with $m_I \leq m_P$ and $e_i \in [d]_0$
- B starts with guessing nodes $v_1, v_2, \ldots, v_{m_P}$.
- For all $1 \leq i \leq m_I$ and every $e_i = I[i]$, B returns fake responses to the first e_i queries $\mathsf{encrypt}(*, v_i)$. All other queries are responded with real answers. In the end A outputs bit $b \in \{0, 1\}$.
- Define $\mathsf{T}^{(1)}, \ldots, \mathsf{T}^{(m_P)}$ from $P, v_1, \ldots, v_{m_P}$ as in Line 7.
- If v_1 is the challenge node and for $i > 1$, v_i well-divides $\mathsf{T}^{(i)}$, then return b; otherwise return 0.

Fig. 5: Definition of Game_I^P for the general-tree case.

---

an IND-CPA distinguisher. We lost another factor $1/(3n)^{\lceil \log \ell \rceil}$. In total, starting from a GSD adversary with success probability δ, we obtain an IND-CPA distinguisher with probability

$$\delta / \left(dn((2d+1)n)^{\lceil \log s \rceil} (3n)^{\lceil \log \ell \rceil} \right) \ .$$

To formalize this approach, in Figure 5 we formally define a hybrid game Game_I^P, where P determines the nodes we guess in the graph and I determines the queries (involving the guessed nodes) that are replied to with a fake answer. In contrast to Section 4, P and I are ordered *lists* rather than sets. We explain our definition in the following.

Checking whether we've guessed correctly. We start with explaining how in Game_I^P we check whether our guesses were well-dividing nodes. As an example, consider Figure 3. H_0 corresponds to $\mathsf{Game}_{(0,0)}^{(0,1)}$ (the superscripts mean that we must guess the challenge v_1 (the 0 in $(0,1)$) and a splitting node v_2 for the tree $\mathsf{T}_{v_1[1]}$ that ends in the *first* (denoted by 1 in $(0,1)$) edge going into the challenge; the subscripts mean that there are 0 fake edges going into v_1 and 0 fake nodes going into v_2. Guesses $v_1 = 7$ and $v_2 = 9$ would be correct, since A queried $\mathsf{challenge}(7)$ and $v_2 = 9$ well-divides the first subtree going into the challenge, that is T_{v_1}.

Now consider H_3 and H_4, corresponding to $\mathsf{Game}_{(0,3)}^{(0,1)}$ and $\mathsf{Game}_{(1,3)}^{(0,1)}$, respectively. The next guess v_3 should split the tree $\mathsf{T}_{v_1} \setminus \mathsf{T}_{v_2}$ (obtaining games $\mathsf{Game}_{(0,3)}^{(0,1,0)}$ and $\mathsf{Game}_{(1,3)}^{(0,1,0)}$). If we were to further recurse between these games, we thus never have to consider T_{v_2} anymore, because in all games all edges going into v_2 are fake edges.

The tree $\mathsf{T}^{(i)}$, which v_i should well-divide is thus defined for $i = 1, \ldots$ as follows: as long as there have not been any $a_i = P[i]$ with $a_i = 0$ (that is, we considered a subtree below the splitting node), $\mathsf{T}^{(i)}$ is the tree ending in the a_i-th parent of v_{i-1} (that is, $\mathsf{T}_{v_{i-1}[a_i]}$). When we have $a_i = 0$ for the first time then the lower subtree is chosen, that is $\mathsf{T}^{(i)} = \mathsf{T}^{(i-1)} \setminus \mathsf{T}_{v_{i-1}}$. In any further recursion, we need not consider $\mathsf{T}_{v_{i-1}}$ anymore, so we store it as $\mathsf{R}^{(i)}$ (for "removed"). If $a_{i+1} \geq 1$ then v_{i+1} must well-divide $\mathsf{T}_{v_i[a_{i+1}]} \setminus \mathsf{R}^{(i)}$. We need to apply the same method for the rest of the guessed nodes as well. Meaning we let $\mathsf{R}^{(i)}$ be the set of all the removed edges in the first ith steps and we first remove $\mathsf{R}^{(i)}$ from $\mathsf{T}_{v_{i-1}[a_i]}$ or $\mathsf{T}^{(i-1)}$ for $a_i \geq 1$ and $a_i = 0$ respectively, and what is left must be well-divided by the new guessed node. Formally,

$$\mathsf{R}^{(1)} = \emptyset \ , \ \mathsf{T}^{(1)} = \mathsf{T}_{v_1} \ , \quad \begin{array}{l} \text{if } a_i = 0 \text{ then } \mathsf{R}^{(i)} = \mathsf{R}^{(i-1)} \cup \mathsf{T}_{v_{i-1}}, \ \mathsf{T}^{(i)} = \mathsf{T}^{(i-1)} \setminus \mathsf{R}^{(i)}, \\ \text{if } a_i \geq 1 \text{ then } \mathsf{R}^{(i)} = \mathsf{R}^{(i-1)}, \qquad\quad \mathsf{T}^{(i)} = \mathsf{T}_{v_{i-1}[a_i]} \setminus \mathsf{R}^{(i)} \end{array} \quad (7)$$

Set I and the hybrid games. In Section 4 the set I was a subset of P and for every x in I the query $\mathsf{encrypt}(*, x)$ was responded with a fake answer. Unlike before, here we can have multiple queries of form $\mathsf{encrypt}(*, x)$, therefore I needs to determine which queries are fake for each node in P. To do so, the i-th element in I, e_i, determines that the first e_i queries $\mathsf{encrypt}(*, v_i)$ are responded with fake answers and the rest of queries are responded with real ones. If $e_i = 0$ then all such queries are replied with real answers.

Previously, for each new element z added to P, and for any game Game_I^P, we defined two new hybrid games, one with a fake response and one with a real response to the query $\mathsf{encrypt}(*, z)$. Thus for any pair of games Game_I^P and $\mathsf{Game}_{I'}^P$, with $|I \triangle I'| = 1$ we ended up with 4 games which built up 3 pairs of games satisfying property 1. Here, for a newly guessed node v_i, we can define many new games with $P_{new} = P \parallel (a_i)$ and $I_{new} = I \parallel (e_i)$, for $a_i, e_i \in [d]_0$. However not all of such games are needed. First consider two games Game_I^P and $\mathsf{Game}_{I'}^P$ with $|P| = |I| = m - 1$ such that they satisfy Property 1, meaning

- I and I' are identical except at one position i, we have $I[i] = e_i$ and $I'[i] = e'_i$, and $|I| = |I'| = m - 1$,
- $e'_i < d$ and $e_i = e'_i + 1$,
- If $i < m - 1$ then, $a_{i+1} = e_i$, $e_{i+1} = d$ and $\forall j$, $i + 1 < j < m$, $e_j = 0$, $a_j = d$.

The differentiating step in these two games is $\mathsf{encrypt}(v_i[e_i], v_i)$. Therefore the new element (to be guessed) must be in the subtree that ends with $v_i[e_i]$. That's why if the $(i+1)$-th element is already in the set P then it must be $a_{i+1} = e_i$, and to add a new element in P we must set $a_m = 0$. On the other hand, if a_i is the last element in S, the new element must be $a_m = e_i$. Thus we only need to define $2(d + 1)$ games which create $2d + 1$ pairs of games satisfying Property 1, as listed below.

$$\left(\mathsf{Game}_{I\parallel(e_m)}^{P\parallel(a_m)}, \mathsf{Game}_{I\parallel(e_m+1)}^{P\parallel(a_m)} \right) \text{ for } 0 \le e_m < d,$$

$$\left(\mathsf{Game}_{I'\parallel(e_m)}^{P\parallel(a_m)}, \mathsf{Game}_{I'\parallel(e_m+1)}^{P\parallel(a_m)} \right) \text{ for } 0 \le e_m < d,$$

$$\left(\mathsf{Game}_{I\parallel(d)}^{P\parallel(a_m)}, \mathsf{Game}_{I'\parallel(d)}^{P\parallel(a_m)} \right) \ .$$

Finally, one of the pairs of games satisfying property 1, must be $\delta/n(2d+1)$-distinguishable, if Game_I^P and $\mathsf{Game}_{I'}^P$ are δ-distinguishable.

In Lemma 3 we formalize this argument. Let us first show that well-dividing nodes always exist.

Lemma 2. *Let T be a tree of n nodes, s sources, $s \ge 2$ and let x be a sink in T with strictly more than one source reaching x. Then there exists a node x^* in T such that x^* well-divides T_x.*

Proof. [of Lemma 2] We proof Lemma 2, by giving an algorithm that finds x^* in T_x in Figure 6.

Let $S(x)$ denote the number of sources in the subtree T_x and let $x[1], \dots, x[d]$ be all nodes where there is an edge $x[i] \to x$ in T. As long as $S(x) \ge 2$, the following algorithm will output a node x^* in tree T_x such that x^* well-divides T_x.

TERMINATION. In each iteration of the while loop either x is replace with one of its parents, which can happen at most n times before reaching a source, or the algorithm terminates. Therefore there can be at most be n iterations of the while loop.

CORRECTNESS. The algorithm starts from the sink x of the tree and it traverses the tree up until it finds a well-dividing node. In each iteration, the variable m keeps track of the number of sources in the subtree $\mathsf{T} - \mathsf{T}_x$, where x is the node at hand and T is the tree it started with. In other words, at each iteration of the **while** loop, $m + S(x)$ is the total number of sources in the tree. Note that at most one of the parents of x can satisfy the predicate of the **if** statement and if such a parent exists then that node has more than half of the sources of the entire tree, so we update m and x accordingly and search for the well-dividing node in the new T_x. On the other hand, if no such a parent exists, it means that after removing all the ingoing edges of x from the tree all the subtrees ending in one of the parents of x have less sources than half of the sources of the entire tree. Also note that m is never greater than $s_\mathsf{T}/2$, since m is changed only when $S(x[i]) > s_\mathsf{T}/2$. Therefore the tree $\mathsf{T} \setminus \mathsf{T}_x$ has less than $s_\mathsf{T}/2$ sources as well, meaning x is the well-dividing node.

```
Well-Dividing Node(T, x)
    m ← 0
    s_T ← S(x)
    found ← false
    while found = false do
        if ∃i with S(x[i]) > s_T/2 then
            m ← s_T − S(x[i])
            x ← x[i]
        else
            found ← true
        end if
    end while
    output x
```

Fig. 6: Algorithm finding a well-dividing node in a tree

Lemma 3. *Let P and I be sets of integers as defined in Figure 5 and $|P| = |I| = m − 1 \geq 0$, and I' be identical to I except for i^*-th element, we have $I[i^*] = e_{i^*}$ and $I'[i^*] = e'_{i^*} = e_{i^*} − 1$. Then there are 2 cases: either*

- *$i^* = m − 1$, then let $a_m := e_{i^*}$; or*
- *$i^* < m − 1$, $a_{i^*+1} = e_{i^*}$, $e_{i^*+1} = d$ and $\forall j$, $i^* + 1 < j < m$, $a_j = 0$, $e_j = d$, then let $a_m := 0$.*

If Game_I^P and $\mathsf{Game}_{I'}^P$ are (t, ϵ)-distinguishable then

- *if $m = \lceil \log s \rceil + 1$ then, $(\mathsf{Enc}, \mathsf{Dec})$ is not $(t + Q_{\mathsf{Adv}}T_{\mathsf{Enc}} + \tilde{O}(Q_{\mathsf{Adv}}), \epsilon/(3n)^{\lceil \log \ell \rceil})$-IND-CPA-secure.*
- *Otherwise, at least one of the following pairs of games is $(t, \epsilon/n(2d + 1))$-distinguishable*

$$\mathsf{Game}_{I\|(e_m)}^{P\|(a_m)} \text{ and } \mathsf{Game}_{I\|(e_m+1)}^{P\|(a_m)} \quad \text{for } 0 \leq e_m < d$$
$$\mathsf{Game}_{I\|(d)}^{P\|(a_m)} \text{ and } \mathsf{Game}_{I'\|(d)}^{P\|(a_m)}$$
$$\mathsf{Game}_{I'\|(e_m)}^{P\|(a_m)} \text{ and } \mathsf{Game}_{I'\|(e_m+1)}^{P\|(a_m)} \quad \text{for } 0 \leq e_m < d$$

Proof. [of Lemma 3] According to P, I and I', the two games only differ in the response to e_{i^*}-th query of the form $\mathsf{encrypt}(v_{i^*}[e_{i^*}], v_{i^*})$. If $m < \log s$, we add another element to P. The m-th guessed node must well-divide subtree $\mathsf{T}^{(m)} := \mathsf{T}_{v_{i^*}[e_{i^*}]} - \sum_{i \in \{2,3,\cdots,m-1\}:P[i]=0} \mathsf{T}_{v_{i-1}}$ in both games. If such a node exists, the probability of guessing the node correctly is $1/n$, thus $\mathsf{Game}_I^{P\|(a_m)}$ and $\mathsf{Game}_{I'}^{P\|(a_m)}$ are $(t, \epsilon/n)$-distinguishable. Moreover, by the triangle inequality we have

$$\sum_{e_m=0}^{d-1} \Delta_t \left(\mathsf{Game}_{I\|(e_m)}^{P\|(a_m)}, \mathsf{Game}_{I\|(e_m+1)}^{P\|(a_m)} \right) + \Delta_t \left(\mathsf{Game}_{I\|(d)}^{P\|(a_m)}, \mathsf{Game}_{I'\|(d)}^{P\|(a_m)} \right)$$

$$+ \sum_{e_m=0}^{d-1} \Delta_t \left(\mathsf{Game}_{I'\|(e_m)}^{P\|(a_m)}, \mathsf{Game}_{I'\|(e_m+1)}^{P\|(a_m)} \right) \geq \Delta_t \left(\mathsf{Game}_I^{P\|(a_m)}, \mathsf{Game}_{I'}^{P\|(a_m)} \right) > \frac{\epsilon}{n} .$$

Consequently at least one of the $2d+1$ pairs of games, on the left hand side of the inequality above, must be $(t, \epsilon/n(2d + 1))$-distinguishable. On the other hand if no such a node exists we simply ignore the rest of the guessed nodes in the game. Then $\mathsf{Game}_{I\|(d)}^{P\|(a_m)}$ and $\mathsf{Game}_{I'\|(d)}^{P\|(a_m)}$ are as distinguishable as Game_I^P and $\mathsf{Game}_{I'}^P$. Therefore the claim of the lemma holds in this case as well.

Since every guessing step reduces the number of sources in the subtree reaching the challenge into at most half of what it was before, after $m - 1 = \lceil \log s \rceil$ steps, we know for certain that there is at most one source left in the subtree reaching the challenge, in which case we can apply Lemma 1 at most $\lceil \log \ell \rceil$ times to conclude the proof.

Guessing the challenge node and defining games $\mathsf{Game}^{(0)}_{(0)}, \dots, \mathsf{Game}^{(0)}_{(d)}$ loses a factor $n \cdot d$. Applying Lemma 3 repeatedly for $\lceil \log s \rceil$ times, we get the proof of Theorem 2.

6   Conclusions and Open Problems

We showed a quasipolynomial reduction of the GSD game on trees to the security of the underlying symmetric encryption scheme. As already discussed in the introduction, it is an interesting open problem to extend our reduction to general (directed, acyclic) graphs or to understand why this is not possible. This is the second result using the "nested hybrids" technique (after its introduction in [FKPR14] to prove the security of constrained PRFs), and given that it found applications for two seemingly unrelated problems, we believe that there will be further applications in the future.

One candidate is the problem of proving security under selective opening attacks [DNRS99] [FHKW10,BHY09], where one wants to prove security when correlated messages are encrypted under different keys. Here, the adversary may adaptively chose to corrupt some keys after seeing all ciphertexts, and one requires that the messages in the unopened ciphertexts are indistinguishable from random messages (sampled so they are consistent with the already opened ciphertexts). This problem is notoriously hard, and no reduction avoiding complexity leveraging to IND-CPA security of the underlying scheme is known.

Acknowledgements

References

[ABBC10] Tolga Acar, Mira Belenkiy, Mihir Bellare, and David Cash. Cryptographic agility and its relation to circular encryption. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 403–422. Springer, May 2010.

[BGI14] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In Hugo Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 501–519. Springer, March 2014.

[BH93] Donald Beaver and Stuart Haber. Cryptographic protocols provably secure against dynamic adversaries. In Rainer A. Rueppel, editor, *EUROCRYPT'92*, volume 658 of *LNCS*, pages 307–323. Springer, May 1993.

[BHY09] Mihir Bellare, Dennis Hofheinz, and Scott Yilek. Possibility and impossibility results for encryption and commitment secure under selective opening. In Antoine Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 1–35. Springer, April 2009.

[BRS03] John Black, Phillip Rogaway, and Thomas Shrimpton. Encryption-scheme security in the presence of key-dependent messages. In Kaisa Nyberg and Howard M. Heys, editors, *SAC 2002*, volume 2595 of *LNCS*, pages 62–75. Springer, August 2003.

[BW13] Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 280–300. Springer, December 2013.

[CFGN96] Ran Canetti, Uriel Feige, Oded Goldreich, and Moni Naor. Adaptively secure multi-party computation. In *28th ACM STOC*, pages 639–648. ACM Press, May 1996.

[CL01] Jan Camenisch and Anna Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 93–118. Springer, May 2001.

[DNRS99]   Cynthia Dwork, Moni Naor, Omer Reingold, and Larry J. Stockmeyer. Magic functions. In *40th FOCS*, pages 523–534. IEEE Computer Society Press, October 1999.

[FHKW10] Serge Fehr, Dennis Hofheinz, Eike Kiltz, and Hoeteck Wee. Encryption schemes secure against chosen-ciphertext selective opening attacks. In Henri Gilbert, editor, *EURO-CRYPT 2010*, volume 6110 of *LNCS*, pages 381–402. Springer, May 2010.

[FKPR14] Georg Fuchsbauer, Momchil Konstantinov, Krzysztof Pietrzak, and Vanishree Rao. Adaptive security of constrained PRFs. In Palash Sarkar and Tetsu Iwata, editors, *ASI-ACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 82–101. Springer, December 2014.

[GGM86]   Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986.

[KPTZ13] Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 13*, pages 669–684. ACM Press, November 2013.

[MP06]   Daniele Micciancio and Saurabh Panjwani. Corrupting one vs. corrupting many: The case of broadcast and multicast encryption. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *ICALP 2006, Part II*, volume 4052 of *LNCS*, pages 70–82. Springer, July 2006.

[Pan07]   Saurabh Panjwani. Tackling adaptive corruptions in multicast encryption protocols. In Salil P. Vadhan, editor, *TCC 2007*, volume 4392 of *LNCS*, pages 21–40. Springer, February 2007.

[SS00]   Ali Aydin Selçuk and Deepinder P. Sidhu. Probabilistic methods in multicast key management. In *Information Security, Third International Workshop, ISW 2000, Wollongong, NSW, Australia, December 20-21, 2000, Proceedings*, pages 179–193, 2000.

[WGL00]   Chung Kei Wong, Mohamed Gouda, and Simon S. Lam. Secure group communications using key graphs. *IEEE/ACM Transactions on Networking*, 8(1):16–30, February 2000.